Università degli Studi di Padova

Dipartimento di Ingegneria dell'Informazione

Probabilistic Active Sensing in Multi-Sensor Applications

Ph.D. candidate
Luca Varotto

Advisor
Prof. Angelo Cenedese

Director & Coordinator
Prof. Andrea Neviani

Ph.D. School in
Information Engineering

Department of
Information Engineering
University of Padova
2021

University of Padova
Department of Information Engineering

Università
degli Studi
di Padova

Ph.D. course in: Information Engineering
Curriculum: Information Science and Technology

Probabilistic Active Sensing in Multi-Sensor Applications

Director: Prof. Andrea Neviani
Advisor: Prof. Angelo Cenedese

Ph.D. candidate: Luca Varotto
Cycle: 34$^{\text{th}}$ (2018-2021)

Year: 2021

"Questions you cannot answer are usually far better for you than answers you cannot question."

Yuval Noah Harari

Abstract

In the recent years, the advances in high performance embedded microcontrollers, optimized signal processing techniques and reliable communication protocols have fostered the development of pervasive sensing technologies. By acquiring spatially distributed and information-rich data, pervasive sensing has become fundamental in several applications. In industrial settings it is used to predict failures, to reduce slowdowns, and to improve workers safety and energy efficiency; in urban scenarios pervasive sensing spots potential security threats, while in critical environments it enables continuous monitoring at previously inaccessible locations. In robotics, the aggregation of information from multiple spatially distributed sources allows the robot to learn faster, and to plan more robust and more adaptive strategies under dynamic and uncertain conditions. Recently, the research in machine learning and automatic controls has contributed to a new generation of smart pervasive sensors, where connected devices become autonomous. Smart sensors can decide how to gather data within a set of resource constraints; thus they represent a new frontier for data acquisition systems in a wide range of applications including agriculture, surveillance, post-disaster assessment, and manufacturing.

The main research interest of this Thesis concerns Multi-Sensor Probabilistic Active Sensing. Active sensing schemes are responsible for the generation of control and optimization strategies in resource-constrained perception and data collection processes; hence, they are of paramount importance in the decision making process of smart devices. At the same time, probabilistic approaches accommodate for noisy sensors and unknown environments, while multi-sensor solutions introduce more accuracy, reactivity and robustness in intelligent operations like planning, learning, and reasoning. Furthermore, multi-sensor architectures allow to capture environmental properties that single-sensor counterparts usually miss.

The main contribution of this Thesis is the development of a mathematical framework for multi-sensor probabilistic active sensing; the proposed scheme is modular and applicable for general active sensing tasks. Indeed, the formulation is derived in the context of active position estimation, but it is then validated under four distinct applications: active localization, active tracking, energy-constrained video surveillance, and connected vehicles.

Extensive simulation and experimental studies have been conducted to evaluate the proposed approaches, examining the performance against original baselines. The main findings suggest that active sensing, coupled with Bayesian inference and information fusion, enables self-supervised learning techniques, allows to account for realistic perception uncertainties and unmodeled dynamics, and introduces faster, more accurate, and more robust estimation processes.

Sommario

Negli ultimi anni, i numerosi progressi nel campo dei sistemi embedded ad alte prestazioni, l'ottimizzazione delle tecniche di elaborazione dei segnali, e la maggiore affidabilità nei sistemi di comunicazione, hanno portato allo sviluppo di tecnologie di sensoristica pervasiva. Tramite l'acquisizione di dati spazialmente distribuiti e ricchi di informazioni, la sensoristica pervasiva è diventata fondamentale in diverse applicazioni. In ambito industriale viene utilizzata per prevedere guasti, ridurre i rallentamenti e migliorare la sicurezza dei lavoratori e l'efficienza energetica; negli scenari urbani la sensoristica pervasiva individua potenziali minacce alla sicurezza, mentre in ambienti critici consente il monitoraggio continuo in luoghi precedentemente inaccessibili. Nella robotica l'aggregazione di informazioni provenienti da più fonti spazialmente distribuite, consente ai robot di apprendere più velocemente e di pianificare strategie di movimento più robuste e più adattive, anche all'interno di ambienti variabili e poco noti. Recentemente, la ricerca nel settore del machine learning e dei controlli automatici ha favorito lo sviluppo di una nuova generazione di sistemi di sensoristica intelligenti, in cui i sensori, oltre ad essere connessi tra di loro, diventano autonomi. Essi possono infatti prendere decisioni sul processo di raccolta dati, nel rispetto delle risorse a disposizione. Per questo motivo, le tecnologie di sensoristica intelligente rappresentano oggi una nuova frontiera per i sistemi di acquisizione dati in numerose applicazioni tra cui l'agricoltura, la videosorveglianza, il monitoraggio in ambienti critici, e nel campo manifatturiero.

Questa Tesi è dedicata alle tecniche di sensoristica attiva probabilistica e multi-sensoriale. La percezione attiva si occupa di generare strategie di controllo e di ottimizzazione nei confronti del processo di raccolta dati e sulla base delle risorse a disposizione; per questo motivo, essa è fondamentale nel processo decisionale dei sistemi di sensoristica intelligente. Allo stesso tempo, gli approcci probabilistici possono operare in presenza di sensori rumorosi e ambienti sconosciuti, mentre le soluzioni multi-sensore introducono maggiore precisione, reattività e robustezza in fase di pianificazione ed apprendimento. Inoltre, le architetture multi-sensore consentono di catturare proprietà ambientali che le controparti a sensore singolo non riescono a cogliere.

Il contributo principale di questa Tesi risiede nello sviluppo di un framework matematico per la percezione attiva probabilistica e multi-sensore; lo schema proposto è modulare e applicabile in diversi ambiti legati alla sensoristica attiva. La formulazione viene dapprima proposta nel contesto della stima di posizione attiva, e viene poi convalidata in quattro applicazioni distinte: localizzazione attiva, tracciamento attivo, videosorveglianza con vincoli di energia, e veicoli connessi.

Al fine di valutare gli approcci proposti, sono stati condotti ampi studi sia numerici che sperimentali; in particolare, le performance sono state confrontate con delle baseline di riferimento. I risultati principali mostrano che la percezione attiva, quando viene integrata con sistemi di inferenza Bayesiana e con tecniche di sensor fusion, consente di progettare metodi di apprendimento auto-supervisionato, di modellare le incertezze tipiche dei sistemi di percezione reali, e di introdurre maggiore rapidità, accuratezza e robustezza nei processi di stima.

Acknowledgments

During my life, and during this three-year course, I had the opportunity to live and interact with people that influenced the person that I am; they gave me the support and the tools to improve and become a better person. It is therefore compulsory to thank all of them.

I would first like to thank my supervisor, Professor Angelo Cenedese, whose expertise was invaluable in formulating the research questions and methodology. Thank you for all the support, patience and encouragement you gave me.

I would also like to deeply thank Professor Andrea Cavallaro for having supervised my stay at the Queen Mary University of London and for having been a valuable mentor. Your insightful feedback pushed me to sharpen my thinking and brought my work to a higher level.

Thanks to all the members of the Centre for Intelligent Sensing (QMUL), in particular Alessio for the fruitful discussions about computer vision and London life.

Thanks to my London mate Raphael. I will always be grateful for your politeness and hospitality. Cheers!

Thanks to all the professors, researchers and students of the University of Padova that I have met during this eight-year journey. Thank you Iulian, the time spent with you changed my perspective on life.

I must mention my colleagues of the SPARCS research group: Giulia, Marco, Riccardo, Nicola, Beniamino, Pietro, and Federico. It has been an honor to share ideas with brilliant and passionate researchers like you. Thanks to Marco and Nicola to be good friends before than co-workers.

My gratitude extends to all the other people who I worked with in these years; in particular, I would like to thank Digital Lighting Srl for all the opportunities they offered me, and for feeding my research activity with real-world problems.

I am eternally grateful to my family, especially my parents and my sister. Dear mom and dad: there are no words to say thank you enough for everything you have done for me; first you gave me birth, and then you gave me life. Thank you Elena, because I know I can always count on you.

A special thanks goes to the friends ever. Thank you Andrea: you taught me how to smile under the rain. The world would be a better place if there were more people like you.

Thanks to Yianny, my sweet home. You're really special to me.

Acronyms & Abbreviations

2D	two-dimensional
3D	three-dimensional
ADAS	Advanced Driving Assistance Systems
AI	Artificial Intelligence
AL	Active Localization
Alg.	Algorithm
AoA	Angle of Arrival
APE	Active Position Estimation
AS	Active Sensing
ASe	Active Search
AT	Active Tracking
AUV	Autonomous Underwater Vehicle
BLE	Bluetooth Low Energy
BO	Bayesian Optimization
Chapt.	Chapter
Def.	Definition
DRL	Deep Reinforcement Learning
ECDF	Empirical Cumulative Distribution Function
e.g.	exempli gratia
Fig.	Figure
FIM	Fisher Information Matrix
FoV	Field of View
GP	Gaussian Process
GPR	Gaussian Process Regression
GPS	Global Positioning System
HAZMAT	Hazardous Material
i.e.	id est
IEEE	Institute of Electrical and Electronics Engineers
i.i.d.	independent and identically distributed
IoT	Internet of Things
KLD	Kullback-Liebler divergence
MARL	Multi-Agent Reinforcement Learning
MAS	Multi-Agent System
MC	Monte Carlo
MCTS	Monte Carlo Tree Search

MDP	Markov decison process
MI	Mutual Information
MM-PAS	Multi-Modal Probabilistic Active Sensing
MPC	Model Predictive Control
MS-MP	Multi-Sensor Multi-Platform
MS-SP	Multi-Sensor Single-Platform
MTS	Minimum Time Search
PAS	Probabilistic Active Sensing
Ph.D./PhD	Philosophiæ Doctor
PLM	path-loss model
POD	Probability of Detection
PTS	Probabilistic Target Search
PTZ	Pan-Tilt-Zoom
QoS	Quality of Sight
RBE	Recursive Bayesian Estimation
RF	Radio Frequency
RMSE	Root Mean Squared Error
RPA	Robotic Process Automation
RSSI	Received Signal Strength Indicator
Rx	receiver
SAR	Search and Rescue
SAT	Search and Tracking
Sec.	Section
SNR	Signal-to-Noise Ratio
SPARCS	Space and Aerial Control Systems
STE	Source Term Estimation
TDOA	Time Difference of Arrival
Tx	Transmitter
UAV	Unmanned Aerial Vehicle
UGV	Unmanned Ground Vehicle
VSN	Visual Sensor Network
w.l.o.g.	without loss of generality
w.r.t.	with respect to
WSN	Wireless Sensor Network

Notation

All the statements belonging to the latter have to be intended of global usage for the whole document. In the following lines, the most used symbols together with further basic notation and operators are reported as a preliminary.

Symbols In this manuscript the following notation is adopted. The symbols $\mathbb{N}, \mathbb{Z}, \mathbb{R}, \mathbb{C}$ respectively denote the set of natural, integer, real and complex numbers. The symbols $\mathbb{R}^n$ and $\mathbb{R}^{n\times m}$ indicate the spaces of n-dimensional vectors and $(n \times m)$-dimensional matrices having real entries, respectively. The notation $\{0,1\}^N$ is used to denote the space of N-dimensional boolean vectors. Lowercase italic letters and greek symbols, e.g., x, v, ϵ, denote scalar values. Lowercase roman bold letters and bold greek symbols, e.g., $\mathbf{x}, \mathbf{v}, \boldsymbol{\epsilon}$, denote vectors assumed by convention to be column vectors. Uppercase roman bold letters, e.g., $\mathbf{X}, \mathbf{A}, \mathbf{C}$ denote matrices. Uppercase calligraphic letters, e.g., $\mathcal{X}, \mathcal{A}, \mathcal{C}$ denote sets. The symbol $\mathbf{I}_n$ denotes the $(n \times n)$-dimensional identity matrix, while $\mathbf{0}_n$ is the zero vector of dimension n. The symbol $\varnothing$ indicates the empty set.

Regarding the statistical distributions used in this Thesis, $\chi^2(n)$ denotes the chi-squared distribution with n degrees of freedom; $\mathcal{N}(x|\mu, \sigma^2)$ is the Gaussian distribution over the random variable x with expectation μ and variance σ^2 (when it is not necessary to explicit the random variable, the notation $\mathcal{N}(\mu, \sigma^2)$ is used); a Bernoulli distribution with parameter $p \geq 0$ is denoted as $\mathcal{B}(p)$; a Binomial distribution with parameters n and p is denoted as $\mathcal{B}in(n, p)$; uniform distributions over the interval $[a, b]$ are denoted as $\mathcal{U}(a, b)$.

The shorthand notation $z_{t_0:t_1}$ indicates the sequence of variables from time instant t_0 to t_1, namely the set $\{z_k\}_{k=t_0}^{t_1}$ (i.e., the subscripts denote a temporal evolution of the same variable). Otherwise, the shorthand notation $z_{1:N}$ indicates the sequence of variables from index 1 to N, namely the set $\{z_i\}_{i=1}^{N}$ (i.e., the subscripts denote a collection of N different variables); if these N variables can evolve in time, the notation $z_{t,i}$ identifies the i-th variable at time t.

Operators Given a scalar x, $|x|$ denotes its absolute value. Analogously, given a n-dimensional vector $\mathbf{v}$, $|\mathbf{v}|$ represents its absolute value, while $\|v\|_p$ is its p-norm. The (row) vector $\mathbf{v}^\top$ indicates the transpose of $\mathbf{v}$ and the scalar value $v(i)$ specifies its i-th entry In addition, $\text{diag}(\mathbf{v})$ denotes the $(n \times n)$ diagonal matrix whose main diagonal consists of the elements of $\mathbf{v}$.

Given a matrix $\mathbf{A}$, $\|\mathbf{A}\|_p$ and $\mathbf{A}^\top$ denote the same operations of norm and transpose as before. The symbol $a(i, j)$ indicates the ij-th scalar entry of the matrix.

Given a set $\mathcal{C}$, the symbol $|\mathcal{C}|$ denotes its cardinality, i.e., the number of elements belonging to the set. Given two sets $\mathcal{A}$ and $\mathcal{C}$, the relation $\mathcal{C} \subset \mathcal{A}$ ($\mathcal{C} \subseteq \mathcal{A}$) indicates that $\mathcal{C}$ is a proper (non necessarily proper) subset of $\mathcal{A}$. The symbols $\cap$, $\cup$ and $\setminus$ represent intersection, union and complement of sets, respectively.

Contents

1

Towards pervasive and smart sensing

The significant factors impacting the growth of the global smart sensor industry include a surge in demand from the automotive and domotics sectors, a growing trend toward the Internet of Things, and the development of smart city scenarios. In this context, this Chapter provides an overview of the technological trends that support the main topics of this Thesis: smart sensing, information fusion and intelligent control.

Contents

1.1 Autonomous systems and ubiquitous computing

Autonomous systems are composed by one or more agents capable to act without external intervention, possibly under uncertainties and against external disturbances; in this regard, each agent autonomously senses and interacts with the environment to accomplish a set of goals (predefined or computed online, collective or personal) Chen, Abbod, and Shieh (2019a). According to literature Chen et al. (2019a), the fundamental properties of autonomous systems are self-regulation and self-diagnosis (to operate continuously and fulfil repetitive tasks), self-adaptation (to act under dynamic environmental conditions), and self-organization (to guarantee optimal configuration in relation to the goals to be achieved). Examples of autonomous systems are unmanned robots Meera, Popović, Millane, and Siegwart (2019) and smart sensor systems Varotto and Cenedese (2021b). The past several decades have witnessed incredible breakthroughs in the field of control engineering and autonomous systems, enabling high levels of autonomy and integration with modern computing architectures Chen et al. (2019a). More specifically, the maturation of robotic, machine learning and automation disciplines, coupled with the rapid advances in ubiquitous computing technologies, have enabled autonomous systems to increasingly penetrate across multiple industrial and civil domains Chen et al. (2019a); Popovic (2019). At the same time, with their rising accessibility, robots and smart sensors are growing more commonplace in our workplaces, in our private homes, and in our society as a whole. The main motivation comes from the fact that, as autonomous intelligent machines become more accurate, they can introduce remarkable time, labor, and cost savings across all sectors of the economy. Hence, they are promising to revolutionize the ways in which we work and live. In fact, the potential applications are extensive and they expand at an extraordinary rate, especially in tasks that involve repetitive procedures, demand high levels of accuracy or precision, and transpire in risky and hazardous environments.

Recently, new research topics related to autonomous systems problems have started to gain popularity among the robotics, automation and control communities. Some examples are given by ubiquitous robotics, pervasive computing, intelligent and distributed edge computing systems, smart energy systems, security control, and privacy of cyber-physical systems. All these scientific fields are fostered by the advances in communication technology (mobile communication, secure broadband transmission), microelectronics (new materials, smaller and more powerful devices), energy supply, and sensor systems; an important role is also played by the rise of the Internet of Things (IoT) paradigm, which allows the interconnection of different networked embedded devices used in the everyday life and integrated into the Internet Li, Da Xu, and Zhao (2015). An additional contribute has been given by the increased performance of production processes, better methods to create software, more efficient programming languages and operating systems, together with innovative concepts for man-machine interaction Friedewald and Raabe (2011). Among all the aforementioned research lines, this Thesis will cover topics related to ubiquitous robotics and pervasive sensing.

Ubiquitous robotics (also referred to as ubirobotics or networked robotics), cloud robotics and multi-agent systems are crosscutting technologies that utilize the whole range of mod-

ern information and communication technologies (ICT) with representative application in systems engineering domains Chibani, Amirat, Mohammed, Matson, Hagita, and Barreto (2013); Fabris (2021); Herrera, Pérez-Hernández, Kumar Parlikad, and Izquierdo (2020); Miki, Popović, Gawel, Hitz, and Siegwart (2018). In the ubirobotics framework, multiple autonomous agents communicate, interact and cooperate to accomplish a common goal or to satisfy a set of concurrent tasks Krasniqi and Hajrizi (2016); Masehian, Jannati, and Hekmatfar (2017). Ubirobots overcome the limitations of stand-alone robots by integrating them with web services and ambient intelligence technologies Chibani et al. (2013); moreover, they acquire richer functionalities and open the way for the composition of a variety of services, leveraging on semantic perception, reasoning and actuation Chibani et al. (2013). Ubirobotics plays an important role also within the emergent framework of robotic process automation (RPA) Hofmann, Samp, and Urbach (2020), where collaborative software agents (i.e., bots) run on intelligent digital services Avula, Chadwick, Arguello, and Capra (2018). In line with the IoT and ambient intelligence concepts, these agents have the ability to coordinate autonomously their activities with other physical or logical entities, in order to provide assistive and monitoring services. Moreover, the emergence of cloud computing allows to enrich the agents' cognitive capabilities and share their knowledge by connecting to web infrastructures Ferrarotti, Breschi, and Bemporad (2021).

Most ubiquitous robotics concepts evolve from the application in robotics scenarios of the ubiquitous computing paradigm (also known as *pervasive computing*). Pervasive computing promotes the idea of a physical world that is richly interwoven with sensors, actuators, displays, and computational elements, embedded seamlessly in everyday objects, and connected through a communication network Weiser, Gold, and Brown (1999). In other words, pervasive computing is a general concept referring to the ubiquity of information technology and computing power, and to the integration of internet connected computers, microprocessors and sensors into the environment Friedewald and Raabe (2011). More specifically, ubiquitous computing systems are characterized by the following features Friedewald and Raabe (2011); Kim, Kim, and Lee (2004):

- *modularity, decentralization and networking*: all devices are capable to share information, but their computational and storage capabilities make them self-reliant, so that they do not require any central control unit;
- *integration* of ambient intelligence sensors and actuators in objects of daily use;
- *ubiquity*: mobile support for the user through information services anywhere and any time;
- *accessibility*: computers should be accessible at anytime and at anyplace;
- *adaptation*: context awareness and adaptation of the system with respect to changing requirements;
- *self-organization*: automatic recognition and autonomous processing of repetitive tasks without user intervention.

The principal goal of pervasive computing is that of providing support to everyday tasks in an automated fashion, from people home assistance to the continuous optimization

of economic, industrial and social processes. In the long term, ubiquitous computing can pervade all spheres of life: it promises to increase comfort and energy efficiency in the private home area; it may increase roads safety, if applied to smart transportation systems; it could raise work productivity in the offices through RPA and adaptive personal assistance systems; finally, in the medical field implantable sensors and micro-computers may be used to monitor the health of the user.

The principal research trends in ubirobotics and ubiquitous computing concern context and activity recognition Jiang, Cai, Ma, Yang, and Liu (2018), semantic reasoning Sa, Popović, Khanna, Chen, Lottes, Liebisch, Nieto, Stachniss, Walter, and Siegwart (2018), multi-agent coordination Fabris (2021), and privacy-by-design Osia, Shahin Shamsabadi, Sajadmanesh, Taheri, Katevas, Rabiee, Lane, and Haddadi (2020). From a more industrial and technological perspective, the main challenges regard the design of software and middleware tools capable to creare ubirobot services as plug-and-play applications, to guarantee interoperability among robots and smart devices, and to extend ubirobots intelligence and knowledge from the cloud.

1.2 Pervasive sensing

When ubiquitous computing meets sensing devices (e.g., RFIDs Deyle, Nguyen, Reynolds, and Kemp (2010), audio microphones Li, Girin, Horaud, and Gannot (2016), cameras SanMiguel and Cavallaro (2014)), it is referred to as *pervasive sensing* Bordel, Alcarria, Robles, and Martín (2017); Park, Chao, Hussain, and Yen (2014). Facilitated by the increasing availability of inexpensive, wireless sensors and by the cost-effective and straight-forward installation, pervasive sensing has recently become a fundamental technology in several applications. In industrial settings it provides a variety of measurable and significant improvements, from process and worker safety, to regulatory compliance, from equipment reliability, to energy efficiency. Pervasive sensing is also used to monitor and track processes performance so that productivity levels can be optimized[1]; in addition, the large amount of data collected are used to predict failures and reduce slowdowns and shutdowns. In urban scenarios, instead, pervasive sensing spots potential security threats, while in critical environments (e.g., oil refineries) it enables continuous, on-line monitor at previously inaccessible locations and at an affordable cost; this makes it possible to detect and respond to hazards early. In the robotics community, the emergence of pervasive sensing is strongly contributing to extend ubirobots' perception capabilities. As a matter of fact, spatially distributed sensing platforms can acquire large amount of data from the physical world; hence, a robot surrounded by internet-connected sensing devices has access to a larger and richer amount of information compared to the localized data collected by its on-board sensors Bordel et al. (2017). Consequently, by aggregating information from multiple sources, the robot attains faster, more robust and more adaptive planning under dynamic and uncertain environments. Ubiquitous computing not only improves the decision making process of the robot, but it is also useful to increase its cooperative capabilities. In fact, the storage and communication

[1] https://www.emerson.com/documents/automation/-advantages-of-pervasive-sensing-en-us-94136.pdf

capabilities embedded in modern pervasive sensing platforms allow to use these devices as communication bridge between two or more robots involved in the same task.

Despite the increasing interest of several industrial sectors towards pervasive sensing applications, there are still many open questions to be addressed in research. These are mostly related to data management criteria (e.g., how to deal with data accumulation and limited storage capabilities), data stream processing (e.g., how to deal with continuous incoming data flows), distributed processing techniques and computational intelligence (e.g., speech recognition, image analysis, text mining, information retrieval, gesture recognition), as well as artificial intelligence applied to embedded devices Park et al. (2014).

1.3 Smart sensing

Recently, pervasive sensors have become smart Zhang, Gu, Vlatkovic, and Wang (2004) thanks to the advances in digital and embedded technologies, and from the research in the fields of machine learning, automatic controls, sensor fusion and signal processing. *Smart sensors* are static Varotto and Cenedese (2021c) or mobile Varotto, Cenedese, and Cavallaro (2021) devices endowed with on-board decision-making capabilities; hence, a smart sensor is required to actively control the data acquisition process, according to the incoming data, the sensing mission specifications, and its current state (e.g., energy available, position in space, internal configuration parameters). To accomplish this task, the main functional properties of any smart sensor are self-organization and self-reconfiguration, self-awareness and self-diagnostics (i.e., knowledge of the current state of the system), self-calibration and adaptation Rinner, Esterle, Simonjan, Nebehay, Pflugfelder, Domínguez, and Lewis (2015); SanMiguel, Micheloni, Shoop, Foresti, and Cavallaro (2014). For this reason, they can be treated as autonomous intelligent agents (see Sec. 1.1), whose functionalities are supported by processing, communication, and sensing modules.

Sensing module The core unit of any measurement system is the sensing module, which is characterized by several device-dependent parameters: sensing range, Field of View (FoV, or sensing domain), accuracy, information content, and sampling time. In the following, a qualitative introduction of these sensor properties is provided; for a more formal discussion, the reader is addressed to Chapter 3.

- Sensing range: maximum operational distance at which data can be collected.
- Field of View: the portion of environment from which data can be collected.
- Accuracy: quality of the information collected; it is often non-homogeneous over the FoV.
- Information content: type of information collected. Sensors can be roughly classified into five main classes Michieletto (2018): (i) environmental sensors, which retrieve information about a specific quantity of the environment (e.g., temperature, humidity, pressure, light); (ii) proximity sensors, which reveal the presence of objects of interest (e.g., people, obstacles); (iii) visual sensors, which capture frames of a

certain scene in a continuous or sampled manner; (iv) range sensors that provide distance measurements; and (v) bearing sensors that sense the relative direction.

- Sampling time: data collection rate.

Communication module Smart sensors are capable to propagate information to other agents[2]. As mentioned in Sec. 1.2, the communication capabilities of pervasive sensors can be used to establish cooperation among different sensing units and with surrounding connected robots. Therefore, smart sensors are fundamental building blocks in sensor networks and multi-agent systems (see Sec. 1.4).

In real-world scenarios, any communication-reliant application (e.g., multi-agent cooperation) is affected by several limitations. At first, communication may be slow, unavailable and intermittent Rondón, Mahmood, Grimaldi, and Gidlund (2019). Secondly, any communication device is generally constrained to a certain range; hence each agent can only interact with a limited set of neighbours. To address the aforementioned issues, it is important to regulate the transmission/reception of the messages through protocols and algorithms that tolerate and mitigate communication faults and limited communication coverage. In this regard, coordination among agents can follow a centralized, decentralized or distributed scheme. With a centralized coordination, the agents exchange information with a central node or ground station. As discussed in Chapter 3, the central node is a single point of failure and might not receive complete and updated information due to sensing and communication limitations. With decentralized coordination there are multiple leader agents that act as central nodes for smaller groups of agents. Each leader then coordinates with other leaders. Finally, agents with a sufficient amount of memory and processing power can coordinate in a distributed manner. In distributed coordination each robot decides independently, even with limited available information. Distributed algorithms enable individual robots to operate with partially available information and are therefore only marginally affected by communication problems.

Computational module In smart environments scenarios, the trend is to avoid centralized architectures, where a single high powerful control unit (e.g., a server) elaborates all the data received from the sensors in order to produce the required results. It is instead preferred the adoption of a decentralized (or distributed) paradigm, according to which the storage and processing resources of the agents are exploited to elaborate information locally. In this case, only distilled, high-level, information is delivered to the central unit, or to the end users. For this reason, smart sensors are generally equipped with a microprocessor with a suitable memory buffer and with the computational resources necessary to provide a local contribution to the distributed task solution. Some implementations exploit multi-core architectures which are often equipped with dedicated cores for signal processing and system management. In this work, a sensor endowed with a computational module is considered to be *smart*, since it is capable to implement intelligent control[3], decision-making and planning strategies.

[2]In this work the term *agent* is used to identify any autonomous entity, either a human being, a robot, or a smart sensor.

[3]The term *intelligent control* is used when any kind of artificial intelligence (neural networks, Bayesian inference, machine learning, evolutionary algorithms, reinforcement learning) is present in the control loop Blondin, Sáez, and Pardalos (2019); Jafari and Xu (2018); Rodríguez-Molina, Mezura-Montes, Villarreal-Cervantes, and Aldape-Pérez (2020).

Memory bandwidth and capacity are further crucial resources for smart sensors, since signal processing tasks typically deal with huge amounts of data.

1.3.1 Applications

Smart sensing is an emerging technology that is becoming a turning point in several industrial settings, especially when large amount of information is used to control the manufacturing quality, reduce unnecessary maintenance, and improve system reliability. In most industrial contexts, system reliability is as important as the maintenance costs (e.g., in large scale smart grids Li et al. (2015)); hence, sensors connectivity, self-diagnosis and self-configuration are critical requirements.

Smart sensors automatize and optimize data acquisition campaigns, where traditional systems may require excessive human labor, in terms of supervision and maintenance. Thanks to the on-board decision-making capabilities, smart sensors represent the intelligent units of smart cities, smart buildings, smart grids, and smart factories Jaradat, Jarrah, Bousselham, Jararweh, and Al-Ayyoub (2015). Other relevant use-cases include agriculture Sa et al. (2018), urban surveillance Bernardini, Fox, and Long (2017), post-disaster assessment Queralta, Taipalmaa, Pullinen, Sarker, Gia, Tenhunen, Gabbouj, Raitoharju, and Westerlund (2020), and environmental monitoring Popovic (2019). In these scenarios, and many more, intelligent systems enable gathering valuable information at unprecedented levels of resolution and autonomy.

1.3.2 Use case: smart cameras

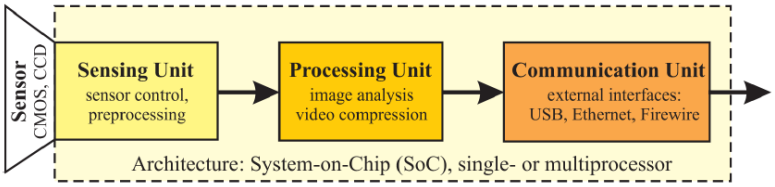

Figure 1.1. A generic architecture of a smart camera Rinner and Wolf (2008).

Many sensors collect and process scalar values, such as temperature or light measurements. The main benefit of scalar quantities is that they require limited memory and processing resources; nevertheless, the amount of information carried by scalar sensors is often inadequate for many applications. On the other hand, cameras collect richer information, at the cost of higher processing demand both for image enhancement and for high-level features extraction (e.g., object bounding boxes in recognition tasks).

Recently, the convergence of advances in computer vision, embedded computing, and communication technologies has brought to the evolution of smart cameras Rinner and Wolf (2008). Smart cameras represent a prominent example of pervasive sensing devices, by integrating visual sensing with embedded processing and communication capabilities Rinner and Wolf (2008), as depicted in Fig. 1.1. Smart cameras are used to autonomously perform vision tasks through on-board resources; then, aggregated high-level knowledge (e.g., objects bounding boxes, spatio-temporal trajectories of tracked

people, action recognition results) is delivered to third parties (data centers, end-users, other cameras). The data source of the processing pipeline in a smart camera is the image sensor, implemented either in complementary metal–oxide–semiconductor (CMOS) or in charge-coupled device (CCD) technology. The sensing unit collects the raw data from the image sensor; this propagates the information towards the processing unit, which is responsible for both image processing and decision-making tasks. At this point, the obtained abstracted/semantic data are eventually transferred to third parties through the communication unit. Smart cameras represent a fertile research field with a variety of different applications; it also conveys a number of heterogeneous methodological areas such as signal processing, machine learning, task and resource optimization, embedded computing, predictive control, and planning under uncertainty. For this reason, a lot of attention has been given to this research field over the last few years, both from academia and industry, as demonstrated by the increasing number of dedicated publications Esterle, Lewis, McBride, and Yao (2017); Kyung et al. (2016); Magno, Tombari, Brunelli, Di Stefano, and Benini (2013); Rinner, Winkler, Schriebl, Quaritsch, and Wolf (2008); Rinner and Wolf (2008); Varotto, Fabris, Michieletto, and Cenedese (2019a). More specifically, as mentioned in Sec. 1.3, smart sensors are not composed by the sensing module only; the intelligence of smart devices resides upon a computational unit, which enables complex data processing algorithms. In smart cameras, many research efforts have been devoted to the migration of computer vision on embedded platforms; this has facilitated the implementation of a variety of novel algorithms in real-life applications, from real-time quality assessment Chang, Mazzon, and Cavallaro (2018) and privacy protection Tadesse, Bent, Marcenaro, Weldemariam, and Cavallaro (2020) in wearable cameras, to mobile object detection, classification and tracking Meera et al. (2019). The computational module of smart cameras is also responsible for controlling and regulating the device configuration and sensing parameters for performance optimization. Smart cameras can be static, PTZ (Pan-Tilt-Zoom) Micheloni, Rinner, and Foresti (2010), or mobile Rinner et al. (2008): static cameras are mounted at fixed locations Aghajan and Cavallaro (2009); PTZ visual sensors can modify the focal length and rotate the optical axis Saragih, Kinasih, Machbub, Rusmin, and Rohman (2019a) (Fig. 1.2(a)); mobile cameras are usually mounted on ground or aerial robots, thus they can move in the space Schwager, Julian, Angermann, and Rus (2011) (Fig. 1.2(b)). Therefore, from the control perspective, smart cameras can be treated as robotic platforms characterized by a number of degrees of freedom and composed by internal sensing and physical configuration parameters and (e.g., the sampling rate and the PTZ parameters). Hence, the perception process can be cast to an optimization (control) problem that, given present and past measurements, regulates the camera degrees of freedom; this is the core idea behinds the active sensing paradigm as discussed in Chapt. 3.

Applications Smart cameras are currently used in many applications. Intelligent Video Surveillance Systems (IVSS) is a very active research area whose fundamental goal is to detect abnormal behaviors. This requires complex image analysis, starting from motion detection to segmentation, feature extraction, and classification. Smart cameras are also employed in Intelligent Transportation Systems (ITS), which can be divided into infrastructure-based systems and vehicle-based systems. The former are closely related to IVSS, but the main objective is related to traffic monitoring. On

(a) (b)

Figure 1.2. Examples of PTZ (a) and mobile (b) cameras.

the other hand, in vehicle-based systems the cameras are used inside automobiles to monitor the vehicle's surrounding environment, as well as the driver's state of attention while driving. These are fundamental building modules in adaptive cruise control and in forward collision-warning systems. Smart cameras can be used also in many medical applications like in robotic surgery, to automate experiments, or to monitor patients and medical personnel inside hospitals. Intelligent cameras have been extensively deployed also in machine vision systems, which deals with the application of computer vision methods to manufacturing, inspection, and robotics. Finally, there are many important applications for smart cameras in the fields of entertainment and smart environments (e.g., smart cities). All these applications often pose strong requirements on the processing speed and robustness of the vision methods, which are therefore embedded in the overall application.

1.3.3 Challenges and Trends

Future pervasive smart sensor systems will be targeted at end-user applications, where special emphasis is put on privacy and security, cost-efficiency, as well as simple deployment, supported by self configuration capabilities. Considering the recent advances of smart sensors in research and industrial practice, it is possible to identify several trends and research challenges to be addressed in the future.

- **Power consumption**: A core feature of smart sensors is that they are designed to run on battery power; hence, one of the main challenges is to find a reasonable trade-off between computing power, memory resources, communication capabilities, system size, and power consumption. These aspects can be regulated through the design of efficient logical architectures, as well as middleware algorithms that account for the power consumption of each operational unit. At the same time, in the area of embedded systems, advances in chip manufacturing technologies are expected to lower the power consumption and, therefore, increase the overall operating time of the system.
- **Adaptability and self-configuration**: The main advantage of smart sensing is the presence of embedded processing, which enable autonomous decision making

strategies; hence, smart sensors are devices with high levels of automation and adaptation capabilities. These properties are fundamental to simplify system deployment and minimize the human intervention: at first, the self-organizing functionalities make installation and integration of new applications easy also for non-expert users; in addition, autonomous sensor nodes can self-regulate their energy consumption, and they can also self-calibrate to adapt themselves in case of environmental changes.

- **Limited computational capabilities**: Many real-life applications (e.g., search and rescue Queralta et al. (2020)) require real-time processing speeds; nonetheless, self-operating agents, like UAVs, carry only a limited amount of payload Michieletto (2018); hence, it is often necessary to use cloud offloading for more computational power. This solution, however, slows down the computing process Krasniqi and Hajrizi (2016). This problem is particularly critical for deep learning, which is one of the most studied fields in machine perception. State-of-the-art deep learning models often lead to heavy and slow methods, but recent research has also focused towards the development of lighter and faster models, able to operate in real-time with limited hardware resources Moons, Bankman, and Verhelst (2019).

- **Privacy and security**: To manage the massive amounts of data generated from smart visual sensors, it is possible to exploit smart sensors' connectivity capabilities to transfer information to other sensor nodes, to human-readable end services, or to cloud infrastructures. The main problem of cloud computing is that it introduces several latency and privacy concerns. One approach to improving privacy is to not let raw data leave the sensors, delivering only encrypted information. Nonetheless, this makes it much more difficult to ensure the trustworthiness of the overall system Osia et al. (2020); hence, preserving individuals' privacy versus providing detailed data analytics faces a dichotomy in most pervasive sensing applications. The future trend is to implement pervasive computing and IoT infrastructures to the edge, rather than using the centralized cloud computing models. The main concept behind the edge computing paradigm is that computation and data storage are brought closer to the sources of data through distributed schemes (see Fig. 1.3). Edge computing is expected to improve response times and save bandwidth Yu, Liang, He, Hatcher, Lu, Lin, and Yang (2017); furthermore, by leveraging on distributed architectures, it also increases privacy by minimizing the transmission of sensitive information to the cloud.

- **Data accumulation**: The IoT technologies integrated in smart sensing devices produce large volume of data. Some data are generated directly by the sensor, such as measurements, device configuration, components status, log files, post-processing information (e.g., segmented images in computer vision applications), and, possibly, deep neural network models. The remaining data come from external sources, such as end-users or other connected devices (according to the ubuquitous computing paradigm). Such huge amount of data could overwhelm existing processing and storage techniques and systems. Hence, smart sensors must have the software and hardware capabilities to store, manage, and process the collected data efficiently. This requirement collides with the fact that many off-the-shelf devices have limited

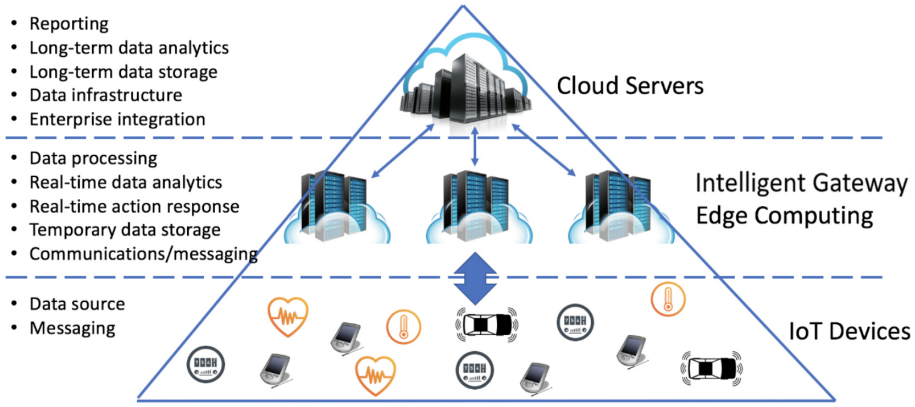

Figure 1.3. The basic edge computing infrastructure Yu et al. (2017).

storage resources, which impose severe constraints on the data accumulation capabilities. At the same time, excessive data accumulation affects the interpretability of a service. For these reasons, future research efforts should be devoted to the design of efficient data aggregation and selection criteria to prevent information overload and data accumulation, and to improve the interpretability of the sensing results.

- **Wearability**: Modern wearable devices can incorporate sensing, actuating, communication and energy storage functionalities; hence, they play a fundamental role within the IoT framework Balsamo, Merrett, Zaghari, Wei, Ramchurn, Stein, Weddell, and Beeby (2017). In particular, body worn sensors are becoming popular in sport, entertainment, medicine and security Brutti and Cavallaro (2017); Davison, Mayol, and Murray (2003); Gu and Ren (2015); Jiang et al. (2018) (Fig. 1.4). Despite this, there are still many significant research challenges around achieving robust solutions: the hardware miniaturization and integration should be designed to maximize the usability and the ergonomics; the materials should be flexible, stretchable and washable; the electronic components have to be robust and with low energy operation; software algorithms should mitigate environmental disturbances and ego-noise from the wearer (for instance, videos captured with body-worn cameras are often affected by motion blur Chang et al. (2018)).

Figure 1.4. Example of body-worn camera.

1.4 Multi-sensor systems

When networked or integrated, smart sensors can cooperate in order to carry out more complex tasks and achieve collective goals (e.g., maximise the number of tracked objects in visual sensor networks) Micheloni et al. (2010); Park and Oh (2020). A *multi-sensor system* integrates a number of sensors to achieve a specific task. Multi-sensor systems offer several advantages over single-sensor counterparts: at first, the on-board processing and storage resources allow to exploit information fusion techniques to refine the incoming data and, therefore, to attain better estimation and control performance Rao (2001); secondly, different sources of information enable high levels of parallelisation and specialisation Esterle et al. (2017), induce inherent robustness and complementarity, and allow inferences that are not possible with single-sensor measurements Nakamura, Loureiro, and Frery (2007). In general, multi-sensor systems are composed by sensors that can be homogeneous or heterogeneous, spatially distributed or embedded in a unique sensing platform, where, in this work, a *sensing platform* is any physical device endowed with one or more sensors on-board[4]. According to this definition, multi-sensor systems can be divided into single-platform and multi-platform solutions[5].

Multi-Sensor Single-Platform A Multi-Sensor Single-Platform system refers to a unique sensing platform comprising more than one sensor on-board (i.e., sensor-rich platform). The platform can be static Varotto and Cenedese (2021b) or dynamic Varotto et al. (2021), while the sensors can be homogeneous (e.g., multi-camera robots) or heterogeneous (e.g., robots endowed with audio-visual sensing capabilities). The former are referred to as *uni-modal platforms*, while the latter are called *multi-modal platforms*.

Heterogeneous perception capabilities is an increasingly prevalent property in cyber-physical systems Lewis, Goldingay, and Nallur (2014), due to the fact that multiple sensing modalities induce inherent robustness and complementarity (i.e., different properties of the environment can be perceived); moreover, the aggregated multi-modal data allow inferences that are not possible with uni-modal measurements. Nevertheless, heterogeneous data sources suffer from different types of noise, they might produce conflicting features, and they often require different calibration procedures, each of which might involve extensive human intervention (see Chapt. 3 for further details). Examples of multi-modal platforms are audio-visual sensing devices, composed of microphone array and a monocular camera Qian, Brutti, Lanz, Omologo, and Cavallaro (2021a,b); Wang, Sanchez-Matilla, and Cavallaro (2019): the camera provides information-rich data, while audio signals allow to better handle occlusions and to collect information from a wider sensing domain. Another example is provided by the turtlebot in Fig. 1.5; it allows to integrate a variety of different sensors (cameras, LiDAR, IR) into a single robotic platform.

[4]This definition can be eventually extended to the RPA framework, where a sensing platform is not a physical device, but a digital entity; in this context, it is possible to extend also the definition of sensor, from a device that measures physical properties of the environment, to software tools that extract information from files or cloud databases.

[5]This Chapter provides a preliminary and qualitative categorization; for a more formal discussion, the reader is addressed to Chapter 3.

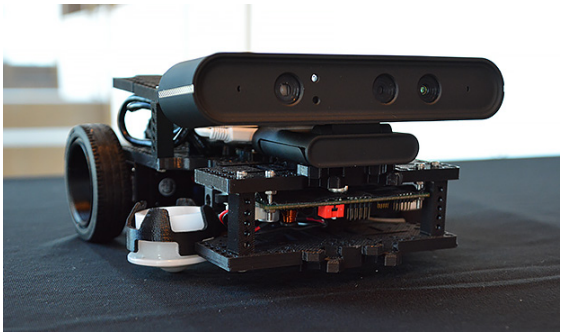

Figure 1.5. Turtlebot3 is a typical multi-modal platform. It is a small ground robot set up for a variety of USB cameras, RGB-D sensors, lidars, IR and ultrasonic sensors as well.

Multi-Sensor Multi-Platform Multi-Sensor Multi-Platform systems exploit the communication capabilities of pervasive sensors to create *sensor networks*, which consist of spatially distributed devices that are employed to monitor environmental conditions (e.g, temperature), detect anomalous events, or track the trajectory of people or objects.

A sensor network is said to be *wireless* (WSN) if each node in the network is equipped with wireless communication technologies (e.g., bluetooth or WiFi Zafari, Gkelias, and Leung (2019)), as well as with an on-board energy storage system (e.g., a battery). A WSN is then said to be *smart* when any of its node is a smart sensor Zhang et al. (2004). A WSN is *cooperative* when its agents (i.e., the sensor nodes) share information with the purpose of taking collective decisions; the final aim is the improvement of the perception capabilities in terms of quality and speed Hu, Xie, Xu, and Xu (2014). Thanks to their cost-effective and easily deployable nature, wireless sensors are often employed in *large scale networks*; these require the design of efficient information fusion and decision making strategies to manage the sensors' energy consumption, as well as the communication and processing workload over the network. The evolution in robotics and embedded systems have enabled the introduction of *mobile sensor networks* (see Fig. 1.6). These are composed of smart sensor nodes with additional degrees of freedom which allow them to move in space; examples are sensorized ground, underwater or aerial autonomous robots[6]. Swarms of mobile sensors have the ability to simultaneously gather information from disjoint locations, enlarging the spatial and temporal coverage of the network. Fig. 1.6(a) shows an aerial mobile camera network that attains optimal area coverage through the control of the UAVs position and the cameras PTZ parameters. Finally, a sensor network is *multi-modal* if it is composed by heterogeneous sensors. In addition to perception heterogeneity, certain applications require the simultaneous use of robots with *actuation and morphological heterogeneity* (see Fig. 1.6(b)); this is common in cooperative localization, surveillance, assistive navigation, exploration and mapping Chen, Xu, Ren, and Chirarattananon (2021)[7]. For instance, smaller platforms have access to narrow passages, but their payload is limited Chen et al. (2021); therefore, it usually advisable to employ larger robots that carry application-specific payloads (e.g.,

[6]In this work, the emphasis is on autonomous perception systems; hence, a sensorized robot is considered as a mobile sensing platform.

[7]To avoid confusion, the term *heterogeneous robots* is used when the focus is on actuation and morphological properties; otherwise, when the emphasis is on the perception capabilities, the term *heterogeneous sensors* is used.

long-duration batteries, heavy sensors or actuators) Rizk, Awad, and Tunstel (2018).

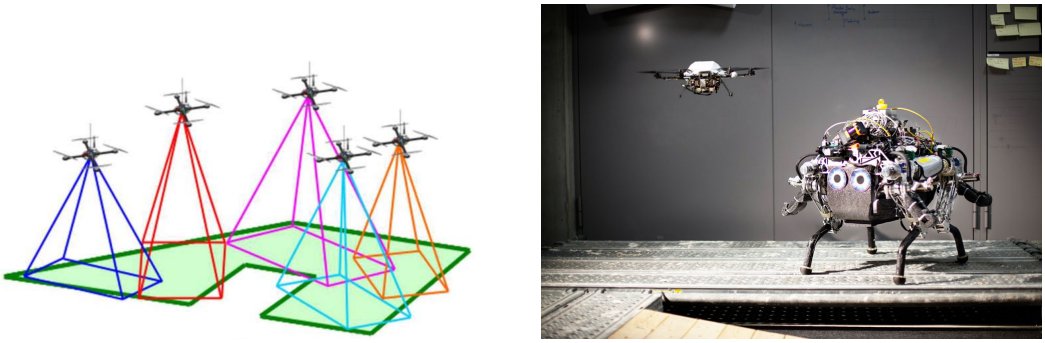

(a) (b)

Figure 1.6. From a perception point of view, a group of connected sensorized robots is a mobile WSN. (a) Swarm of camera-equipped UAVs representing a homogeneous and uni-modal mobile camera network Schwager et al. (2011); (b) two heterogeneous mobile robots, where the heterogeneity is mainly related to the morphological and operational characteristics of the robots.

1.4.1 Applications

Multi-platform solutions are often employed to compensate the limited computing and sensing capabilities of the individual sensors; in addition, sensor networks are more fault-tolerant with respect to systems relying on a single-platform, even if they come at a cost of higher logical and structural complexities. Multi-sensor systems are common in a variety of applications, ranging from civil and private contexts to emergency missions in post-disaster scenarios. In industrial environments, collaborative fleets of sensor-rich mobile robots are being utilized in the manufacturing sector and in logistics warehouses. The high covering capabilities of large scale WSNs are often exploited to monitor widespread areas for the surveillance of urban streets Varotto, Zampieri, and Cenedese (2019b), airports, and shopping centres. Nowadays, modern sensor fusion techniques simplify the deployment of sensor networks also in ambient intelligence and Intelligent Transportation Systems (ITS), especially regarding tasks like localization and tracking Varotto et al. (2021); Varotto and Cenedese (2021c), advanced driver-assistance systems Varotto and Cenedese (2021a), and access control systems Varotto and Cenedese (2021b). At the same time, cooperative multi-robot systems have been shown to significantly improve the efficiency of Seach and Rescue (SAR) missions, with faster search of victims, initial assessment and mapping of the environment, and real-time monitoring Queralta et al. (2020). Finally, smart sensor networks provide numerous opportunities for smart grid applications including power monitoring, demand-side energy management, coordination of distributed energy storage, and integration of renewable energy generators Jaradat et al. (2015). In smart grids, intelligent monitoring and sensing capabilities are necessary to insure real-time response from the power grid. In particular, the growing electrification process resides on the pervasiveness and efficiency of electric energy supply systems; hence, a reliable and real-time monitoring is highly required to provide solutions quickly, especially when natural accidents severely affect the transmission lines status. In this

context, wireless sensors are also used to collect real-time weather information, so that it is possible to accurately predict the energy availability in the near future; these forecasts are particularly crucial for energy scheduling models relying on renewable energy sources, like wind energy and solar energy.

1.4.2 Use case: smart camera networks

Smart camera networks are a prominent example of multi-sensor multi-platform systems. As discussed in Sec. 1.3.2, a smart camera can be either fixed, PTZ or mobile. Therefore, from the perception point of view, a smart camera network is obtained from the integration of a number of fixed or PTZ smart cameras in a common network, as in motion capture systems (see Fig. 1.7); alternatively, smart camera networks can also be composed by a set of cameras mounted on interconnected ground or aerial vehicles (see Fig. 1.6(a)). In both cases each camera has reasonable computing and communication capabilities, therefore such networks of smart cameras can also be treated as systems with distributed sensing and image processing resources. These systems can be organized either in a centralized or decentralized manner. In the former case, some sort of pre-processing is done on the individual cameras, but the most of the computations are performed at a central server, which merges information provided by individual cameras. On the other hand, in decentralized architectures cameras organize themselves and collaborate on a certain task; video streams can be analyzed independently and only abstracted information is delivered to other nodes alternatively, videos can be jointly and cooperatively processed by sharing information in a distributed manner. In general, the embedded intelligence of each camera allows to integrate adaptivity and autonomy to the network, and it is fundamental to stimulate the emergence of collective and cooperative behaviors Micheloni et al. (2010).

The ultimate vision is to provide a service-oriented network which is easy to deploy, adapts to changes in the environment, and provides various customized services to users. Clearly, this increases the complexity of the system as future smart camera networks will operate on a vast scale, drawing on mobile, internet-connected resources. They will comprise heterogeneous and decentralised aggregations of visual sensors with other non-visual sensors. Furthermore, the potential to include citizen-contributed mobile streaming, videos coming from body-worn and robot-mounted cameras, alongside more traditional fixed or PTZ cameras, leads to a number of difficult and important challenges Esterle et al. (2017).

Applications The aim of distributing smart cameras is mainly enhanced reliability, coverage of larger areas and the ability to resolve occlusion Mavrinac and Chen (2013). These properties are integral part of applications like unmanned search and rescue operations, disaster management, wildlife conservation, environmental mapping, and multi-camera target tracking. In healthcare instead, smart camera networks can assist the monitoring of at-home patients. Many of these applications require cameras to develop autonomous, cooperative and collective intelligence; this is useful to address unforeseen situations, predict the presence of events of interest, manage fluctuating resources, and regulate the network parameters for task optimization. Some examples of collective intelligence are given by cameras selection for video surveillance systems Varotto et al.

Figure 1.7. Motion Capture Systems are typically composed by several cameras to resolve occlusion problems (https://www.vicon.com).

(2019b), pan-tilt control for minimal FoV intersection in environmental monitoring, and multi-task allocation Esterle et al. (2017); SanMiguel et al. (2014).

1.4.3 Challenges and Trends

Multi-sensor systems represent a multidisciplinary research field, at the frontier between robotics and information technology. In particular, the high complexity of smart sensor networks poses new challenges that are based upon the system robustness and reliability, real-time operation, cooperative self-reconfiguration and self-healing capabilities. The next steps to the future development of intelligent sensor networks are as follows.

- **System reliability and fault tolerance**: if a single node of the network goes down, the network should continue working without interruptions. Moreover, it must provide actions to detect and recover from the failure Khalastchi and Kalech (2019).
- **Communication robustness**: network bandwidth, latencies, limited communication ranges, and communication faults represent the bottleneck in the implementation of most cooperative algorithms, especially in mobile robotics. For this reason, multi-robot coordination in communication-constrained environments will be a fertile research field in the future Kemna, Rogers, Nieto-Granda, Young, and Sukhatme (2017).
- **Deployment in real scenarios**: little attention is given to the analysis of multi-sensor systems in real-life scenarios. Therefore, more research (and development) is needed on software and hardware components in order to facilitate inter-node communication and to simplify the deployment procedures. To this aim, ubiquitous smart sensor nodes should form dynamic networks, where each node can be easily removed and added by non-expert users, supported by the self-organizing functionality of the network. In this direction, it is fundamental to design service-oriented architectures that facilitate operations like maintenance, installation and integration.
- **Multi-modality**: researchers are expected to start integrating different sensors (audio, seismic, thermal, visual, radio) into distributed multi-sensor networks. By

Figure 1.8. Examples of heterogeneous cameras (fixed, PTZ, mobile, and body-worn) Esterle et al. (2017).

fusing data from multiple sensors, the system exploits the distinct characteristics of the individual sensors, resulting in an enhanced overall output.

- **Distributed intelligence and collective learning**: multi-agent systems are mostly used to carry out tasks where the emergence of collective behaviors provides some benefit. Hence, one of the main future challenges in machine learning and robotics research is the design of multi-agent architectures where each device contributes to the collection of data used to increase the intelligence of the entire system, considered as a whole. Some of the main methodological concepts related to this task are multi-agent reinforcement learning Hernandez-Leal, Kartal, and Taylor (2019), game theory Grünwald, Dawid, et al. (2004) and federated learning Li, Sahu, Talwalkar, and Smith (2020).

- **Heterogeneity**: today's technology is mature to envision large networks composed of heterogeneous smart devices (static, wearable, and mobile). These may interact and coordinate autonomously, being aware of their sensing, physical and operational differences. Fig. 1.8 shows examples of heterogeneous cameras that can cooperate within the same network[8]. In this network, each device is specialized into a specific task: fixed cameras constantly monitor environmental sectors of high relevance, PTZ cameras are regulated to focus on sporadic event of interest, body worn cameras are used to take detailed close-up videos, mobile cameras exploit their dynamic nature to provide a panoramic overview of the environment, or to accomplish tasks that require robotic actuation capabilities (e.g., object picking).

- **Real-time requirements and synchronization**: multi-sensor systems are designed to work with large amount of data and processing workloads; at the same time, it is necessary to deal with network and device limitations, as well as strict synchronization among all network nodes. As a consequence, working under real time requirements is a difficult constraint to be guaranteed.

[8]Note that, in this case, the network is homogeneous from the sensing perspective (i.e., uni-modal), but heterogeneous from the operational point of view.

- **Resource and task allocation, data dissemination**: one of the main advantages of multi-agent systems is that the nodes can cooperate to optimize the shared usage of processing and memory resources, so that the network lifetime is maximized. However, in large networks, task allocation problems often lead to an intractable combinatorial explosion of possible configurations. Thus it is necessary to introduce some automation, either through a-priori optimisation, or by online machine learning techniques.

- **Collaborative self-organization**: future research in control theory and machine learning will guarantee more efficient inter-node collaboration for supporting self-organizing and self-configuring networks. In particular, smart nodes are expected to adapt themselves to environmental changes and non-fixed network topologies. The former are very common in real-life applications and are mainly caused by external events, such as illumination and weather conditions in video surveillance systems, dynamic obstacles in smart warehouses, time-varying occlusion, and cluttering in domestic and industrial environments. Topology variations, instead, are often due to changes in the network behavior. For instance, mobile robots introduce continuous topological fluctuations, communication faults affect the inter-node communication topology, while camera pan and tilt rotations modify the sensing topology in visual sensor networks (i.e., the FoVs projected into the environment are time-varying). To deal with these scenarios, smart sensors should be endowed with self-calibration SanMiguel et al. (2014); Varotto et al. (2019a), self-localization, and self-reconfiguration capabilities SanMiguel et al. (2014) In this regard, one of the main challenges to be tackled is the fact that self-organisation in sensor networks often requires a high degree of interaction among the individual devices; this can be error prone and may result in higher communication demands.

1.5 Research activity outline and structure of the manuscript

The focus of this Thesis is on the design of autonomous data acquisition and perception schemes for multi-sensor applications. In particular, the dissertation explores a variety of use-cases, with the aim of analyzing and solving some of the open questions related to robotic perception systems and smart connected sensors. After this Chapter where general concepts have been introduced as regards to pervasive and smart sensing technologies, the rest of the Thesis in organized into four main parts.

1. Chapt. 2 provides the background and the main motivations towards the Multi-Sensor Probabilistic Active Sensing (MS-PAS) framework.

2. Chapt. 3 introduces the first contribution of this work, by formalizing the MS-PAS paradigm (under the context of active position estimation).

3. Chapters 4-7 represent the Thesis core, since they describe the predominant research studies carried out during the Ph.D. activity. The objective of these Chapters is to investigate the applicability and effectiveness of the proposed MS-PAS framework in some application scenarios; these include Localization (Chapt. 4), Tracking

(Chapt. 5), energy-constrained video surveillance systems (Chapt. 6), and IoT for connected vehicles (Chapt. 7). In each Chapter, the devised techniques are validated through comparison with other literature solutions, according to specific performance metrics; the results are supported with theoretical statements, numerical experiments of realistic case studies, and, when available, with real data.

4. Chapter 8 summarizes the principal achieved results, draws the general conclusions and presents potential future works for the contents of this dissertation.

The manuscript ends with the appendix in Chapter A, regarding Gaussian Process Regression, Bayesian Optimization, and Recursive Bayesian Estimation.

1.5.1 List of publications

The following list provides the full catalog of the author's publications. The items marked with a blue asterisk (*) are considered relevant for the Ph.D. activity and hence for this Thesis. Other publications concerning these topics are envisaged.

*Varotto and Cenedese (2021a) L. Varotto and A. Cenedese, "*Online and Adaptive Parking Availability Mapping: An Uncertainty-Aware Active Sensing Approach for Connected Vehicles*". IEEE 32nd Intelligent Vehicles (IV2021) - 2nd workshop on Online Map Validation and Road Model Creation (MAVROC), 2021 [accepted]

*Varotto and Cenedese (2021c) L. Varotto and A. Cenedese, "*Transmitter Discovery through Radio-Visual Probabilistic Active Sensing*". IEEE 25th International Conference on Methods and Models in Automation and Robotics (MMAR), 2021 [accepted]

*Varotto and Cenedese (2021b) L. Varotto and A. Cenedese, "*Probabilistic RF-Assisted Camera Wake-Up through Self-Supervised Gaussian Process Regression*". IEEE 29th Mediterranean Conference on Control and Automation (MED), 2021 [accepted]

*Varotto et al. (2021) L. Varotto, A. Cenedese and A. Cavallaro, "*Probabilistic Radio-Visual Active Sensing for Search and Tracking*". European Control Conference (ECC), 2021 [accepted]

Cenedese and Varotto (2019). Angelo Cenedese and Luca Varotto. "*A Distributed Approach to 3D Reconstruction in Marker Motion Capture Systems*". ACM 13th International Conference on Distributed Smart Cameras (ICDSC), 2019, Article 17, pp. 1–6, doi:10.1145/3349801.3349818

Varotto et al. (2019b). L. Varotto, A. Zampieri and A. Cenedese, "*Street Sensors Set Selection through Road Network Modeling and Observability Measures*". IEEE 27th Mediterranean Conference on Control and Automation (MED), 2019, pp. 392-397, doi: 10.1109/MED.2019.8798593

Varotto et al. (2019a) L. Varotto, M. Fabris, G. Michieletto and A. Cenedese, "*Distributed Dual Quaternion Based Localization of Visual Sensor Networks*". IEEE 18th European Control Conference (ECC), 2019, pp. 1836-1841, doi: 10.23919/ECC.2019.8796165.

2

Multi-Sensor Probabilistic Active Sensing: Framework and Contributions

Smart sensors are endowed with decision-making and communication capabilities; thus, they can exploit self-organizing schemes to control and optimize their perception and data collection process. This Chapter formalizes the paradigm of Probabilistic Active Sensing for multi-sensor scenarios; it serves as reference for the application-specific variations that will be presented in the successive Chapters.

Contents

2.1 Active Sensing

An important component of any robotic and data acquisition system is the processing and interpretation of the available sensory information. At the lowest level, the sensor data are used to derive regulatory and actuation signals, while at a higher level they are used to create models of the environment Papanikolopoulos (1996). In this regard, to fully exploit the potential of modern smart data acquisition systems, many open challenges remain to be addressed. In particular, as mentioned in Chapt. 1, the era of ubiquitous and mobile computing has opened up many exciting research directions in robotic perception, sensor fusion, and information control Bajcsy, Aloimonos, and Tsotsos (2018). In this regard, *active sensing* (AS) Bajcsy et al. (2018) represents a research framework that stands at the nexus between automation and robotics research, leveraging the intimate interplay among estimation, perception, and control. According to Bajcsy et al. (2018), active sensing is the problem of *intelligent control strategies applied to the data acquisition process*; hence, it consists in the control of a dynamical system with actuation and sensing capabilities (e.g., eye-in-hand manipulators Radmard and Croft (2017)). In line with this definition, an agent is an active perceiver if it knows why it wishes to sense, it chooses what to perceive, and determines how, when and where to achieve that perception Bajcsy et al. (2018). The objective of AS is to automate the perception process and maximize its efficiency, in terms of acquired information or task completion Bajcsy et al. (2018). If follows that any mature active sensing platforms should offer a controllable, flexible and adjustable sensory and processing apparatus. Second, the system must have exploratory capabilities, that is, it must also be able to evaluate the information gain obtained with each new measurement. Third, the system must be selective in its data acquisition, in order to be able to act upon the perceptual information in real time. Fourth, the system must be able to learn Bajcsy et al. (2018). Given these premises, most AS schemes consider a smart sensor as an intelligent agent that can reason and decide on the next sensing actions, with the aim to maximize a set of target criteria, given its current state; thereby, AS allows to collect the most useful data, given platform-specific constraints. In doing so, AS can trade off between the level of task completion and the finite quantity of sensing resources (e.g., energy consumption, mission time, travel distance). For instance, in object recognition tasks not all input images contribute equally to the quality of the output model, since several images may often contain redundant visual information; therefore, image selection saves unnecessarily increased processing times. On the other hand, a certain degree of redundancy can help to improve the reconstruction in more difficult regions. It is therefore necessary to propose AS schemes which result in a balanced image selection process Hornung, Zeng, and Kobbelt (2008).

Fig. 2.1 provides an example of an active sensing system (eye-in-hand manipulator), compared to a passive counterpart (eye-to-hand manipulator) within a visual servoing application. The eye-to-hand scheme is composed by a sensor-free robotic arm that needs to grasp some objects inside a box; the control is responsible for getting the end-effector closer to the objects inside the box, possibly under some safety and performance guarantee (e.g., collision avoidance, fast task execution). The control feedback exploits

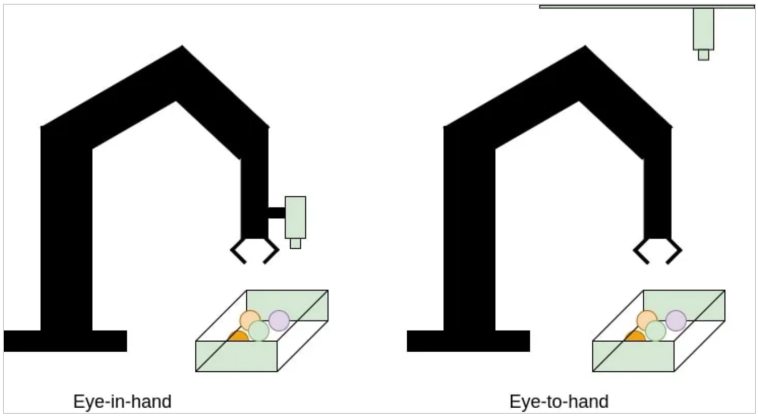

Figure 2.1. Examples of active (a) and passive (b) sensing systems in robotics.

measurements coming from a static camera, which determines the pose (position and orientation) of the end-effector and the position of the objects to be grasped. On the other side, the eye-in-hand manipulator configuration has a camera directly attached to the end-effector; hence, the arm becomes a mobile sensing platform (with visual sensing capabilities). Similarly to the passive case, the camera observations are used to feed the robot controller, but the main difference is that the control inputs modify both the end-effector location and the sensor configuration (i.e., camera pose and Field of View). Therefore, each sensor observation is affected by the past and present control actions. In conclusion, the fundamental difference between active and passive sensing is that the former provides the possibility to dynamically modulate the overall agent's behavior by controlling the measurement process and acting on the sensor configuration; in the latter, instead, the sensor parameters are constant and the observations depend only on the properties of the monitored environment Bajcsy et al. (2018).

2.1.1 Motivations

To introduce the main benefits and motivations towards active sensing approaches, we still refer to Fig. 2.1. As mentioned in the previous Section, active sensing provides the possibility to act on the measurement process; in turn, this provides the following advantages.

- **Performance optimization**: suppose that the robot in Fig. 2.1 needs to collect all the objects inside the box; if some objects are occluded from the eye-to-hand camera, or if the resolution is not sufficient to detect all the elements inside the container, the task can not be completely accomplished and only a sub-optimal result can be attained. On the other hand, the possibility to move the eye-in-hand camera allows to better explore the environment and, therefore, to handle possible bad initial resolutions or to tackle occlusion issues. In general, active sensing allows to implement a proactive sensor behavior, which seeks to attain the best performance possible by self-regulating its configuration. Without these capabilities, the sensor (and, thus, the overall system) is typically more bond to the initial setup conditions.

- **Adaptation to dynamic and unstructured environments**: if the container in Fig. 2.1 is movable (e.g., mounted over a ground robot), it can leave the camera FoV, and the grasping task may be compromised. In this case, a mobile visual sensor is able to track the box and guarantee a certain task continuity level. In general, a traditional sensor with fixed structure is often inadequate to perceive features in dynamic and uncertain environments; conversely, dynamic reconfigurable sensors can better fit the environmental properties at runtime.
- **Multi-tasking**: suppose now that the objective is to estimate the objects shape and pick only a subset of them, according to the estimation result. To this aim, the camera needs to take several pictures at different perspectives for each object. Obviously, this is not possible in the eye-to-hand scenarios (especially if the robot can not manipulate the box), while the eye-in-hand setup allows to generate suitable control inputs so that the final estimation result is accurate enough to proceed with the successive phase. In general, active sensing allows to combine powerful and versatile exploration techniques (e.g., object shape estimation) with more executive ones (e.g., objects picking).
- **Larger number of degrees of freedom**: active sensing is the control framework that allows to exploit the degrees of freedom of modern smart sensing platforms. As a matter of fact, PTZ and mobile cameras offer their main benefits only when their configuration is regulated and modified over time.

2.1.2 Applications

In the previous Sections visual servoing has been used as use case to introduce the concept of active sensing and its main advantages over passive sensing. Indeed, the main active sensing applications are related to the robotic field, both for research and industrial purposes. In particular, by enabling autonomous perception in robotic systems, active sensing has been successively employed in object recognition, classification and manipulation Calli, Caarls, Wisse, and Jonker (2018); Patten, Martens, and Fitch (2018), collaborative mobile robotics Carron, Todescato, Carli, Schenato, and Pillonetto (2015); Haugen and Imsland (2015); Mathew, Smith, and Waslander (2015); Nagaty, Thibault, Trentini, and Li (2015); Yu, Budhiraja, and Tokekar (2018); Zhou (2012), active sensor calibration Papanikolopoulos (1996), environmental mapping Masehian et al. (2017); Popovic (2019), and exploration Bai, Wang, Chen, and Englot (2016); Ghassemi and Chowdhury (2020); Mavrommati, Tzorakoleftherakis, Abraham, and Murphey (2017).

Active sensing is also fostered by the higher performance guarantees with respect to passive data collection systems Czarnetzki, Kerner, and Kruse (2010); Patten et al. (2018); hence, it has been used for optimal sensor coverage Fabris and Cenedese (2019), intelligent video surveillance systems SanMiguel et al. (2014), and target tracking Varotto et al. (2021). The capability to balance task execution and knowledge refinement has been proven to be fundamental in applications like search and rescue missions Du, Zhang, Ling, and Zheng (2019); Lin and Goodrich (2009); Murphy, Sreenan, and Brown (2019), wildlife monitoring Cliff, Saunders, and Fitch (2018); Katenka, Levina, and Michailidis (2013); Rodríguez, Negro, Mulero, Rodríguez, Hernández-Pliego, and Bustamante (2012), and

autonomous source term estimation Bourne, Pardyjak, and Leang (2019); Hutchinson, Liu, Thomas, and Chen (2019b); Hutchinson, Oh, and Chen (2017); Masson, Bechet, and Vergassola (2009); Park and Oh (2020); Ristic, Angley, Moran, and Palmer (2017); Ristic, Skvortsov, and Gunatilaka (2016). The high flexibility of the active sensing paradigm is proven by the fact that it can be employed in digital games and virtual reality. For instance, the authors of Mousavi, Liu, Yuan, Takáč, Muñoz-Avila, and Motee (2019) cast an image classification problem into a navigation and planning mechanism, leveraging on visual attention and deep reinforcement learning; they design a control architecture that learns to complete a task and, at the same time, it collects only the necessary observations relevant for the classification problem. Finally, thanks to the high performance guarantee in terms of sensor management and information processing, active sensing is today an important player in many IoT scenarios Varotto and Cenedese (2021a) and smart factories Zhang, Wang, Du, Qian, and Yang (2018).

2.1.3 Use case: Active Vision

Active vision [1] applies the active sensing paradigm to camera-based scenarios; in doing so, active vision combines modern computer vision techniques with intelligent control algorithms. Active vision involves the continuous interaction between smart cameras and the environment to decide what to observe and how to prioritize observation tasks. This information is then used to actively change smart cameras intrinsic and extrinsic parameters (i.e., pose and PTZ), with the aim to adapt their FoVs, according to specific (possibly time-varying) tasks. By combining adaptive control techniques with computer vision algorithms, the active vision framework introduces new capabilities, such as the derivation of depth maps from controlled motion, the development of efficient vision-assisted grasping, the design of active calibration and visual inspection techniques Häne, Heng, Lee, Fraundorfer, Furgale, Sattler, and Pollefeys (2017); Papanikolopoulos (1996); Varotto et al. (2019a), the minimization of the detection time in objects search Ryan (2008), the maximization of accuracy and efficiency in localization and mapping problems Popovic (2019), the gaze control for simultaneous target tracking and area coverage Lathuilière, Massé, Mesejo, and Horaud (2019), and the adaptation of camera configuration to dynamic lighting conditions, to name a few. In this regard, pan and tilt movements are often used for viewpoint optimization, while zoom is usually regulated to increase the image resolution and, therefore, to maximize the quality of visual data. If the platform is mobile, PTZ movements can be substituted Dogancay (2012) (or integrated Varotto et al. (2021)) with spatial degrees of freedom. In particular, zoom-in and zoom-out actions are equivalent to a platform coming closer and further; on the other hand, pan and tilt movements are often replaced by the roll-pitch-yaw orientations in aerial robotic platforms. More seldom, active vision is also used to tune internal camera parameters, with the purpose of enhancing the image quality; for instance, it is possible to control the camera exposure Kim, Cho, and Kim (2018a) in order to prevent image saturation caused by dynamic lighting conditions.

Applications Active vision algorithms are utilized in industrial inspection, object

[1] The reader is addressed to Chen, Li, and Kwok (2011) for a survey on active vision applications and techniques.

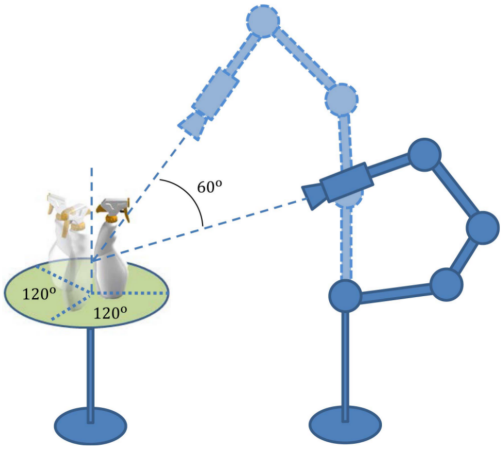

Figure 2.2. Active vision system used to estimate the shape of an object.

recognition and reconstruction, security and surveillance, site modeling and exploration, mapping, navigation, tracking Chen et al. (2011). In robotics, active vision is used to supply better visual data for the execution of a given task by systematically changing the viewpoint of the robot's sensor Calli et al. (2018). In the object grasping and manipulation domain, the goal of the active vision system is to increase the success rate of grasp synthesis algorithms for manipulating objects whose models are not available a priori Calli et al. (2018); Pang, Xompero, Oh, and Cavallaro (2021). The first step consists in the reconstruction of some properties of the target object (e.g., shape); to this aim, the camera is utilized to collect images at multiple viewpoints, since a global description of objects often can not be reconstructed from only one viewpoint, due to occlusion, measurements uncertainty and ambiguities (e.g., from a single image it is not possible to retrieve the distance of the object). The estimation process is then used to control the movements of the robotic arm and fingers in order to achieve a successful grasp. In this case, the camera degrees of freedom allow the reformulation of the manipulation task as an optimal control problem, where the viewpoint optimization algorithm maximizes the quality of the grasp and, therefore, its probability to be successful. In particular, the majority of the strategies design the optimization problem such that each new camera action provides images that minimize the uncertainty on the target physical properties and for which the confidence in generating a good grasp is the highest (see Fig. 2.2). In some circumstances it is also helpful to integrate terms related to occlusion minimization and collision avoidance Ji, Huang, and Huang (2019). In any case, it is necessary to adopt an eye-in-hand or a PTZ eye-to-hand configuration (see Sec. 2.1).

2.2 Probabilistic Active Sensing

Probabilistic AS (PAS) exploits incoming data to generate a belief map Queralta et al. (2020). This encodes the knowledge gathered during the sensing mission and it is used to guide the platform decision-making process for the generation of the successive actions. For instance, in sensor planning for object search (i.e., the problem of finding an object in an unknown space) Liu, Li, and Hedrick (2015), the belief map models the target

probabilistic location Liu and Hedrick (2017) and allows to implement selective attention mechanisms. The search task is a nondeterministic polynomial-time (NP)-complete problem, whose goal is to maximize the probability of detecting the target with minimum cost; hence, the next best view is usually chosen according to an optimization problem over the probabilistic map. This means that the camera FoV at each step is selected based on statistical metrics regarding the quality of the observations (e.g., the information gain, or the probability of detection) [2]. More in general, Fig. 2.3 depicts a generic bi-modal probabilistic map, which may represent the spatial concentration of gas particles, the probabilistic location of an object, or the uncertainty level associated to a navigation and mapping task. As shown in the Figure, the platform moves from its initial position towards the highest peak of the map, following a non-linear trajectory.

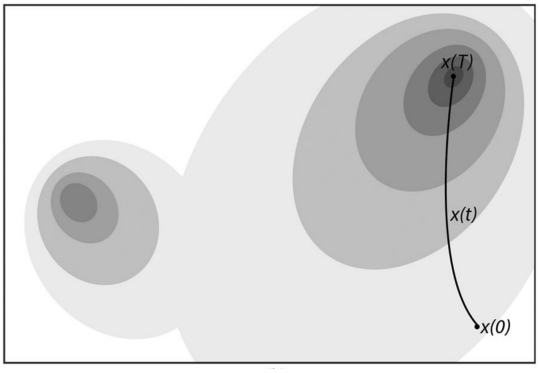

Figure 2.3. Example of platform trajectory guided by a probabilistic map, according to the PAS paradigm.

Major benefits of PAS versus AS reside in the fact that probabilistic approaches account for *realistic perception uncertainties* Radmard and Croft (2017); hence, they are suitable to manage real-world (noisy) scenarios, unmodeled dynamics, and sensing nuisance. Furthermore, probabilistic decision making has high *adaptivity* properties Popovic (2019), it allows to include *a-priori knowledge* Ramirez, Doucette, Curtis, and Gans (2014), and it *scales well* with high-dimensional control spaces (e.g., in environment exploration Bai et al. (2016)). On the contrary, under these conditions, deterministic approaches often become computationally intractable Trummel and Weisinger (1986) and ineffective.

For the aforementioned reasons, PAS is expected to be a flourishing research area with a direct application in environmental mapping and exploration, autonomous source term estimation, search and rescue, and collaborative mobile robotics.

2.3 Multi-Sensor PAS

Novel technologies and the evolution in embedded systems have enabled the integration of different types of sensors to perform autonomous sensing, among which video Pérez-Carabaza, Besada-Portas, Lopez-Orozco, and Pajares (2019a), radio Koohifar, Guvenc, and Sichitiu (2018), and acoustic Haubner, Schmidt, and Kellermann (2019a) sensors

[2] The reader is addressed to Chapters 3-5 for a deeper discussion on PAS applied to target search and tracking.

play a major role, paving the way towards *multi-sensor* PAS (i.e., MS-PAS) Lathuilière et al. (2019); Rajasekaran, Ahmed, and Frew (2020); Varotto et al. (2021). The coupling and integration of different information sources (either homogeneous or heterogeneous) opens up new perspectives in scene perception, since the uncertainty of measured data is mitigated through sensor fusion Khaleghi, Khamis, Karray, and Razavi (2013). Moreover, the aggregated data allow inferences that are not possible with single-sensor measurements: multi-sensor platforms enable parallelisation and specialisation Esterle et al. (2017), while sensors heterogeneity induces inherent robustness and complementarity (i.e., different properties of the environment can be perceived) Varotto et al. (2021).

With the aim of underling the importance of sensor fusion in real-life applications, two prominent use cases are provided in the following: Sec. 2.3.1 discusses the role of active vision in smart camera networks and multi-modal camera networks; Sec. 2.3.2 highlights the role of active sensing and multi-sensor data aggregation in robotic object manipulation and grasping.

2.3.1 Use case: collaborative active vision

Static camera networks allow to cover wide areas in surveillance and monitoring activities; nevertheless, their fixed parameters may represent an important limitation in terms of flexibility and adaptability. More specifically, high coverage capabilities and good object recognition performance are two conflicting requirements in multi-camera systems: the former demands wide-camera angles, while the latter necessitates for high image resolutions. To compensate this problem, camera networks that require both panoramic and detailed images, are often deployed with a certain degree of redundancy, namely, with overlapping FoVs between different sensors. This approach introduces robustness with respect to occlusions and obstacles (since any point in the environment belongs to two or more FoVs), but, at the same time, it also poses severe economical and operational limitations. In fact, the FoV overlapping constraint requires a huge amount of video sensors; therefore, a camera network composed by specialized sensors does not scale well with the dimension of the monitored environment, with the number of parallel tasks to be accomplished, and with the level of environmental variability.

The introduction of collaborative PTZ cameras brought new capabilities to surveillance networks, so that higher flexibility and scalability performance criteria can be reached. PTZ cameras can adapt their FoVs with the aim to focus the attention on areas of interest. Hence, some problems related to non-overlapping FoVs can be overcome (e.g., occlusions, localization and low target resolution) SanMiguel et al. (2014). In particular, smart camera networks implement the paradigm of *collaborative active vision* by integrating intelligent visual nodes that self-regulate their parameters according to the tasks to be accomplished, the overall network status (e.g., power consumption, bandwidth usage, number of sensors switched on), and the network actuation capabilities. In addition to the usual advantages of standard active vision solutions (see Sec. 2.1.1), collaborative camera networks typically provide higher levels of parallelisation Rinner et al. (2015): optimization methods enable topology reconfiguration; in this way, the areas of major importance are covered and the events of interest are tracked, or further analyzed, via zoom-in resolutions. Another advantage of collaborative camera networks over single-

camera active vision systems is that the data exchange capabilities improve self-diagnostic and fault tolerant algorithms Khalastchi and Kalech (2019); Yang, Han, Ge, Ding, Xu, Jiang, and Zhou (2019a). Finally, the use of sensor fusion techniques generates information that can not be obtained with monocular vision only, such as depth estimation Aghajan and Cavallaro (2009).

Collaborative active vison for tracking The following case studies show that target tracking is one of the main applications that takes advantage from the deployment of collaborative smart camera networks.

- *Tracking with handover*: a single PTZ camera is usually not sufficient to cover large-scale environments, due to the physical saturation limits of the pan and tilt rotations; this becomes particularly critical in presence of obstacles and occlusion. Hence, if the task is to track a target, it is not possible to provide any continuity guarantee to maintain the target always inside the FoV. In this case, a multi-camera setup expands the sensing coverage and, through collaboration, it is possible to decide which camera will be the most suitable to track the target, as it moves across different FoVs (handover). Active vision comes into action when the FoVs are disjoint and, in order to synchronize the target hand-off, it becomes necessary to cooperatively control the pan and tilt parameters of each camera.

- *Tracking with coverage*: most surveillance systems are expected to satisfy multiple, possibly conflicting, objectives Ding, Song, Morye, Farrell, and Roy-Chowdhury (2012). For example, in event recognition tasks, it is often necessary to continuously perform area coverage and to capture close-up pictures of events of interest (if any). This is only possible with a self-organizing multi-camera architecture that actively selects which cameras have to narrow their FoV to focus on the target (zoom control); other cameras must then cooperate to compensate for this lack of information by changing their parameters and modifying their FoVs accordingly. Overall, the system must achieve an optimal balance between area coverage maximization and information gathering. Notably, this task usually involves precise and efficient pan-tilt control algorithms, either for the reconfiguration of the coverage FoVs, or for the active tracking mission.

- *Multi-target tracking*: if the number of targets is large, the level of parallelisation requires a camera network. The optimization process computes a feasible configuration to satisfy the Quality of Sight (QoS) for each target; this might also include distributing tasks among active cameras while meeting network constraints.

Multi-modal smart camera networks As already mentioned, camera networks are useful to cover and monitor wide areas; nonetheless, it is sometimes not practical, in terms of deployment and maintenance costs, to have video cameras that completely cover the entire region of interest. Therefore, an alternative is to deploy sparse camera networks, where cameras are distributed over the most interesting regions and then complemented by additional sensors with different modalities. It has been shown that these methods lead to a more efficient usage of the camera resources, by focusing on the most important parts of the scene: they allow to save power, bandwidth and cost, and they also reduce concerns of privacy Nayak, Gonzalez-Argueta, Song, Roy-Chowdhury, and Tuncel (2008).

2.3.2 Use case: multi-modal grasping

Robots must be equipped with different sensing modalities to be able to operate in unstructured environments Luo, Bimbo, Dahiya, and Liu (2017). In this regard, the process of active perception includes the fusion of different data sources, which allows to generate more meaningful, higher-level representations of the world. Furthermore, redundant information is useful to reduce the uncertainty on the quantities to be measured or estimated; at the same time, it increases the reliability of the system in case of failure. Combining different sensing modalities provides also complementary information and has a synergistic effect, especially for those environmental features that can not be perceived using each sensor separately.

Recent studies have shown that grasping and object manipulation are among the principal robotics applications that benefit from the aggregation of visual cues with other sensing modalities, ranging from radio and audio signals, to pressure, tactile and proximity sensors Milighetti, Emter, Kuntze, Bechler, and Kroschel (2006); Papanikolopoulos (1996); Xue, Wang, Ma, Liu, Pan, and Han (2020). Indeed, the information richness of visual sensors may be impaired by occlusions, FoV directionality, and illumination changing Mavrinac and Chen (2013). In contrast, radio and acoustic signals are less energy-harvesting and provide safer privacy guarantees; however, their noise level requires robust signal processing techniques when working in adverse environments Milighetti et al. (2006). In particular, most radio and acoustic localization systems have omnidirectional domains , so that it is possible to sense the full scene without scanning. Nonetheless, these systems suffer from relatively low accuracy and have severe problems with background noise. On the other side, proximity sensors are specifically used to detect the presence of objects of interests, or obstacles along the robot path. Finally, the sense of touch in robots can play a fundamental role in enhancing perceptual capabilities; specifically, vision is complemented by tactile perception when an agent must discriminate different material properties through physical interaction Luo et al. (2017).

Human-to-robot handover Robots operating in dynamic and unstructured environments are sometimes required to exhibit advanced forms of interaction with humans. Human-to-robot handovers are important for daily household activities, and for industrial manufacturing. The objective is to deliver an object from the robot to the human, or viceversa, with strict safety and precision constraints. Thus, human-to-robot handovers of unknown objects require accurate estimation of hand poses and object properties, such as shape and weight Pang et al. (2021). To this aim, existing perception algorithms use multi-modal data (e.g., audio-visual perception Ishikawa, Nagao, Hachiuma, and Saito (2021); Liu, Feng, Lan, and Chan (2021b)), which provide the best performance in terms of object properties estimation Iashin, Palermo, Solak, and Coppola (2020) and overall task safeness Pang et al. (2021).

2.4 Open research questions and Contributions

Active sensing was first introduced in 1985 by the authors of Bajsy (1985); since then, the AS framework has received a considerable attention, both from academic and industrial communities. Nonetheless, a number of open research questions still need to be addressed. In the following, the most relevant for this dissertation are reported.

1. In Probabilistic Active Sensing the actions are taken according to the environmental model learnt from data. These actions affect the model itself; thus, it is fundamental to understand how to rigorously formulate and conceptualize the PAS framework.

2. Once the PAS paradigm has been formalized, it is necessary to define the methodological tools that allow to efficiently generate and maintain the probabilistic map over time.

3. In Multi-Sensor PAS it is fundamental to rely on efficient, accurate and robust techniques that collect and aggregate different information sources to generate the belief map. Moreover, decentralized multi-platform architectures require to distribute the information over the agents, in order to have a global synchronization on the same version of the map. Finally, multi-platform coordination poses severe challenges regarding system integration and real-time data communication. In this context, it is of paramount importance some research on how to apply data fusion and knowledge encryption techniques within the PAS framework. This problem becomes particularly interesting in multi-platform and multi-modal scenarios, where each sensing modality and each sensing platform has its own characteristics.

4. Some of the fundamental issues of the active sensing research concern the next movement and the realization of the motion of the sensor. According to previous studies Kreucher, Kastella, and Hero Iii (2005b), the computational complexity of most sensor planning problems is NP-hard; thus, it is necessary to design control algorithms that guarantee sub-optimal solutions within predefined external specifications.

5. The majority of AS applications deal with harsh settings where cluttering, environmental noise, and time-varying conditions complicate the effectiveness of control and estimation tasks. Thus, AS algorithms must be robust to real-life conditions.

6. The research community has to show that it is possible to build robotic systems with active perception capabilities; namely, robots that are adaptive to varied environments and tasks, that can deal with unexpected, novel situations and that can learn from experience. Finally, their performance must be robust, reliable and safe Bajcsy et al. (2018).

As already mentioned in Chapt. 1, the focus of this Thesis is on the active sensing paradigm applied to multi-sensor architectures. In particular, the objective is to address the research challenges enlisted above, by providing the following contributions.

1. The major comprehensive contribution that has been devised in the entire dissertation rests upon the combination and exploitation of distinct disciplines with the purpose to consolidate both the theoretical and practical frameworks in which multi-sensor PAS plays a crucial role.

2. The literature on MS-PAS still misses a unified methodological framework; hence, this Thesis proposes a general and modular scheme to incorporate multi-sensory data into any probabilistic active sensing problem. The proposed framework serves as methodological reference for future research and can be applied to a number of different scenarios, including position estimation, video surveillance, and IoT applications.

3. The third contribution is the validation of the proposed scheme in four different real-world case studies: localization, tracking, video surveillance, and parking estimation for connected vehicles.

4. Another contribution is the development of control techniques for data collection systems, designed to fit real-world requirements and conditions, such as computational efficiency, robustness to noise, and accuracy. Throughout this dissertation, a particular attention is devoted the management of limited resource capabilities, and to the presence of measurement uncertainties and unstructured environments Calli et al. (2018), which are handled through Bayesian reasoning and learning techniques.

5. Multi-modal PAS schemes are a promising trend in intelligent perception systems, thanks to the increasingly prevalence of heterogeneity in robotic systems Lewis et al. (2014). Along this line, a further contribution stems on the exploitation of multi-modal sensor fusion techniques, used to refine incoming measurements, learn the sensors observation models in self-supervised manners, and keep under control the platform resource usages.

3

Active Position Estimation: methodological framework

This Chapter proposes a reference methodological framework for PAS and MS-PAS solutions, within the context of Active Position Estimation (APE), which is the task of localizing one or more targets using one or more sensing platforms. The fundamental elements and the taxonomy of any AS scheme are introduced by exploring the state of the art on autonomous perception systems for localization tasks.

Contents

3.1 Active Position Estimation

Active Position Estimation (APE) is a specific active sensing task where one Lathuilière et al. (2019) or more Pack, DeLima, Toussaint, and York (2009) platforms monitor and localize one Shahidian and Soltanizadeh (2017) or more Chung and Furukawa (2009) static Fink and Kumar (2010) or dynamic Ryan and Hedrick (2010) targets (humans Murphy et al. (2019), animals Cliff et al. (2018), or robots Koohifar et al. (2018)), which may be active Twigg, Fink, Paul, and Sadler (2012) or passive Liu and Hedrick (2017). Actuated sensing platforms drive the perception process towards observations that are informative optimal Shahidian and Soltanizadeh (2017) or efficient in terms of task completion Lathuilière et al. (2019). In the APE framework, information-seeking strategies Aoki, Bagchi, Mandal, and Boers (2011a) adopt explorative policies that consist of making decisions for next actions in order to maximize the uncertainty reduction on the target position; this leads to better estimation performance with respect to passive sensor systems Denzler, Zobel, and Niemann (2003); Queralta et al. (2020). On the contrary, task-driven approaches bring the sensing platform closer to its task execution, such as reaching a goal position within certain time/energy tolerances Kreucher, Hero, and Kastella (2005a).

An APE task can be categorized as Search, Localization and Tracking. *Active Search* (ASe) refers to the problem of making data-collection decisions in order to search for objects of interest (i.e., targets) which have not been detected yet Ghods, Durkin, and Schneider (2020). In *Active Localization* (AL), the purpose is to estimate the location of static targets that have already been detected Fink and Kumar (2010); Hutchinson et al. (2019b); Park and Oh (2020); Shahidian and Soltanizadeh (2017); Twigg et al. (2012); *Active Tracking* (AT) differs from Active Localization in that the targets are dynamic Wang, Sanchez-Matilla, and Cavallaro (2018). Furthermore, in a *Search and Tracking* (SAT) task, Active Search precedes Active Tracking Lathuilière et al. (2019); Liu and Hedrick (2017); Pack et al. (2009). The common challenge of all these categories is to design autonomous and, possibly, cooperative decision making strategies to search and localize the targets within a set of given resource constraints (e.g., energy consumption, mission time, travel distance) Popovic (2019). Fig. 3.1 shows the main elements of a generic APE scheme, namely the targets and the sensing platforms.

This Chapter provides a formal definition of the Active Position Estimation problem, with a comparative analysis on the main sensing modalities (Sec. 3.2) and control schemes (Sec. 3.4) involved. One of the main insights of this Thesis is on autonomous and multi-sensor perception systems; along this line, this Chapter highlights the role of information fusion and multi-agent cooperation in APE tasks (see Sec. 3.3). Even though several works discuss the level of cooperation between sensing robots Khan, Rinner, and Cavallaro (2016); Queralta et al. (2020); Robin and Lacroix (2016); Senanayake, Senthooran, Barca, Chung, Kamruzzaman, and Murshed (2016), the importance of multi-modal fusion algorithms is mentioned only in Queralta et al. (2020). However, the authors of Queralta et al. (2020) specifically focus on search and rescue applications, while this Thesis promotes the development of a general framework that can be applied to the entire APE class. Control theory is another relevant part of autonomous perception;

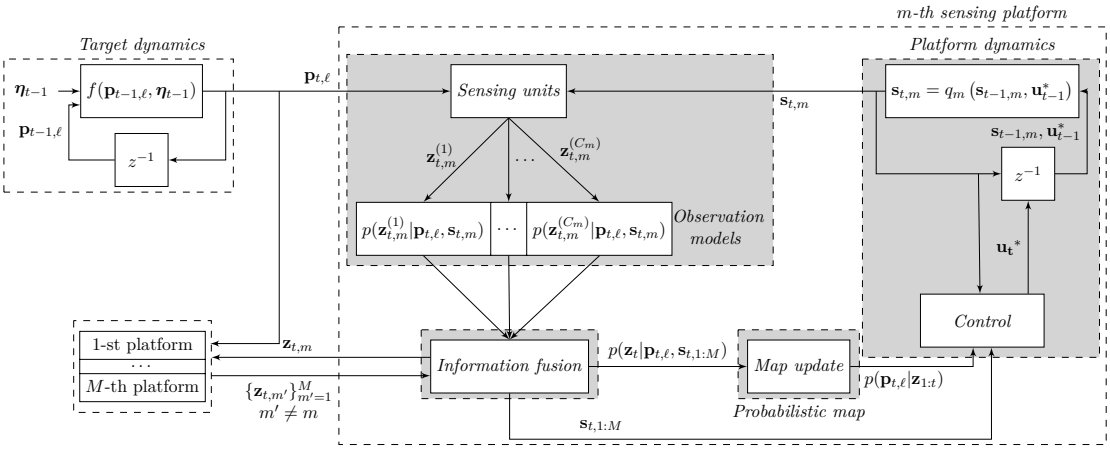

Figure 3.1. Generic APE scheme: different targets properties require different APE algorithms and, therefore, sensing platforms with specific requirements. Perception modules, on-board data management and multi-agent coordination, and platform control are the most important aspects to be considered in the design of any APE mission.

Table 3.1. Comparison of Active Position Estimation methods. Key – SAT: Search and Tracking; AL: Active Localization; AT: Active Tracking; A: Active; P: Passive. N: number of targets; M: number of sensing platforms.

Ref.	Class	Target			M
		motion	type	N	
Liu et al. (2015)	SAT	static	P	1	1
Liu and Hedrick (2017); Ramirez et al. (2014)	SAT	dynamic	P	1	1
Lathuilière et al. (2019)	SAT	dynamic	A	>1	1
Pack et al. (2009)	SAT	dynamic	A	>1	>1
Fink and Kumar (2010); Hutchinson et al. (2019b); Radak, Baulig, Bijak, Schowalter, and Frey (2017) Sun, Xiao, Li, and Cabrera-Mora (2008); Twigg et al. (2012); Varotto and Cenedese (2021c)	AL	static	A	1	1
Dogancay (2012)	AL	static	P	1	>1
Hoffmann and Tomlin (2009); Park and Oh (2020); Shahidian and Soltanizadeh (2017) Vander Hook, Tokekar, and Isler (2015)	AL	static	A	1	>1
Czarnetzki et al. (2010); Denzler et al. (2003)	AT	dynamic	P	1	1
Cliff et al. (2018); Haubner, Schmidt, and Kellermann (2019b); Varotto et al. (2021)	AT	dynamic	A	1	1
Koohifar et al. (2018); Koohifar, Kumbhar, and Guvenc (2016)	AT	dynamic	A	1	>1
Van Nguyen, Chen, Chesser, Rezatofighi, and Ranasinghe (2020)	AT	dynamic	A	>1	1
Meera et al. (2019)	PTS	static	P	>1	1
Katsilieris, Boers, and Driessen (2012)	PTS	dynamic	P	1	1
Bourgault, Furukawa, and Durrant-Whyte (2003)	PTS	dynamic	P	1	>1

hence, the state of the art is reviewed according to the platforms degrees of freedom, the control methods adopted for the platform actuation, and the main criteria that guide APE tasks (see Sec. 3.4).

Tab. 3.1 organizes the reviewed articles according to the Active Position Estimation method applied, the characteristics of the targets, as well as the number of sensing platforms involved.

3.1.1 The targets

Let N be the number of targets to be searched and tracked[1]. The movements of *target* ℓ at time instant t can be described by a stochastic Markovian state transition model Katsilieris

[1]In practice, the number of targets is not always known beforehand (e.g., in search and rescue applications); moreover, it may also be time-varying (e.g., in video surveillance systems). In these cases, the uncertainty of the APE problem increases; hence, ad-hoc and more complex planning algorithms are required Yi, Morelande, Kong, and Yang (2012)

et al. (2012):

$$\begin{aligned} &\mathbf{p}_{t+1,\ell} = f_\ell(\mathbf{p}_{t,\ell}, \boldsymbol{\eta}_{t,\ell});\ \ell = 1, \ldots, N \\ &\text{s.t. } \mathbf{p}_{0,\ell} \sim \mathcal{P}_0 \\ &\quad \mathbf{p}_{t,\ell} \in \Pi,\ t \geq 0 \\ &\quad \boldsymbol{\eta}_{t,\ell} \sim \Xi_\ell \end{aligned} \tag{3.1}$$

where $\mathbf{p}_{t,\ell}$ is the target position, $\mathbf{p}_{0,\ell}$ is the initial condition, realization of a random vector with probability density $\mathcal{P}_0$, $\Pi \subset \mathbb{R}^d$ ($d \leq 3$) is the workspace (i.e. the environment where targets move); and $\boldsymbol{\eta}_{t,\ell}$ is the process noise, realization of a random vector with probability density Ξ_ℓ and support in $\mathbb{R}^d$. More specifically, the bias of Ξ_ℓ (i.e., $\mathbb{E}[\boldsymbol{\eta}_\ell]$) models the target driving input, while the variance (i.e., $var\,[\boldsymbol{\eta}_\ell]$) accounts for the model inaccuracies and the state uncertainty. The process model is described by $p(\mathbf{p}_{t+1,\ell}|\mathbf{p}_{t,\ell})$, which is defined by $f_\ell : \Pi \times \mathbb{R}^d \to \Pi$, a (possibly) non-linear function that describes the dynamics of target ℓ Katsilieris et al. (2012). Typically, (3.1) can be simplified by assuming the same dynamic behavior and process noise distribution for all targets Wong, Bourgault, and Furukawa (2005), namely

$$f_\ell(\cdot) = f(\cdot),\ \Xi_\ell = \Xi;\ \ell = 1, \ldots, N. \tag{3.2}$$

A typical choice for Ξ is the Normal distribution Katsilieris et al. (2012),

$$\boldsymbol{\eta} \sim \mathcal{N}\left(\boldsymbol{\mu}_{\boldsymbol{\eta}}, \boldsymbol{\Sigma}_{\boldsymbol{\eta}}\right) \tag{3.3}$$

where $\boldsymbol{\mu}_{\boldsymbol{\eta}}$ and $\boldsymbol{\Sigma}_{\boldsymbol{\eta}}$ are the bias and variance of Ξ, respectively.

In the APE literature, targets are often classified as either *active* or *passive* Deak, Curran, and Condell (2012).

Active targets We say that a target is active when it occasionally Koohifar et al. (2018) or continuously Shahidian and Soltanizadeh (2017) releases information on their presence Khan et al. (2016) (e.g., through radio ID Cantón Paterna, Calveras Auge, Paradells Aspas, and Perez Bullones (2017b)) or on their position (e.g., GPS coordinates Rodríguez et al. (2012)). To this aim, target and (robotic) sensing platforms may establish some form of direct communication, for instance via radio Cliff et al. (2018) or acoustic Wang et al. (2018) devices. Alternatively, active targets propagate information on their position via radio broadcasting Fink and Kumar (2010) or by spreading specific materials in the surroundings (e.g. hazardous chemicals, fluids, or gaseous particles Hutchinson et al. (2019b)). The presence of active targets is usually captured by robots equipped with ad-hoc sensing capabilities (i.e., radio receivers Van Nguyen et al. (2020), microphone arrays Haubner et al. (2019b), gas detectors Park and Oh (2020)); once the target is perceived, its position is inferred through the extraction of specific features from the collected data (e.g., the signal strength in radio communication Li, Geng, Ye, Xu, Lin, and Pang (2018b)) Active targets often carry electronic devices that send information to the APE system (e.g., portable WiFi transmitters Carpin, Rosati, Khan, and Rimoldi (2015)). Despite this, it is important to remark that a target is not required to carry an emitting device to be active: it can be a device-free speaker Lathuilière et al. (2019) or a radiation source in critical environments Gao, Wang, Zhu, Huang, Wu, and

Du (2018). In addition, it should be noticed that the definition of active target does not include any awareness or collaboration guarantee from the target to the sensing platform.

Passive targets As opposed to active targets, passive ones neither send their location information to robots nor hide from spread substances Deak et al. (2012); Khan et al. (2016); thus, they can only passively reflect the electromagnetic waves emitted from external sources (e.g., the platform transmitters Rosić, Simić, and Pejović (2020)). For this reason, passive targets are usually detected by sensors like cameras Ryan and Hedrick (2010) or RADARs White, Williams, and Hoffensetz (2008), via video processing techniques Redmon, Divvala, Girshick, and Farhadi (2016); Viola and Jones (2001) or through the variation of a measured signal (e.g., Differential Air Pressure Deak et al. (2012)).

3.1.2 The sensing platforms

As discussed in Chapt. 2, active sensing can only be accomplished with one or more sensing platforms with self-regulation capabilities. Let M be the number of sensing platforms (robots) whose state is $\mathbf{s}_{t,m} \in \mathcal{S}_m$, where $\mathcal{S}_m$ is the state space of the m-th robot[2]. The state space represents the platform degrees of freedom. It is *Euclidean*, $\mathbb{R}^d$, when the platform translates in a one Noori, Renzaglia, Vander Hook, and Isler (2016) (i.e. $d = 1$), two Park and Oh (2020) (i.e. $d = 2$) or three Ramirez et al. (2014) (i.e. $d = 3$) dimensional space; in this case,

$$\mathbf{s}_{t,m} = \mathbf{c}_{t,m} \in \mathbb{R}^d, \tag{3.4}$$

where $\mathbf{c}_{t,m}$ denotes the platform position. The state space is a *rotation manifold* when the platform can only rotate, either in one direction Katsilieris et al. (2012) (pan/tilt only platform) or in two Lathuilière et al. (2019) directions (pan-tilt platform); it follows that

$$\mathbf{s}_{t,m} = \begin{bmatrix} \alpha & \beta \end{bmatrix}^\top \in \mathbb{S}^2, \tag{3.5}$$

where α and β are the pan and tilt angles, respectively; $\mathbb{S}^2$ is the sphere in $\mathbb{R}^3$. Finally, the state space is *Riemannian* Tron, Vidal, and Terzis (2008) when the state embeds both translations and rotations Pérez-Carabaza et al. (2019a). There are also high-dimensional robotic systems, like eye-in-hand manipulators Radmard and Croft (2017), where the state represents the joint parameters vector and encodes a large amount of degrees of freedom.

In APE tasks, the detection event is the process of collecting informative measurements on a specific target SanMiguel and Cavallaro (2017). To this aim, each sensing platform is endowed with detection capabilities (e.g., object detectors in camera systems Redmon et al. (2016), or gas detectors in source term estimation[3] Park and Oh (2020)). Hence, we can characterize the detection event of target ℓ from platform m, at time t, as a

[2]The same considerations made in Sec. 3.1.1, regarding the number of targets, can be repeated for the number of platforms, since M is not always known beforehand and it may change over time.

[3]Source term estimation (STE): estimation of location, emission rate and other variables needed to describe, through dispersion models, the spread of the hazardous materials (e.g., gaseous pollutants) into the ambient environment Hutchinson et al. (2019b).

Bernoulli random variable $D_t^{m,\ell} \in \{0,1\}$ with success probability Varotto et al. (2021)

$$p(D_t^{m,\ell} = 1 | \mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}) = P_d(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}). \tag{3.6}$$

The target detectability depends on both the target position and the sensor state.

3.1.3 Active Search and Probabilistic Target Search

Active Search of target ℓ is accomplished before the target detection event has happened yet in *any* platform, that is

$$D_k^{m,\ell} = 0;\ k = 0, 1 \ldots, t;\ \forall m \in [1, M]. \tag{3.7}$$

In other words, Active Search (ASe) starts with no target position information available and may employ a deterministic Mohamed and El-Hadidy (2013), random Cao, Tan, Li, Gu, and Wang (2006); Gage (1994); Isler, Kannan, and Khanna (2005) or probabilistic Katsilieris et al. (2012) search. Deterministic search follows predetermined patterns, namely a systematic traversal of the whole environment along a predefined direction[4]. Examples of deterministic patterns are the spiral search Bernardini et al. (2017); Fricke, Hecker, Griego, Tran, and Moses (2016) and the lawn-mower pattern Bopardikar, Bullo, and Hespanha (2007); Hutchinson, Liu, and Chen (2019a); Ousingsawat and Earl (2007). On the contrary, random search is based on a completely random choice of the next platform actions. Randomization is more powerful than deterministic strategies, since it allows to solve problems that are not solvable by deterministic algorithms Isler et al. (2005); in spite of this, random search is not always efficient and usually serves as a lower bound when the objective is to minimize the expected time until target detection Chung et al. (2011).

Probabilistic Target Search (PTS) Chung et al. (2011); Furukawa, Mak, Durrant-Whyte, and Madhavan (2012); Katsilieris et al. (2012) accounts for target motion and sensing uncertainties Katsilieris et al. (2012). In doing so, PTS typically formulates the search task within a Bayesian probabilistic framework (e.g., Recursive Bayesian Estimation Arulampalam, Maskell, Gordon, and Clapp (2002); Hue, Le Cadre, and Pérez (2002); Smith (2013), or Bayesian Optimization Brochu, Cora, and De Freitas (2010); Frazier (2018); Meera et al. (2019); Snoek, Larochelle, and Adams (2012)). In this way, it is possible to encode and keep updated the knowledge about potential target locations as a probability distribution, also referred to as belief or probabilistic map Liu et al. (2015); this is done by treating no-detection observations (i.e., measurements with no information on target position) as negative likelihood (see Sec. 3.2.1). PTS methods are based on decision-making modules that rely on the optimization of a search objective, built upon the belief map Chung et al. (2011); some examples are the probability of detection Katsilieris et al. (2012), time to detection Pérez-Carabaza et al. (2019a), information gain Cliff et al. (2018), or distance to the target Chung and Burdick (2011); Hasanzade, Herekoğlu, Yeniçeri, Koyuncu, and İnalhan (2018); Kreucher et al. (2005a). The reader is addressed to Sec. 3.4.4 for a comparative analysis on the main APE criteria. PTS has its origins

[4]The reader may refer to Chung, Hollinger, and Isler (2011); Robin and Lacroix (2016) for a broader discussion on all possible active search classes.

in the research done by the USA Anti-Submarine Warfare Operations Research Group (ASWORG) during World War II Koopman (1946) and it is currently widely used in both research and industrial settings. The main motivation stems on the fact that probabilistic approaches are suitable to real-world scenarios, especially when resource consumption (energy and time) is critical Pérez-Carabaza et al. (2019a); this is due to the use of stochastic target motion models Li and Jilkov (2003), combined with the capability of representing realistic perception uncertainties Popovic (2019); Radmard and Croft (2017). Moreover, Bayesian-based control has been proven to be faster, more robust, and more effective compared to traditional control methods that rely on direct (or filtered) sensor feedback Bourne et al. (2019). In addition, relying on probabilistic strategies is necessary when dealing with large-scale search spaces; indeed, the inherent NP-hard complexity of the searching task Hollinger and Singh (2008); Pérez-Carabaza et al. (2019a); Pérez-Carabaza, Scherer, Rinner, López-Orozco, and Besada-Portas (2019b); Popovic (2019); Trummel and Weisinger (1986) induces deterministic strategies to be computationally intractable for very large-scale environments Shubina and Tsotsos (2010), due to their polynomial Park, Lee, and Chwa (2001) or exponential Hollinger, Singh, Djugash, and Kehagias (2009); Schlotfeldt, Atanasov, and Pappas (2019) scalability with the size of the environment. Finally, compared to deterministic path planning, PTS strategies are more capable to adjust the platform behavior in response to new information Ghods et al. (2020); Popovic (2019); this higher adaptivity properties are proven to attain more accurate target position estimates and lower search times (i.e., higher efficiency) Bourne et al. (2019); Hutchinson et al. (2019b).

Applications Active Search is widely used in the field of *search and rescue* missions, where aerial Du et al. (2019), ground Rodríguez, Al-Kaff, Madridano, Martín, and de la Escalera (2020), or underwater Stone, Keller, Kratzke, and Strumpfer (2011) robots, embedded with visual Lin and Goodrich (2009), acoustic Sibanyoni, Ramotsoela, Silva, and Hancke (2018) or radio Murphy et al. (2019) devices, are employed to detect and localize people in distress or imminent danger.

3.1.4 Active Localization and Tracking

Active Tracking starts the moment k_0 in which condition (3.7) breaks, namely

$$k_0 = min\left\{k \in \mathbb{N} \text{ s.t. } D_k^{m,\ell} = 1\right\}. \tag{3.8}$$

Active Tracking is referred as Active Localization when targets are static, that is

$$\mathbf{p}_{t,\ell} = \mathbf{p}_\ell, \ \forall \ell = 1, \ldots, N. \tag{3.9}$$

Sec. 3.2 provides an overview of the main sensing modalities adopted in AL and AT, including vision Ryan and Hedrick (2010), RADAR White et al. (2008) and LiDAR Kumar, Lee, Hwang, Park, Youn, and Kwon (2020), audio Haubner et al. (2019b), and radio signals Shahidian and Soltanizadeh (2017). As discussed in Sec. 3.3, AL and AT use measurements provided by single Haubner et al. (2019b) or multiple Koohifar et al. (2018) on-board sensors, to control the platform while estimating the target position. Measurements may be aggregated through sensor fusion, which can be uni-modal Koohifar

et al. (2018) or multi-modal Lathuilière et al. (2019), cooperative Park and Oh (2020) or non-cooperative Shahidian and Soltanizadeh (2017). Similarly to the PTS case, in probabilistic AL and AT the sensing platform selects the information to extract from the belief map based on a predefined criterion (e.g. reaching a goal position Hasanzade et al. (2018), maximizing the information gain Cliff et al. (2018); Hutchinson et al. (2019b)); this information is then injected into the controller, in order to generate the next sensor action.

Applications Accidental releases of liquid or gaseous pollutants (i.e., hazardous materials - *HAZMAT*) into the ambient environment can occur at industrial facilities such as petroleum refineries, petrochemical plants, natural gas processing plants, oil and gas transportation pipelines, and chemical plants. HAZMAT pose both an immediate and chronic risk to environment and human health; hence, a prompt and accurate STE is important to predict the spread of hazardous material and enable appropriate mitigation steps. In fact, the ultimate goal of autonomous STE is to drive the robot to the location of an emitting source without endangering the personnel. Therefore, the target in STE problems is represented by the HAZMAT source (usually static), while the platforms are typically unmanned ground or aerial vehicles, equipped with specific HAZMAT detectors (e.g, atmospheric sensors Hutchinson et al. (2019b)). According to the concentration of the HAZMAT, it is possible to retrieve the source position Park and Oh (2020) by leveraging on specific dispersion models Hutchinson et al. (2017). In fact, the dispersion rate and density of emitted substances provides information on the position of the HAZMAT source Hutchinson et al. (2017).

Wildlife monitoring is essential for understanding animal movement patterns, habitat utilisation, population demographics and poaching incidents. Radio transmitters are usually attached to animals, so that autonomous robotic platforms (e.g., UAVs) embedded with radio receivers can track their movements Cliff et al. (2018). Automated wildlife tracking is necessary because robot systems can access rugged areas that are difficult for humans to traverse; moreover, manually locating radio transmitters with handheld equipment is labor-intensive.

The recent development in aerial mobile robotics[5] has revolutionized several applications that require live and dynamic camera viewpoints, such as *autonomous cinematography* Bonatti, Ho, Wang, Choudhury, and Scherer (2019); Jeon, Shim, and Kim (2020) and *entertainment* Kim, Jeong, Park, Ryu, and Oh (2018b). In films production, flying agents are employed to follow actors motion, moving through an unknown environment at high speeds. At the same time, they are required to map the environment to provide artistic and technical feedback to the producers. In this contexts, to plan smooth, collision-free trajectories while avoiding occlusions, is extremely challenging even for experienced pilots. For this reason, the current trend is to exploit autonomous active tracking modules Bonatti et al. (2019).

Localization and tracking is fundamental in multi-robot cooperative systems Carron et al. (2015); Haugen and Imsland (2015); Nagaty et al. (2015); Rizk, Awad, and Tunstel (2019), where teams of robots work together to complete a mission Aranda, López-Nicolás, Sagüés, and Mezouar (2015); Chen, Zhang, Xin, and Fang (2015); Lissandrini, Michieletto, Antonello, Galvan, Franco, and Cenedese (2019). In *autonomous warehouse* inventory

[5] https://store.dji.com/guides/how-to-use-the-djis-return-to-home

schemes Guérin, Guinand, Brethé, Pelvillain, et al. (2016), UGVs work as carries to bring aerial platforms among rows of racks; then, the UAVs are used to scan or pick goods that are not accessible to the ground robots. From the UGV perspective, the objective is to locate and reach a desired warehouse planar position. On the other hand, the target of the UAV is a specific good, after the take-off, and the UGV carrier during landing. Both tasks must be accomplished in an autonomous manner, relying on APE strategies based on visual, acoustic or radio signals.

Another application of multi-robot cooperation is rendezvous planning for *mobile recharging* stations Mathew et al. (2015); Yu et al. (2018). The task consists of a cooperative replenishment strategy for a team of battery-powered aerial robots. To this aim, one or more mobile ground robots are employed as recharging stations, being equipped with a payload of batteries and automated battery swap systems. The goal is to schedule paths for charging robots such that they rendezvous with the UAVs, as replenishment is needed, during the mission.

3.1.5 Search and Tracking

When Active Search precedes Active Tracking, it is referred as Search and Tracking (SAT). The problem is to continuously search for (known or unknown) targets and to track already located ones, which are two competing tasks Khan et al. (2016). In fact, while tracking allows to reduce the uncertainty of a detected target, the search mode is more explorative and gives the possibility to find previously undetected targets; however, some targets are left unobserved and their location uncertainty grows. For this reason, tracking and searching modes are often mutually exclusive; thus, typical SAT solutions Frew and Elston (2008); How, Fraser, Kulling, Bertuccelli, Toupet, Brunet, Bachrach, and Roy (2009); Lathuilière et al. (2019); Pack et al. (2009) are designed so that each robot does not lock its operation in a given mode, thus facilitating keeping a balance between search and track operations Frew and Elston (2008); Furukawa, Bourgault, Lavis, and Durrant-Whyte (2006). This is often implemented through switching policies that leverage on the number of detected targets Lathuilière et al. (2019), the uncertainty level, the accuracy associated to each tracked target Frew and Elston (2008), and the cost of changing from one mode to the other Pack et al. (2009).

Applications *Collaborative robotics* Vicentini (2021) aims at designing robotic platforms capable of cooperating with human beings. Most applications in this field require to actively search, localize and track people by using on-board sensors. To this aim, novel technologies have been proposed to extract contextual information regarding user behaviors (e.g., health monitoring of elderly people, gesture recognition for smart home applications) Jiang et al. (2018); this data can be used to activate positioning tasks in the context of *assisted living* and *ambient intelligence* Sokullu, Balcı, and Demir (2019). For instance, if a fall detection system Jiang et al. (2018) raises a health monitoring alarm, then a mobile platform might be used to search for the user and visually verify their conditions Lewis and Beard (2016).

3.2 The sensing units

The focus of this Thesis is on probabilistic active perception systems. The main steps involved into the design of any PAS scheme are:

1. modeling of the perception process;
2. in case of multi-sensor PAS, definition of a sensor fusion technique to aggregate data coming from different platforms;
3. definition of a statistical signal processing and filtering technique to generate the probabilistic map from collected (and aggregated) data;
4. definition of the mission objective function, together with a control strategy for the decision-making module that optimizes the cost function.

Within the APE context, this Section covers the first and third aspects; sensor fusion is the topic of Chapt. 3.3, while the platform dynamics and control are treated in Sec. 3.4.

3.2.1 The observation model

Each platform may have one Lathuilière et al. (2019) or more Katsilieris et al. (2012) sensors. In this Section only the single-sensor case is considered, while the multi-sensor setup is discussed in Sec. 3.3.

Each platform sensor is characterized by a *Field of View* (FoV), $\Phi(\mathbf{s}_{t,m})$, that is the portion of space where the target *Probability of Detection* (POD), P_D, is positive Furukawa et al. (2012); SanMiguel and Cavallaro (2017), namely

$$\Phi(\mathbf{s}_{t,m}) = \left\{ \mathbf{x} \in \mathbb{R}^3 \mid 0 < P_D(\mathbf{x}, \mathbf{s}_{t,m}) \leq 1 \right\}. \tag{3.10}$$

As (3.10) highlights, the FoV is a finite solid in $\mathbb{R}^3$, whose shape depends on the POD; this, in turn, is a function of the sensor state and might be non-constant over $\Phi(\mathbf{s}_{t,m})$. In fact, $\Phi(\mathbf{s}_{t,m})$ is related to the sensor physical properties, such as the border effects, intrinsic and extrinsic parameters in cameras Furukawa et al. (2012); Radmard and Croft (2017); SanMiguel and Cavallaro (2017); transmittance in RADARs Skolnik et al. (1980); geometrical configuration of microphone arrays in acoustic systems Farmani, Pedersen, Tan, and Jensen (2015); Martinson, Apker, and Bugajska (2011); Vargas, Brown, and Subr (2018); antenna radiation pattern in radio signals Stoyanova, Kerasiotis, Prayati, and Papadopoulos (2007); Stoyanova et al. (2007); Zanella (2016). The POD is related also to the *range* of a sensor, which determines the largest distance at which a target is detectable (over all possible sensor's configurations); formally, for the sensor of the m-th platform

$$R_m = \max \left\{ \|\mathbf{x}\|_2^2 \ : \ \mathbf{x} \in \mathbb{R}^3; \ P_D(\mathbf{x}, \mathbf{s}_{t,m}) > 0, \ \forall \mathbf{s}_{t,m} \in \mathcal{S}_m \right\} \tag{3.11}$$

By definition, for the m-th platform in position $\mathbf{c}_{t,m}$ at time t, it follows that

$$\|\mathbf{c}_{t,m} - \mathbf{p}_{t,\ell}\|_2^2 > R_m \Rightarrow P_D(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}) = 0, \tag{3.12}$$

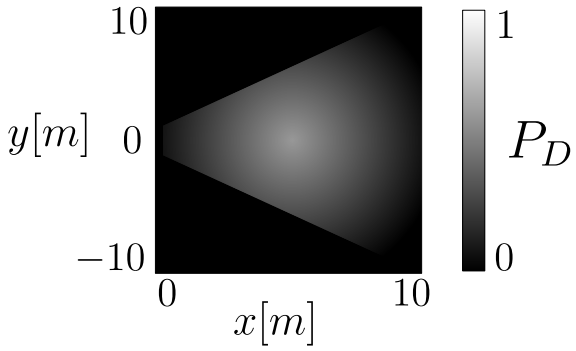

Figure 3.2. Probability of Detection of a visual sensor Furukawa et al. (2012): the camera can not detect any object outside the FoV, while the detection capabilities are maximized at the center of the sensing domain. An object is unlikely to be recognized when too far or too close to the camera, due to scale and resolution issues.

namely, a target can not be detected when outside the sensor range. Fig. 3.2 shows the FOV and the POD of a camera Furukawa et al. (2012). The limited angle of view restricts the target detectability; border effects reduce the POD along the sides of the FOV, while depth effects make the target less detectable at larger distances (the range is $R \approx 20\,\mathrm{m}$). Different sensors are characterized by different FoVs and different ranges (3.10)-(3.11); these, in turn, have a direct impact on the target detection capabilities, according to (3.12).

The behavior of a specific sensor is mathematically characterized through a precise observation model. An *observation model* describes the measurements generation process, relates the target state and the sensor configuration with collected observations, and statistically characterizes the noise associated to each measurement. An accurate and reliable observation model is fundamental to apply suitable statistical filtering techniques Arulampalam et al. (2002) that mitigate sensor nuisance, refine collected data, and quantify the uncertainty associated to an estimation process. Furthermore, observation models are often at the basis of the controller synthesis in active sensing schemes (see Sec. 3.4). Finally, observation models affect the accuracy, the computational load and the convergence rate of most localization techniques Arulampalam et al. (2002); therefore, they are crucial and powerful tools for the design of any probabilistic APE solution. The observation model of sensor m w.r.t. target ℓ can be formalized based on the sensor FoV and the target POD as Furukawa et al. (2012)

$$\mathbf{z}_{t,m}(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}) = \begin{cases} h_m(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}, \mathbf{v}_t), & \text{if } D_t^{m,\ell} = 1 \\ \varnothing, & \text{otherwise} \end{cases} \tag{3.13}$$

where $\varnothing$ is an empty measurement, indicating that the observation contains no information on the target. Otherwise, when the target is detected, a (possibly) non-linear output function

$$h_m : \Pi \times \mathcal{S}_m \times \mathbb{R}^p \to \mathbb{R}^p \tag{3.14}$$

transforms the target position into a p-dimensional measurement vector. The platform can therefore extract information on the target position by leveraging on the set of past and current measurements, namely $\mathbf{z}_{1:t,m}$. The analytical form of $h_m(\cdot)$ and the

dimension of $\mathbf{z}_{t,m}$ depend on the specific sensor type; for instance, $\mathbf{z}_{t,m}$ for a visual sensor is the projection of the target position onto the camera image plane, according to the perspective projection equations Aghajan and Cavallaro (2009). The output function depends also on the (current) sensor state, since different sensing configurations have different effects on the observation process. An example is given by anisotropic radio receivers, where the communication quality with the target is strongly affected by the relative position and orientation of the sensor with respect to the target Zafari et al. (2019); alternatively, consider the effect of camera intrinsic and extrinsic parameters in visual observations Aghajan and Cavallaro (2009); SanMiguel and Cavallaro (2017). Measurements are also affected by the noise component $\mathbf{v}_t$, realization of a random vector whose probability density has support in $\mathbb{R}^p$. The noise components induces a stochastic characterization of the observation model; hence, we can equivalently reformulate (3.13) in a probabilistic manner through the *likelihood function*

$$p(\mathbf{z}_{t,m}, D_t^{m,l}=1|\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}) = p(\mathbf{z}_{t,m}|\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}) P_D(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}). \tag{3.15}$$

The likelihood represents the probability of gathering the measurement $\mathbf{z}_{t,m}$, given the platform state and the target position[6]; in particular, $P_D(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m})$ refers to the target detectability, while $p(\mathbf{z}_{t,m}|\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m})$ captures the sensor performance, supposed to be independent w.r.t. the target detectability. In case of no-detection observations (i.e., $\mathbf{z}_{t,m} = \varnothing$), (3.15) is referred to as *negative likelihood* Katsilieris et al. (2012); Koch (2007) and it holds

$$p(\mathbf{z}_{t,m} = \varnothing, D_t^{m,l}=0 \mid \mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}) = 1 - P_D(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}). \tag{3.16}$$

3.2.2 Probabilistic map

An APE task is said to be probabilistic when information about the target location is modeled with a probability map, kept updated through observations collected by all M platforms until time t, that is $\mathbf{z}_{1:t,1:M}$. As Fig 3.1 shows, the map is then used to guide the platforms towards the next actions Bourgault et al. (2003).

Most probabilistic APE solutions adopt Bayesian filtering Chen et al. (2003) or Gaussian Process Regression (GPR) Rasmussen (2003) for the map generation and update process. In the framework of Bayesian filtering, the map is represented by the posterior distribution of the target location given current and past measurements, namely $p(\mathbf{p}_{t,\ell}|\mathbf{z}_{1:t,1:M})$. To compute this posterior, Recursive Bayesian Estimation (RBE) if often applied (e.g., Kalman filtering Shahidian and Soltanizadeh (2017), particle filtering Arulampalam et al. (2002); Smith (2013)). The major benefit of RBE is that it allows to efficiently and recursively update the map, as new measurements are collected. In particular, Sequential Monte Carlo approaches Smith (2013) allow to treat non-linear and non-Gaussian likelihoods; however, their particle-based approximation makes it difficult to compute uncertainty descriptors (e.g., posterior entropy) from the particle-based approximation of the posterior density Boers, Driessen, Bagchi, and Mandal (2010). As discussed in Sec. 3.4.4, this is a central issue in APE literature.

[6]For this reason, the observation likelihood is also called *generative model*: each datapoint $\mathbf{z}_{t,m}$ is a realization of a random vector whose probability distribution function is the likelihood in (3.15)

Some APE problems are formulated as environmental mapping tasks, where the target location is inferred from the estimated and reconstructed spatial distribution of a physical quantity. For instance, a target can be localized by estimating the maximum of a source signal strength Carpin et al. (2015); Ghassemi and Chowdhury (2020)) or from occupancy maps Khan, Yanmaz, and Rinner (2014); Meera et al. (2019). In this framework, a common approach is to tesselate the workspace Π into a grid map and apply RBE to each cell of the grid Khan et al. (2014). The main drawback of this method is the poor scalability property in large-scale scenarios or when high spatial resolutions are required.

As an alternative to tessellation or RBE, the target posterior distribution can be estimated directly over the continuous domain Π via GPR. Indeed, Gaussian Processes encode spatial correlation structures in a probabilistic and non-parametric manner; they are also capable of accounting for prior knowledge and data dependencies. Furthermore, GPs provide predictions with uncertainty measures, useful to guide robots towards more uncertain areas. GPs were initially unpopular for online APE planning applications, due to their computational complexity and non-recursive nature, as well as their bad scaling properties with large datasets. Indeed, a critical problem for the long term tasks is the large-scale accumulated data, which might exceed the capacity of on-board computational hardware. To alleviate the computational bottleneck caused by large-scale data accumulated, some methods for reducing the computing burdens of GPs have been recently proposed (see Capt. 7). For example, sparse GPR Ma, Liu, and Sukhatme (2017) and active sampling (see Chapt. 7) dynamically select only the most informative data, while abandoning the samples that are less informatively novel (i.e., it is retained only a subset of data that provides the greatest contribution). Other approaches suggest efficient and online GP model refinement techniques hinging on Bayesian recursive algorithms Popovic (2019).

3.2.3 Perceiving passive targets

Assuming the target to be equipped with ancillary emitting devices is a constraining hypothesis in many real-life applications. For instance, in search and rescue missions and video surveillance the objective is to detect people (or objects) that, in general, are not able to establish any kind of communication with the sensing platforms (i.e., the targets are passive). Sensors for perceiving passive targets include cameras, RADARs, and LiDARs.

Vision-based control, or visual servoing, is widely used for target search and tracking Chung and Furukawa (2009); Furukawa et al. (2012); Pérez-Carabaza et al. (2019a,b); Radmard and Croft (2017); Ramirez et al. (2014); Ryan and Hedrick (2010). The main motivation stems on the affordable cost of cameras and the high information content of visual data (i.e., images); furthermore, various computer vision applications have greatly benefit from the recent advances in signal processing and artificial intelligence (deep learning, in particular), such as object detection and classification, action and activity recognition, and human pose estimation Voulodimos, Doulamis, Doulamis, and Protopapadakis (2018). In spite of this, occlusions and Field of View (FoV) directionality limit the range, applicability and success of camera-only platforms Mavrinac and Chen (2013). Another downside of this technology lies in the fact that adverse weather conditions, e.g.,

fog or rain, or external light sources, can severely affect its output. Moreover, single monocular cameras can not provide accurate distance measurements.

Visual servoing systems use feedback information extracted from a camera to control the motion of a robot Radmard and Croft (2017), or the parameters of the camera itself (e.g., pan, tilt, zoom, exposure) Czarnetzki et al. (2010). In this latter case, visual servoing is referred as *active vision* Bajcsy (1988); Bajcsy et al. (2018); Chen et al. (2011). To avoid inefficient search strategies, characterized by local minima and unstable behaviors Radmard and Croft (2017), POD modeling is a critical aspect when dealing with vision-based systems Meera et al. (2019); SanMiguel and Cavallaro (2017). A realistic camera observation model should consider

- image edge effects: the image plane has poor detection capability on the edges, due to distortions caused by the camera sensor Chung and Furukawa (2009);
- detection depth range: distant objects are seen at low resolution and their detection probability is low Furukawa et al. (2012); Ghods et al. (2020); Meera et al. (2019);
- scale effects: objects too close to the camera may not be recognized, due to poor scale-invariance properties of many computer vision algorithms.

To account for these phenomena, a possible detection model is Radmard and Croft (2017)

$$\begin{aligned} P_D(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}) &= \Gamma(d_t) \prod_{i=1}^{2} \frac{1}{c(u_t, \ell_u) c(v_t, \ell_v)} \\ c(x, y) &= 1 + e^{(-1)^i \iota (x-y)} \\ \ell_u &= (-1)^i (L_u/2 - \epsilon), \; \ell_v = (-1)^i (L_v/2 - \epsilon) \\ \Gamma(d_t) &= \left[(1 + e^{\beta_s (d_t - d_s)})(1 + e^{\beta_\ell (d_t - d_\ell)}) \right]^{-1} \end{aligned} \tag{3.17}$$

where L_u and L_v represent the image plane dimension; (u_t, v_t) is the target position projected onto the image plane; d_s is the shortest distance at which the target is detectable; d_t is the target distance; ι, ϵ, β_s and β_ℓ are user-defined parameters. Alternatively, the POD can be modeled as a bi-variate Gaussian distribution over the camera image plane, resembling Fig. 3.2 Liu and Hedrick (2017); Liu et al. (2015). The depth effect can be accounted separately by modeling the POD as piecewise continuous function of the sensor altitude Ramirez et al. (2014). Modeling the POD as in (3.17) or through Gaussian functions has a strong intuitive rationale and it often works in practical applications. Nonetheless, these models incur in a formal issue: the infinite support that characterizes $c(x, y)$ and Gaussian functions is not compatible with the limited sensing domain of cameras. In this regard, literature offers more realistic camera models Pérez-Carabaza et al. (2019a) that handle the limited domain of visual sensors, but they often require extensive user-defined parameters tuning.

LiDAR and *RADAR* are remote sensing devices used for detection, tracking, and imaging of various objects, especially in agriculture and advanced driving assistance systems (ADAS) Raj, Hashim, Huddin, Ibrahim, and Hussain (2020). They both use a transmitter to emit a wave (light for LiDAR, radio for RADAR) and, by measuring the

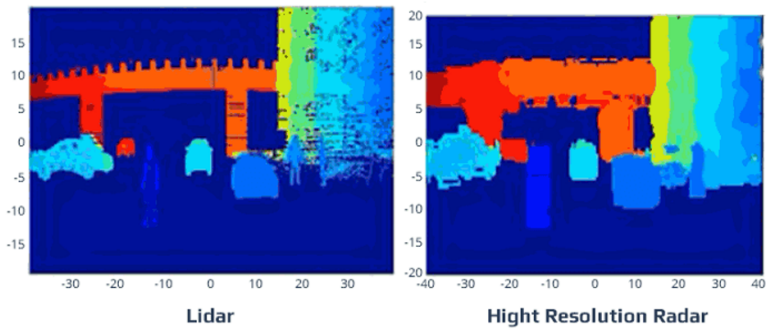

Figure 3.3. A comparison of LiDAR and RADAR accuracy: the image formed through a LiDAR system is much more detailed than the one formed through RADAR.

time it takes the wave to hit an object and bounce back to the receiver, they can measure the distance to the object. Both technologies are highly effective at night, but RADAR scanners Dogancay (2012); Katsilieris et al. (2012) are cheaper, have longer detection ranges (up to 200 m) and allow for measurements with higher reliability and robustness. In particular, radio waves have low attenuation, which allows them to travel minimally disturbed even in poor weather and without being affected by light conditions. However, the main problem of RADAR technologies is the high sensitivity to environmental disturbances, interferences from external sources, and cluttering, as most radio-based detection systems; moreover, the SNR is strongly related to the target distance and to its scattering coefficient (that, in turn, is related to the material reflectivity) Skolnik et al. (1980). Finally, RADARs can not distinguish different types of objects.

On the other side, light signals work in the nanometer range, making LiDAR systems much more accurate and precise than RADAR systems (see Fig. 3.3). A lower wavelength also means that a LiDAR system can detect smaller objects. High resolution and real-time 3D representation (point cloud) of the environment is helpful to obtain shape and distance to surrounding objects, thus facilitating object detection and classification. The downside of LiDAR systems is that they use light waves as a medium, which strongly affects the measurements process. For example, LiDAR systems do not perform as well as RADARs in bad weather.

3.2.4 Perceiving active targets

Active targets propagate in the environment position-related information (e.g., substances, sound, radio frequencies) from which the sensing robots can perceive their presence. Most of the works concerned with active targets typically suppose the target always inside the sensing range; therefore, they neglect the search phase and assume already detected targets Koohifar et al. (2018). Thus, they often differentiate in the way target-informative data are used to perform efficient Active Localization and Tracking tasks.

Concentration sensors such as gas Park and Oh (2020) and atmospheric Hutchinson et al. (2019b) particles detectors are embedded in autonomous sensing vehicles employed in STE tasks. Traditionally, a network of static sensors on the ground are used to estimate the source term; however, it is infeasible to cover all regions of importance with

sensors dense enough to determine the source before it has spread significantly. Hence, recently mobile unmanned sensors have been equipped for STE. Mobile sensors provide the additional ability to perform boundary tracking of the contaminant and source seeking to aid in the emergency response. Nevertheless, using mobile sensors for STE introduces an additional area of research, concerning how to optimally move each sensor in order to produce the best estimate of source parameters within the minimum amount of time, or effort. The solution requires an interdisciplinary fusion of a number of robotics research areas such as autonomous search, multiple robot cooperation, informative path planning and control.

Acoustic signal-based localization techniques often leverage on the energy of the incoming signal, or on the time-related information (i.e., time stamps) contained in modulated acoustic signals; from this information, the microphones array mounted on the robot can estimate the transmitter-receiver relative distance or bearing Cobos, Antonacci, Alexandridis, Mouchtaris, and Lee (2017). MUSIC Schmidt (1986) or SRP-PHAT DiBiase (2000) algorithms are typically used to estimate the relative Angle of Arrival (AoA) between the robot and the source. Alternatively, AoAs can also be measured sequentially, for instance by a moving robot equipped with microphone arrays Haubner et al. (2019b). On the other hand, inferring distance information from microphone data is generally difficult in reverberant rooms and often relies on labeled training data Brendel and Kellermann (2018). In general, audio signals enable high localization accuracy Haubner et al. (2019b) and are characterized by wider FoVs with respect to visual ones Lathuilière et al. (2019); however, this technology is affected by critical pitfalls, like sound pollution Wang et al. (2019), short reception ranges ($\approx 10\,\mathrm{m}$), extra hardware requirements and, consequently, the need of intensive pre-deployment phases with accurate microphones synchronization. Moreover, acoustic signals are subject to distortions caused by environmental changes, object obstruction, signal diffraction, as well as the ego-noise generated by motors and propellers of the sensing robot itself Wang et al. (2019).

Radio-frequency (RF) signals Zafari et al. (2019); Zanella (2016) represent an alternative to acoustics. RF signals guarantee higher reception ranges ($\sim 100\,\mathrm{m}$) Zanella (2016)[7], low energy consumption Siekkinen, Hiienkari, Nurminen, and Nieminen (2012) and they often come as parasitic in many real-life scenarios, since most of the current smart phones, laptops and other portable user devices are WiFi or Bluetooth enabled. For these reasons, RF-based active localization has been largely used in literature Fink and Kumar (2010); Koohifar et al. (2018, 2016); Pack et al. (2009); Shahidian and Soltanizadeh (2017); Twigg et al. (2012) and in commercial applications[8], especially when the objective is to reduce the deployment, economic and energy impact on the overall setup. In spite of this, latencies Sui, Zhou, Liu, Ma, Pei, Zhao, Li, and Moscibroda (2016); Tosi, Taffoni, Santacatterina, Sannino, and Formica (2017) and environmental interference (e.g., cluttering and multi-path distortions) generate strong signal instabilities and limit the accuracy of RF-based solutions Zafari et al. (2019); Zanella (2016). To increase the SNR (often below 10 dB Qin, Dai, and Mitchell (2013)), complex pre-processing algorithms have been suggested Li et al. (2018b); Radak et al. (2017); Zafari, Papapanagiotou,

[7] https://www.blueupbeacons.com
[8] https://quuppa.com

and Hacker (2018). These, however, are computationally demanding and may need dedicated and energy-harvesting hardware, which collides with the typical use of RF architectures. At the same time, latency effects can impose significant challenges to real-time localization; the main cause of latency resides in the adopted communication protocols, which tune collision phenomena and scanning intervals to achieve balance between energy and accuracy Kindt, Saur, Balszun, and Chakraborty (2017); Rondón et al. (2019); Treurniet, Sarkar, Prasad, and De Boer (2015).

Similarly to acoustics, Angle of Arrival (AoA) Xiong and Jamieson (2013) can be estimated also between radio transmitters and receivers, using antennae arrays at the receiver side. To do this, the time difference of arrival (TDOA) at individual elements of the antennae arrays is computed. To estimate the target (transmitter) position from AoA samples, at least two receivers are required. Then, triangulation or Kalman filter can be applied Pack et al. (2009). The former is a pure geometric approach whose accuracy is limited by the number of sensor platforms and it does not improve with time. On the contrary, Kalman filter is capable of gradually increasing the localization accuracy in time, independently from the number of sensor platforms used Pack et al. (2009). Although AoA can provide accurate estimation when the transmitter-receiver distance is small, it requires complex hardware architectures, as well as decent computation power and energy Zafari et al. (2018). In addition, accuracy deteriorates with increase in the transmitter-receiver distance, where a slight error in the angle of arrival calculation is translated into a large error in the actual location estimation.

To reduce hardware requirements and the computational load, it is possible to exploit side effects of wireless communication. One of the simplest and widely used techniques is to extract the Received Signal Strength Indicator (RSSI) from standard data packet traffic Zanella (2016). The RSSI is the actual signal power strength measured at the receiver side, and, in theory, its value is connected to the distance between transmitter and receiver. More formally, the reference model to describe the relation between the RSSI and the Tx-Rx distance is the so-called *log-distance path-loss model* (PLM) Goldsmith (2005), according to which the received power $r(\cdot)$ at distance d from the transmitter can be expressed (in dBm) as

$$r(d) = \kappa - 10n \log_{10}(d/\delta) + v, \; v \sim \mathcal{N}(0, \sigma^2). \tag{3.18}$$

The zero-mean Gaussian noise v accounts for sensor distortions, model inaccuracies and environmental interference; n is the attenuation gain and κ is the RSSI at the reference distance δ. Given the RSSI-distance relation (3.18), several techniques are applied to estimate target position

- *geometric* approaches leverage on model inversion and trilateration Cantón Paterna, Calveras Augé, Paradells Aspas, and Pérez Bullones (2017a); hence, at least three receivers are required. Being deterministic, their robustness to measurement noise and model inaccuracies is often low. Model inaccuracies are particularly critical, since the PLM is too simple to accurately capture the strong signal instabilities and the environmental sensibility Li et al. (2018b).

- *statistical* strategies better account for both model inaccuracies and measurements noise Coluccia and Ricciato (2014), by treating RSSI observations as random vari-

ables and leveraging on statistical filtering techniques to refine raw datapoints (e.g., Kalman filter Röbesaat, Zhang, Abdelaal, and Theel (2017), particle filter Cenedese, Ortolan, and Bertinato (2010); Patel, Girgensohn, and Biehl (2018), belief function theory Achroufene, Amirat, and Chibani (2018)). Model inversion is not always needed and, under specific conditions, one receiver is sufficient to localize a moving target Varotto et al. (2021).

- in *data-driven* models the RSSI-distance relation is represented by a Deep Neural Network (DNN), whose parameters are optimized during the offline training session Choi, Choi, and Talwar (2019); Li et al. (2018b). The main benefit comes from the theoretical function approximation properties of DNNs Hornik (1991): the degrees of freedom given by DNN hyperparameters (e.g., number of layers, number of nodes per layer) allow to tune the level of complexity of the model; hence, in principle, it is possible to accurately capture the underlying propagation behavior. In practice though, hyperparameters optimization is achieved either via extensive brute-force search or through sub-optimal optimization techniques Wu, Chen, Zhang, Xiong, Lei, and Deng (2019). Another important drawback of deep models is that they do not have any physical consistency (i.e., the predictions ignore the underlying physical laws); this may limit the generalization capabilities Hu, Hu, Verma, and Zhang (2020). In addition, training processes commonly require larger datasets than model-based parametric estimation procedures. Finally, when input noise is high (i.e., low SNR), overfitting issues may arise.

- *model-free* techniques mainly refer to the fingerprinting strategy Lin, Chen, Deng, Hassan, and Fortino (2016); Liu, Darabi, Banerjee, and Liu (2007): the operational space is divided into cells and a database is generated, such that each cell is associated to a specific range of RSSI values. When online RSSI data are collected, matching algorithms are applied to retrieve the target position (at cell resolution). Recently, several works have proven the effectiveness of using Recurrent Neural Networks (RNNs) to guide the fingerprinting system Chen, Zou, Yang, Jiang, and Xie (2019b); Hoang, Yuen, Dong, Lu, Westendorp, and Reddy (2019); Sahar and Han (2018). The main disadvantages of model-free techniques is that they are inherently overfitted, thus they offer poor adaptation capabilities in dynamic and cluttered environments. Furthermore, the database generation process requires cumbersome experimental campaigns, even though distributed procedures have been suggested to tackle this issue Xu, Zhu, and Zhu (2019).

To sum up, RSSI-based localization approaches are simple and cost efficient Radak et al. (2017), but they suffer from poor localization accuracy Konings, Faulkner, Alam, Noble, and Lai (2017). Cluttered scenarios make the log-distance path-loss model (5.9) weakly reliable. Furthermore, multipath fading and indoor noise induce severe RSSI fluctuations, so that the noisy component v hides the informative part on the target distance; in fact, in real-world scenarios $\sigma > 3\,\text{dBm}$ Röbesaat et al. (2017); Zanella (2016). Yet, there are many practical solutions that use RSSI to solve APE tasks (see Tab. 3.2). Moreover, the signal strength can also be used to estimate the relative angle between the receiver and the transmitter (i.e., bearing); however, in this case motorized directional antennas are required Vander Hook et al. (2015); Varotto and Cenedese (2021c).

Table 3.2. Comparison of sensing modalities used for Active Position Estimation tasks.

Ref.	Sensing modality					
	Vision	RADAR	Audio	RF	LiDAR	Others
Lathuilière et al. (2019)	✓		DP-RTF			
Koohifar et al. (2016); Shahidian and Soltanizadeh (2017) Koohifar et al. (2018); Twigg et al. (2012) Fink and Kumar (2010); Radak et al. (2017); Sun et al. (2008); Van Nguyen et al. (2020)				RSSI		
Liu and Hedrick (2017); Liu et al. (2015)						generic sensor with limited sensing domain
Hutchinson et al. (2019b)						atmospheric sensor
Cliff et al. (2018)				- RSSI - AoA		
Park and Oh (2020)						gas sensor
Haubner et al. (2019b)			AoA			
Pack et al. (2009)				AoA		
Katsilieris et al. (2012)		✓				
Czarnetzki et al. (2010); Denzler et al. (2003); Meera et al. (2019); Ramirez et al. (2014)	✓					
Hoffmann and Tomlin (2009)				- RSSI - bearing		magnetic field sensing
Varotto et al. (2021); Varotto and Cenedese (2021c)	✓			RSSI		
Vander Hook et al. (2015)				bearing		
Dogancay (2012)		- AOA - TDOA - Scan-Based				
Bourgault et al. (2003)		✓				
Kumar et al. (2020)	✓				✓	

3.2.5 Discussion

Which sensing modalities to rely on is a critical application-dependent choice. Even though the same APE task can be solved by using different sensing modalities (e.g., RSSI, bearing and magnetic field Hoffmann and Tomlin (2009)), most of the reviewed solutions presented in Tab. 3.2 rely on a single sensing modality; only few of them provide formal or empirical comparative analysis to motivate their choice, or to investigate the main pitfalls of alternative solutions. It follows that literature still misses a comparative analysis on the differences between the various sensors types. In particular, it would be interesting to compare the performance of search and tracking in case of passive and active targets: intuitively, updating the belief map with the negative information provided by passive targets is expected to produce long search times before the target is detected; on the other hand, technologies for active targets (e.g., RF and acoustic signals) are typically sensitive to measurement noise and might struggle in harsh environments.

Regarding RF-based methods, it is still unclear the impact that some structural properties have on APE tasks; some of these are long sampling times Zhuang, Yang, Li, Qi, and El-Sheimy (2016) (usually higher than 100 ms Cantón Paterna et al. (2017b); Radak et al. (2017)), delays and packet losses Rondón et al. (2019), cluttering Stoyanova et al. (2007), orientation of the antenna Zanella (2016), and energy consumption Treurniet et al. (2015). Only few works investigate the influence of these factors on search and tracking Gu and Ren (2015); Stoyanova et al. (2007). Even in this rare cases, the data are often collected from static targets and the analyses are mainly empirical; hence, the results can not be generalized to different working conditions.

3.3 Multi-Sensor Active Position Estimation

Single-sensor strategies can sometimes be inadequate: active targets are perceived at long ranges, but measurement noise represents a critical bottleneck for the localization accuracy; on the other hand, passive targets do not need to use any ancillary equipment,

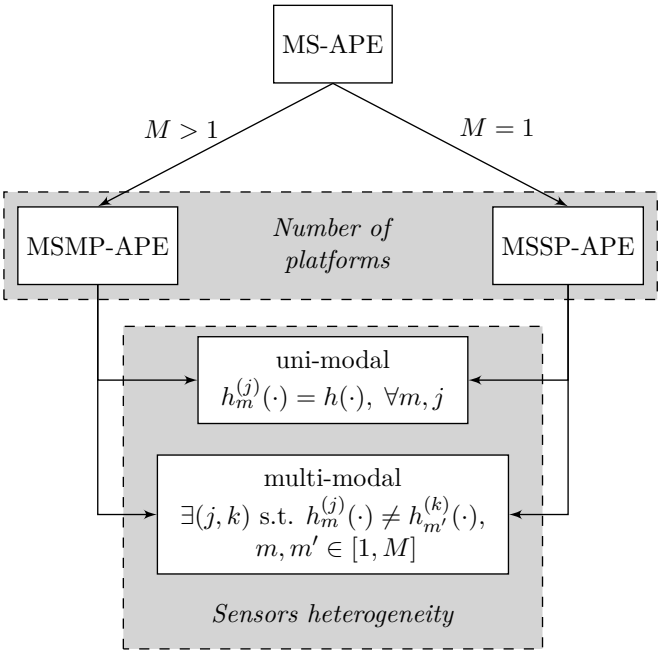

Figure 3.4. MS-APE classification according to the number of platforms and the sensors heterogeneity.

but they are difficult to be detected in highly cluttered environments. For these reasons, some research effort has been devoted to the design of *Multi-Sensor APE* (MS-APE) systems Pack et al. (2009); Park and Oh (2020), where information is collected from multiple sensing units. In MS-APE schemes, the way data is aggregated and manipulated is a fundamental issue; indeed, an efficient information fusion technique can reduce the amount of data traffic, filter noisy measurements, and make better predictions and inferences about the target position. In the last two decades, several information fusion methods, architectures, and models have been proposed. Consequently, the literature on this theme is extremely vast and can not be classified in a unique way Baltrušaitis, Ahuja, and Morency (2018); Khaleghi et al. (2013); Nakamura et al. (2007); Smith and Singh (2006).

In this Section, the works on MS-APE are classified according to the number of platforms, that is, into single Varotto et al. (2021) and multi-platform Koohifar et al. (2018) strategies. These, in turn, are referred to as uni-modal Radmard and Croft (2017) or multi-modal Lathuilière et al. (2019), according to the level of sensors heterogeneity involved (see Fig. 3.4).

3.3.1 Multi-Sensor APE

A multi-sensor APE algorithm exploits multiple sensing units to perform search, localization or tracking tasks. The general formulation considers $M \geq 1$ sensing robots, with the m-th platform endowed with $C_m \geq 1$ sensing units, also referred as channels Jaimes and Sebe (2007). Therefore, a MS-APE system is composed by a total of

$$C_{tot} = \sum_{m=1}^{M} C_m \tag{3.19}$$

sensing units distributed on M spatial locations (i.e., the robot positions).

In this framework, the measurements collected by every channel are aggregated in a unique *multi-sensor observation*

$$\mathbf{z}_t(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,1:M}) = \begin{bmatrix} \mathbf{z}_{t,1}(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,1}) \\ \dots \\ \mathbf{z}_{t,m}(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}) \\ \dots \\ \mathbf{z}_{t,M}(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,M}) \end{bmatrix}, \tag{3.20}$$

where $\mathbf{z}_{t,m}(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m})$ fuses the measurements from the C_m internal channels of the m-th platform, namely

$$\mathbf{z}_{t,m}(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}) = \begin{bmatrix} \mathbf{z}_{t,m}^{(1)}(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}) \\ \dots \\ \mathbf{z}_{t,m}^{(C_m)}(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}) \end{bmatrix}. \tag{3.21}$$

Each channel follows the single-sensor observation model (3.13), that is

$$\mathbf{z}_{t,m}^{(j)}(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}) = \begin{cases} h_m^{(j)}(\mathbf{p}_{t,\ell}, \mathbf{s}_{t,m}, \mathbf{v}_{t,m}^{(j)}), & \text{if } D_t^{m^{(j)},\ell} = 1 \\ \varnothing, & \text{otherwise} \end{cases} \tag{3.22}$$

where $\mathbf{v}_{t,m}^{(j)}$ is the stochastic noise (of given statistics) affecting channel j of platform m, while $D_t^{m^{(j)},\ell}$ is the detection event of target ℓ from channel j of platform m.

Uni-modal and multi-modal MS-APE

A multi-sensor APE can be either uni-modal Radmard and Croft (2017) or multi-modal Lathuilière et al. (2019). A *uni-modal MS-APE* satisfies the condition

$$h_m^{(j)}(\cdot) = h(\cdot),\ j = 1, \dots, C_m,\ m = 1, \dots, M. \tag{3.23}$$

In this case the system is composed only by homogeneous sensors, namely all platforms channels collect measurements of the same sensing modality. Examples are given by swarms of camera-embedded UAVs Pérez-Carabaza et al. (2019a), multi-camera eye-in-hand manipulators Cuevas-Velasquez, Li, Tylecek, Saval-Calvo, and Fisher (2018); Kermorgant and Chaumette (2011), and multi-antenna platforms Jais, Ehkan, Ahmad, Ismail, Sabapathy, and Jusoh (2015); Svečko, Malajner, and Gleich (2015).

On the other hand, a *multi-modal MS-APE* admits the presence of at least two different sensing modalities; formally

$$\exists (j,k) \text{ s.t. } h_m^{(j)}(\cdot) \neq h_{m'}^{(k)}(\cdot),\ m, m' \in [1, M]. \tag{3.24}$$

This is the case of a system composed by heterogeneous sensors. Examples are given by robots equipped with cameras surrounded by a microphone array Lathuilière et al. (2019) (in this case, $m = m'$), or teams of robots where agent m is equipped with different sensors with respect to agent m' (in this case, $m \neq m'$) Doğançay, Hmam, Drake, and Finn (2009).

Multi-Sensor Single-Platform (MSSP) APE

Several APE tasks (e.g., single target tracking Varotto et al. (2021)) do not require multiple sensing robots to be accomplished; hence, the use of a single platform (i.e., $M = 1$) turns out to be more time and cost efficient. As a matter of fact, single-platform solutions do not require any inter-agent synchronization or communication; moreover, efficient multi-agent decision making strategies are often more difficult to design and implement than single-agent counterparts (See Sec. 3.3.1).

As already mentioned, a sensing platform, endowed with multiple sensors on-board, can be either uni-modal or multi-modal. In the former case, it collects the same sensing modality (e.g., visual information Kermorgant and Chaumette (2011)) from all its channels, according to the constraint (3.23). Otherwise, if (3.24) holds, it is referred to as multi-modal. Embedding multiple (potentially different) sources of information on a single platform opens up new perspectives in scene perception. Generally speaking, multi-sensor integration allows to refine single-sensor measurements and obtain more accurate data Rao (2001). More specifically, uni-modal multi-sensor platforms enable parallelisation and specialisation Esterle et al. (2017), while sensors heterogeneity induces inherent robustness and complementarity (i.e., different properties of the environment can be perceived). Notably, the aggregated data (either homogeneous or heterogeneous) allow inferences that are not possible with single-sensor measurements Nakamura et al. (2007).

An example of uni-modal MSSP-APE is given by multi-camera platforms Heng, Choi, Cui, Geppert, Hu, Kuan, Liu, Nguyen, Yeo, Geiger, et al. (2019): the multi-view geometry allows stereoscopic vision from single camera observations; this can then be used for localization and 3D geometric mapping tasks Heng et al. (2019). At the same time, the possibility of acquiring information from different perspectives allows to better handle occlusions, as well as dynamic multi-target scenarios Cuevas-Velasquez et al. (2018); Häne et al. (2017). Besides visual-based positioning, uni-modal MSSP-APE schemes are widely used in audio and radio localization as well, as previously discussed in Sec. 3.2.4. In this case, an array of receivers (either microphones or antennas) is employed to estimate the target position, while this is not feasible with measurements coming from single receivers only Farmani et al. (2015); Xiong and Jamieson (2013). On top of that, the use of multiple receivers has been proven to guarantee higher accuracy and stability in the localization process Kleisouris, Chen, Yang, and Martin (2008); Svečko et al. (2015).

In spite of the advantages introduced by uni-modal MSSP-APE solutions, some applications require a high level of complementarity, which can only be obtained through multi-modal platforms Dore, Cattoni, and Regazzoni (2007a); Lathuilière et al. (2019); Rajasekaran et al. (2020). For instance, RSSI observations are often combined with camera measurements Dore et al. (2007a); Rajasekaran et al. (2020) because visual-based tracking provides high accuracy Aghajan and Cavallaro (2009) and mitigates the inherent ambiguity in RSSI data (see Chapt. 5). Conversely, radio observations are used to detect the target at long ranges and under cluttering conditions, providing a rough position estimate and reducing the typical long search phases of camera-only approaches Ye and Tsotsos (1999). Despite being promising for target search and tracking, the fusion of RF and camera sensor data is still an open problem with sparse documentation Rajasekaran et al. (2020). Finally, it is important to remark that, to achieve multi-modal perception,

it is not always necessary to supply the robot with a specific set of physical sensing units, one for each modality; sometimes it is sufficient to extract different features from the same sensor to get different information cues on the environment. For instance, given an array of radio antennas, it is possible to retrieve both ranging (RSSI) and bearing (AoA) data on the target Dehghan, Farmani, and Moradi (2012).

Applications The main applications that take advantage from multi-sensor integration systems belong to the field of collaborative robotics and human-robot interaction Deyle et al. (2010); Jaimes and Sebe (2007); Lathuilière et al. (2019). Consider, as an example, a robot interacting with surrounding people Lathuilière et al. (2019). In this scenario, the robot should be able to decide where to look in order to understand how many people there are in the environment. To this aim, the combination of audio-visual measurements within a deep reinforcement learning (DRL) framework has been proven to enable the robot to autonomously learn gaze control policies Lathuilière et al. (2019). In this way, the robot decides where to look next, trying to maximize the number of people present in the camera FoV, and favoring people who speak. This application highlights one of the main advantages of audio-visual measurements, that is, the speakers are localized without using any specific device. More specifically, acoustic sources are located via Direct-Path Relative Transfer Function (DP-RTF) Li et al. (2016), which allows to track people outside the camera FoV and, once a person has been visually detected, it is kept within the visual FoV with high accuracy. Another application that benefits from mobile sensor-rich robots is autonomous assistance of motor-impaired patients Deyle et al. (2010). The robot exploits on-board sensors (RFID tags, cameras, and LIDAR) to localize, approach and grasp objects of interest (e.g., a bottle of water, medications, or a mobile phone), successively delivered to the patient Calli et al. (2018). Finally, multi-modality has been found useful also for STE, with atmospheric sensors providing an initial estimate of the source location, further refined through visual observations Monroy, Ruiz-Sarmiento, Moreno, Melendez-Fernandez, Galindo, and Gonzalez-Jimenez (2018).

Probabilistic map in MSSP-APE The stochastic characterization of single channel measurements (3.22) induces $\mathbf{z}_{t,m}$ to be a random vector as well. Hence, the techniques used to manage probabilistic maps in single-sensor APE (see Sec. 3.2.2), can be extended to the multi-sensor (single-platform) case. In particular, once the observation model of each channel is known, it is possible to compute the joint likelihood as Bregonzio, Taj, and Cavallaro (2007)

$$\begin{aligned} p\left(\mathbf{z}_{t,m}|\mathbf{p}_{t,\ell},\mathbf{s}_{t,m}\right) &= p\left(\mathbf{z}_{t,m}^{(1)},\ldots,\mathbf{z}_{t,m}^{(C_m)}|\mathbf{p}_{t,\ell},\mathbf{s}_{t,m}\right) \\ &= \prod_{j=1}^{C_m} p\left(\mathbf{z}_{t,m}^{(j)}|\mathbf{p}_{t,\ell},\mathbf{s}_{t,m}\right), \end{aligned} \tag{3.25}$$

where measurements coming from different sensors are commonly supposed independent. The joint likelihood is then incorporated into a RBE scheme to estimate the target position. Notably, this approach allows to treat uni-modal and multi-modal platforms in the same manner.

Table 3.3. Classification of related works on Multi-Sensor Multi-Platform APE

Ref.	Data type		Coordination		Protocol		
	Uni-modal	Multi-modal	Non-cooperative	Cooperative	Distributed	Decentralized	Centralized
Ghods et al. (2020); Park and Oh (2020)	✓			✓		✓	
Hoffmann and Tomlin (2009); Pack et al. (2009)	✓			✓	✓		
Shahidian and Soltanizadeh (2017)	✓		✓				✓
Koohifar et al. (2018, 2016)	✓			✓			✓
Vander Hook et al. (2015)	✓			✓	✓		
Dogancay (2012)		✓		✓			✓
Bourgault et al. (2003)	✓		✓			✓	

Multi-Sensor Multi-Platform APE

In robotics applications, the presence of multiple sensing platforms (i.e., $M > 1$) increases the redundancy level, making the system less vulnerable to failures and adversarial attacks Queralta et al. (2020). Moreover, swarms of robots have the ability to simultaneously gather information from disjoint locations, while the sensing capability of a single robot is restricted to a limited region (see Sec. 3.2); hence, spatial and temporal coverage problems are reduced when deploying multiple platforms in the monitored environment Chung, Paranjape, Dames, Shen, and Kumar (2018); Nakamura et al. (2007). Finally, multiple platforms are typically employed when the objective includes different concurrent tasks, such as coverage and tracking Mavrommati et al. (2017), or search and tracking Pack et al. (2009).

In the specific case of APE tasks, the cooperation among multiple agents allows to localize multiple targets, to refine the their position estimate, and to accelerate the convergence rate of the localization process. In addition, by jointly performing decision-making strategies, single-robot motion planning can be optimized to account for other agents actions. To this aim, the sensing robots may share their current observations and exchange information on their state (i.e., coordinated APE) Khan et al. (2016).

As outlined in Tab. 3.3, MSMP-APE strategies are here classified according to the type of data to be aggregated (i.e., uni-modal Hoffmann and Tomlin (2009) or multi-modal Doğançay et al. (2009)), the level of robot cooperation (i.e., cooperative Park and Oh (2020) or non-cooperative Shahidian and Soltanizadeh (2017)), and the network communication protocol (i.e., centralized Koohifar et al. (2018), decentralized Park and Oh (2020) or distributed Pack et al. (2009)). The following paragraphs describe the main advantages of coordinated APE and provide an overview on the different communication protocols used in multi-sensor PAS.

Coordination in MSMP-APE Coordination in multi-agent systems implies a minimum level of data exchange. In this work, non-coordinated MSMP-APE solutions (i.e., those where no communication is established among any couple of agents) are not considered; as a matter of fact, any non-coordinated MSMP-APE architecture can be treated as M independent MSSP-APE sub-systems. On the other hand, a coordinated multi-platform APE can be either non-cooperative Shahidian and Soltanizadeh (2017) or cooperative Park and Oh (2020).

- *Non-cooperative* coordination does not involve any negotiation mechanism and decision-makers plan their actions individually, based on their current knowledge of the world (which is represented by the belief map in PAS systems). These kinds of control architectures are commonly used in privacy-preserving applications, where the state of each robot is hidden to the other network agents. Therefore, each agent

trajectory is computed through independent control laws, implemented internally at each sensing platform Bourgault et al. (2003). The only information exchange allowed is on the observations collected by the agents, namely, the measurements can be spread over the network; in this way, all platforms maintain a synchronized global belief map over time, which serves as common knowledge on which local decisions are made.

- *Cooperative* techniques require each agent to take decisions aware of other agents state Park and Oh (2020), usually, with the aim of achieving a common goal. Cooperation allows to take the maximum advantage from a multi-agent system; more specifically, cooperative decision making is useful to have balanced environment exploration and resource utilization Park and Oh (2020), real-time adaptivity levels Rizk et al. (2018), collision avoidance Shahidian and Soltanizadeh (2017), as well as fault detection and mitigation capabilities Yang et al. (2019a). In literature, multi-agent cooperation strategies have been implemented relying on different models and methodological frameworks, including game theory Grünwald et al. (2004); Park and Oh (2020), optimal control Park and Oh (2020), reinforcement learning Buşoniu, Babuška, and De Schutter (2010), finite state machines Pack et al. (2009), Markov decision processes (MDP) Rizk et al. (2018), swarm intelligence Chung et al. (2018), and evolutionary computing Pérez-Carabaza et al. (2019a)[9].

Information fusion protocols Multi-platform systems, either cooperative or not, strongly depend on the information fusion technique adopted to spread robots data over the network Parker, Rus, and Sukhatme (2016). In turn, the functionality of an information fusion strategy greatly depends on robots networking capabilities. Data aggregation can be achieved through different communication protocols, namely, centralized Koohifar et al. (2018, 2016); Shahidian and Soltanizadeh (2017), decentralized Bourgault et al. (2003); Lanillos, Gan, Besada-Portas, Pajares, and Sukkarieh (2014); Park and Oh (2020), or distributed Atanasov, Le Ny, and Pappas (2015); Hollinger, Yerramalli, Singh, Mitra, and Sukhatme (2014); Julian, Angermann, Schwager, and Rus (2012); Pack et al. (2009).

- *Centralized* architectures require the presence of a central base station[10]. Even though this setup may be practical in specific scenarios Mukerjee, Naylor, Jiang, Han, Seshan, and Zhang (2015), many real-life applications does not allow to rely on a central processing unit, capable of communicating with all network devices and of supporting high computational burdens. At the same time, a centralized system leverages on a single point of failure (i.e., the base station) which may compromise the operations of the entire architecture. The central node represents also a communication bottleneck: it might not receive complete or updated information,

[9]The reader may refer to Rizk et al. (2018) for a comprehensive survey on cooperative multi-agent systems and to Khaleghi et al. (2013); Nakamura et al. (2007); Smith and Singh (2006) for surveys on information fusion in APE-related tasks.

[10]The central base station is capable of communicating with all sensing robots and it embeds sufficient memory and computational resources to store and elaborate the large amount of data received from the sensor network. The base station can be either an external unit, or one of the M sensing robots; in this latter case, the platform needs to be *smart* and with enough on-board capabilities to support the processing, communication and storage demand Kyung et al. (2016).

due to sensing and communication limitations, or because of any fault. Finally, centralized protocols suffer from poor modularity and they do not scale well with the network size, since the communication and computation costs increase with the number of agents.

- *Decentralized* architectures are used to increase the scalability and robustness levels with respect to centralized counterparts. Instead of a unique node, decentralized protocols convey the information flow on multiple base stations Ghods et al. (2020). In this way, it is also possible to guarantee higher levels of data redundancy within the network; this, in turn, results in more accuracy, reliability and security. Moreover, the presence of multiple nodes endowed with processing and storage capabilities allows to reduce the computational and memory burden on each single base station.

- *Distributed* protocols require individual robots to operate with local information, that is, their own measurements and the data received from the communicating nodes. To this aim, each agent of the network can be seen as a base station, with sufficient computational and storage resources on-board. Distributed architectures are extremely scalable and fault tolerant. Furthermore, they leverage on partially available information; therefore they are only marginally affected by communication problems Nakamura et al. (2007). However, distributed optimization does not always guarantee optimality bounds as high as centralized methods do Ge, Yang, and Han (2017); Tenney and Sandell (1981); Yang, Yi, Wu, Yuan, Wu, Meng, Hong, Wang, Lin, and Johansson (2019b). The reason is that centralized fusion has a global knowledge (i.e., all measured data is available), whereas distributed fusion has only localized information coming from the set of communicating nodes Nakamura et al. (2007).

3.3.2 Discussion

Teams of cooperative Koohifar et al. (2016) and (possibly) heterogeneous Ferrari, Fierro, and Tolic (2009) robots are used to increase the robustness and the efficiency of APE systems Khan et al. (2016), through multi-robot cooperation and multi-sensor integration Park and Oh (2020). Cooperation is particularly useful when the reconstructed belief map has multiple optima; in this case, coordinated motion planning promotes informative joint observations, which enables quick verification and elimination of unlikely target position hypotheses Bourne et al. (2019). In spite of this, and although decision making in Multi-Agent Systems (MAS) has seen significant improvements in the past decade, there is only limited research on MS-APE, as highlighted in Tab. 3.3. This has led to the presence of numerous open research questions that still need to be addressed Rizk et al. (2018), making MS-APE a potential flourishing research field for the next years. The most interesting research areas regard the design of multi-platform schemes that account for on-board limited resources, network scalability, as well as robots heterogeneity. To foster future developments in these sectors, some considerations and suggestions are provided in the following.

Deployment issues The growth of MAS technologies is mainly limited by deployment issues: as a matter of fact, multi-robot systems require more time-consuming deployment procedures, have additional hardware requirements, and come at the cost of higher computational, communication and storage system resources Hoffmann and Tomlin (2009), especially when cooperative behaviors are required. These factors keep pushing the research and development efforts towards single-platform solutions, when task-compliant. Nonetheless, a single platform might not guarantee the same effectiveness of a multi-platform system; hence, more research is needed on software and hardware tools to facilitate multi-robot communication and multi-sensor integration through plug-and-play solutions.

Limited resources Robots are commonly characterized by limited on-board resources and they might not be able to offload their computations to the cloud, due to bandwidth scarcity, unreliable connectivity, or minimum latency requirements Rizk et al. (2018). Tightly coordinated tasks also increase the computational burden, due to the large amount of communication and data exchange among agents. In addition, most sensing platforms are battery-powered (e.g., UAVs Meera et al. (2019)); hence, their lifetime is related to the presence of energy-harvesting sensors, the communication activity, and the processors usage SanMiguel and Cavallaro (2016). Finally, communication and storage limitations might preclude the possibility to collect (or store) updated information from other robots. In conclusion, limited on-board resources reduce the applicability of MAS schemes and make cooperative algorithms challenging in real-life scenarios. Therefore, future research should focus on the development of accurate models that capture the entire spectrum of resource consumption (e.g., energy, bandwidth, storage) as function of the network activity. This models will allow to apply predictive decision-making strategies and guarantee an optimal balance between resource consumption and network operation.

Scalability Another key issue in multi-platform APE algorithms is the system scalability with respect to the number of platforms. In this regard, decentralized and distributed architectures should be considered; notably, they require agents with minimal on-board processing, communication and storage capabilities, but they do not offer the same optimality guarantees with respect to centralized counterparts Yang et al. (2019b). In mobile robotics, distributed optimization is often used to yield agents coordinated motion. In doing so, each sensing robot locally estimates the utility function to be optimized by aggregating its previous observations with those coming from connected agents; the computational complexity is therefore constant with respect to the number of platforms. This approach is however not cooperative, since each agent only considers its own future actions. To enable a higher level of coordination, each platform might consider the effect of other robots actions in the objective estimation process. It incurs a linear computational expense in the number of vehicles, yet the effect of the approximation error is provably reduced from that of the single-node approximation, allowing coupled effects between mobile sensors to be captured. Indeed, field experiments have shown better results using the cooperative approach Hoffmann and Tomlin (2009). Robots cooperation can be implemented also via finite state machines Pack et al. (2009). A Search and Tracking algorithm can start with the team performing a global search and the next platform action might be chosen by optimizing an objective function that accounts for

both individual and collective goals. For instance, it may foster group dispersion over the environment, while discouraging fuel consumption, or it should balance the revisits of the monitored areas. Once a target has been detected, the search and track phase starts: robots perform task assignment (continue the search or track the detected target), according to the distance to the target, the current number of platforms focused on that target, and the number of targets that have already been detected Pack et al. (2009).

Heterogeneity Robots heterogeneity stems from diverse sensing and actuating capabilities, computing resources, cognitive algorithms, and morphology Parker et al. (2016); Rizk et al. (2018). The literature regarding the utilization of heterogeneous robots has shown the benefits of combining either different types of sensors, different perspectives, or different computational and operational capabilities Queralta et al. (2020); Rizk et al. (2019). More precisely, parallelisation, specialisation and complementarity are important motivations for developing heterogeneity in multi-robot teams Esterle et al. (2017). For these reasons, heterogeneity is an increasingly prevalent property in cyber-physical and multi-robot systems Lewis et al. (2014)

When heterogeneity is related to the perception capabilities, it is referred to as multi-modality, as discussed previously in this Section. Multiple sensing modalities capture different properties of the environment; thus, the aggregated multi-modal data allow inferences that might be not possible with single-sensor measurements. However, multi-modality represents an important challenge to be handled in heterogeneous multi-sensor applications. As a matter of fact, different data sources require different calibration procedures, each of which might involve extensive human intervention Aghajan and Cavallaro (2009); Zanella (2016). Recently, some works have proposed self-supervised calibration procedures requiring minimal human effort. These methods hinge on the correlation between the multi-modal data collected by the platform sensors; in particular, one perception channel is used to produce the supervisory signal for the others, enabling the automatic generation of a labeled dataset Nava, Guzzi, Chavez-Garcia, Gambardella, and Giusti (2019). In this regard, Chapt. 4 and 6 deal with two self-supervised learning procedures exploiting radio-visual data. Calibration is not the only problem to be tackled when dealing with multi-modal systems. Indeed, different data sources suffer from different types and magnitudes of noise, and they might produce unbalanced data. This, in turn, affects the decision making process, which struggles when data have different quality and resolution characteristics Khaleghi et al. (2013). Moreover, unbalanced and missing data might affect the belief map generation procedure Khaleghi et al. (2013); Queralta et al. (2020). Multi-modal data sources might also yield conflicting features. For example, visual sensors and RADARs might detect obstacles at different distances Kumar et al. (2020). All these issues are further complicated when multiple robots are involved: the data to be fused are collected in different physical locations and the sensors move with respect to each other. Thus, it is not trivial to evaluate whether different agents are observing the same objects or not, or how to rank observations from different agents Queralta et al. (2020).

From the operational perspective, certain applications (e.g., cooperative localization, surveillance, assistive navigation, exploration and mapping) may require the simultaneous use of robots with actuation and morphological heterogeneity Chen et al. (2021). For instance, smaller platforms have access to narrow passages, but their payload is

limited Chen et al. (2021); to solve this issue, it is possible to employ larger robots that carry application-specific payloads (e.g., long-duration batteries, heavy sensors or actuators) Rizk et al. (2018). Heterogeneous robotic systems represent a fertile research field, where there are still numerous open questions to be addressed. At first, the wide variety of robots being utilized in APE missions, and the different scenarios in which they can be applied, calls for a certain level of collective and situational awareness. This requires the definition of self-organized swarm intelligence models, which account for robots diversity and take into consideration the surrounding environmental conditions. In this framework, the agents of a heterogeneous multi-robot system should identify the main characteristics and constraints of the current environment, and the conditions of the robots they interact with. In particular, heterogeneity-aware APE strategies in large scale MAS scenarios, with variable number and agents capabilities, require each robot to automatically learn the capabilities of their peers. For instance, a UGV collaborating with other UAVs needs to be aware of the different perspectives that UAVs can bring into the scene, but also of their limitations in terms of operational battery lifetime or payload capabilities. In general, heterogeneity-aware algorithms, and the explicit incorporation of agents diversity, can be advantageous for multi-robot decision making processes Davoodi and Velni (2020). Nevertheless, there is still lack of further research in this area, as most existing projects and systems involving heterogeneous robots predefine the way in which they are meant to cooperate Queralta et al. (2020). This is mainly due to the high complexity associated to the design of heterogeneity-aware APE strategies, where robots are meant to autonomously plan their actions, according to the diverse capabilities of their peers. As a consequence, several works introduce additional simplifications or assumptions, such as the hypothesis that the robots sensing and actuation capabilities are a-priori known over the entire network Chung, Kress, and Royset (2009); others design ad-hoc strategies for a MSMP-APE system where the type and the number of robots is predefined Dogancay (2012). These assumptions limit the scalability and the fault-tolerance of the system, since they do not account for the possibility of malfunctioning or missing robots Maza and Ollero (2007).

3.4 The platform dynamics

Active sensing can be defined as an intelligent data acquisition process where a feedback is performed on sensory data (raw or processed) to satisfy specific perception performance measures, such as scene exploration Mavrommati et al. (2017), target tracking Ferrari et al. (2009), and object classification Patten et al. (2018). According to (3.13), the measurements are affected by the current platform configuration. This implies that any active sensing scheme must include a decision-making process to dynamically modulate the sensing agent behavior through its state Bajcsy et al. (2018). To this aim, the sensing platform is modeled as a dynamical system in most autonomous perception frameworks Liu et al. (2015). In particular, the dynamics of sensing platform m, defined by its state $\mathbf{s}_{t,m}$, is usually described with a deterministic Markovian state transition

model Radmard and Croft (2017):

$$\begin{aligned} &\mathbf{s}_{t+1,m} = q_m\left(\mathbf{s}_{t,m}, \mathbf{u}_{t,m}\right), \ m = 1, \ldots, M \\ &\text{s.t. } \mathbf{s}_{t,m} \in \mathcal{S}_m, \ \mathbf{u}_{t,m} \in \mathcal{A}_m \end{aligned} \tag{3.26}$$

where $\mathcal{A}_m$ is the *control space* (or action set), comprising all possible control actions that can be applied by the platform actuators: a finite set of platform displacements in mobile robots Park and Oh (2020), or a set of Pan-Tilt-Zoom configurations in camera-enabled platforms Denzler et al. (2003); Lathuilière et al. (2019)). The control space represents the platform degrees of freedom (i.e., the controllable state variables), that is, a subset of the state space $\mathcal{S}_m$. The *control input* applied to the platform is denoted as $\mathbf{u}_{t,m}$; finally, a (possibly) is non-linear function

$$q_m(\cdot) : \mathcal{S}_m \times \mathcal{A}_m \rightarrow \mathcal{S}_m \tag{3.27}$$

describes the dynamics of platform m and it mainly depends on the platform structural properties (e.g., aerial, ground or underwater robot) Furukawa et al. (2012); Radmard and Croft (2017). It is worth noting that, in this general setting, each sensing robot has its own control space: as a matter of fact, there may be two agents, say m and m', whose motion capabilities are different (i.e., $\mathcal{A}_m \neq \mathcal{A}_{m'}$). For instance, in heterogeneous camera networks the action set of PT cameras may be different from that of cameras mounted on aerial robots. As introduced in Sec. 3.1.2, the state space can be Euclidean, a rotation manifold, Riemannian or, more in general, a high-dimensional vector. In this regard, it is worth noting that, in practice, dealing with a continuous multi-dimensional state space might be computationally demanding and sometimes unfeasible Ryan and Hedrick (2010). For this reason, $\mathcal{S}_m$ and $\mathcal{A}_m$ are usually discretized Chung et al. (2009); Czarnetzki et al. (2010); Katsilieris et al. (2012); Van Nguyen et al. (2020).

Once the platform dynamic model has been defined, it is important to describe the remaining core elements of any AS algorithm, namely the control law and the criteria adopted to drive the platform behavior.

3.4.1 The control

In Active Sensing, the choice of the control input $\mathbf{u}_{t,m}$ is often cast to an optimization problem of a suitable utility function $J(\cdot)$ Radmard and Croft (2017). This combines the criteria that guide the perception process and it is defined over the control space $\mathcal{S}_m$. In probabilistic AS settings, $J(\cdot)$ depends on some information extracted from the belief map (see Sec. 3.4.4). Sometimes, if the platform is multi-modal and collects observations from $K < C_m$ different sensing modalities, $J(\cdot)$ is computed as the aggregation of K utilities, one for each sensing modality Dogancay (2012).

Additional constraints may be introduced in the optimization problem in order to account for finite resources availability (e.g., time, energy) Nguyen, Lawrance, Fitch, and Sukkarieh (2013); Noori et al. (2016); Pérez-Carabaza et al. (2019a), communication limitations Doğançay et al. (2009), or collision avoidance requirements Pérez-Carabaza

et al. (2019b). The control input is therefore computed as Hutchinson et al. (2019b)

$$\mathbf{u}_{t,m} = \mathbf{u}^*_{t,m} = \underset{\mathbf{u}\in\, \mathcal{A}_m \setminus \mathcal{C}_{t,m}}{\arg\max}\; J\left(\mathbf{u} \mid \mathbf{z}_{1:t,m}, \mathbf{s}_{t,m}\right), \tag{3.28}$$

where (3.28) is computed for each platform, that is for $m = 1, \dots, M$; the notation $J\left(\mathbf{u} \mid \mathbf{z}_{1:t,m}, \mathbf{s}_{t,m}\right)$ says that $J(\cdot)$ is function of $\mathbf{u}$, given the actual platform state and the sequence of observations until time t, namely $\mathbf{z}_{1:t,m} = \{\mathbf{z}_{1,m}, \dots, \mathbf{z}_{t,m}\}$; $(\cdot)^*$ denotes the solution of an optimization problem; $\mathcal{C}_{t,m}$ is the set of constraints for the m-th robot at time t. Obviously, $\mathcal{C}_{t,m}$ is time and platform-dependent; for instance, if a UGV is close to the borders of a room and another one is at the centre of the same room, the former has constrained motion capabilities, while the latter is completely free to move.

In (3.28) each robot exploits its local information (i.e., $\mathbf{z}_{1:t,m}$) and behaves independently to the others, since $J(\cdot)$ depends only on $\mathbf{s}_{t,m}$ (i.e., non-coordinated APE). If a coordinated MSMP-APE is employed, the objective function depends on all measurements collected by other agents (or a subset of them, if the information sharing protocol is not centralized), according to (3.20). The problem (3.28) becomes

$$\mathbf{u}^*_{t,m} = \underset{\mathbf{u}\in\, \mathcal{A}_m \setminus \mathcal{C}_{t,m}}{\arg\max}\; J\left(\mathbf{u} \mid \mathbf{z}_{1:t,1:M}, \mathbf{s}_{t,m}\right). \tag{3.29}$$

Finally, if the MSMP-APE is cooperative Park and Oh (2020), the decision making process is aware of other agents state and actions, namely

$$\begin{aligned} &\mathbf{u}^*_{t,1:M} = \underset{\mathbf{u}_{1:M}}{\arg\max}\, J\left(\mathbf{u}_{1:M} \mid \mathbf{z}_{1:t,1:M}, \mathbf{s}_{t,1:M}\right), \\ &\text{s.t. } \mathbf{u}_{1:M} = \{\mathbf{u}_1, \dots, \mathbf{u}_M\}, \quad \mathbf{u}_m \in \mathcal{A}_m \setminus \mathcal{C}_{t,m} \end{aligned} \tag{3.30}$$

where $\mathbf{s}_{t,1:M} = \{\mathbf{s}_{t,1}, \dots, \mathbf{s}_{t,M}\}$ and similarly for $\mathbf{u}^*_{t,1:M}$.

3.4.2 Control space, constraints and optimization

This Section classifies the APE-related literature according to the control space, the principal platform constraints, and the optimization methods adopted to solve the control problem (3.28)[11]; the objective is to highlight how the same APE task can be tackled by leveraging on platforms with different degrees of freedom and by applying different control techniques. In fact, as Tab. 3.4 suggests, some works exploit approximate Radak et al. (2017), greedy Cliff et al. (2018), or heuristic control laws Pérez-Carabaza et al. (2019b); others adopt optimal and predictive strategies Shahidian and Soltanizadeh (2017), and a few propose statistical black-box optimization methods Meera et al. (2019). This variability comes from the fact that the performance metrics and the optimization strategies vary when operating on different action spaces (e.g., rotation manifolds or canonical Euclidean spaces) Varotto et al. (2019a); furthermore, the choice of the control algorithms depends also on the constraints that characterize the platform movements Pérez-Carabaza et al. (2019b). Finally, the cardinality of $\mathcal{A}_m$ has a direct impact on the computational demand and on the optimality guarantees of a control strategy; in particular, small action

[11] The focus is on the basic optimization problem (3.28); the variation (3.29) involves a larger amount of available information, while (3.30) extends the problem dimensionality.

sets Van Nguyen et al. (2020) and approximation schemes Czarnetzki et al. (2010) are sometimes necessary to meet real-time performance requirements.

Platform position In many cases it is possible to regulate the planar Koohifar et al. (2016) or the 3D Ramirez et al. (2014) position of ground Twigg et al. (2012) and aerial Chung et al. (2018); Koohifar et al. (2016) robots, without any control on the sensors orientation. Some examples are given by fixed radio antennas mounted on UGVs Twigg et al. (2012), atmospheric sensors Park and Oh (2020), or fixed cameras embedded in UAVs (and usually facing downwards) Ramirez et al. (2014). When the cardinality of the action space is small, the optimal control input $\mathbf{u}^*_{t,m}$ can be chosen according to an extensive search over all possible candidate actions (i.e., *grid-search optimization*) Hutchinson et al. (2019b). Grid-search becomes impracticable with continuous action spaces, while *gradient-based controllers* are more suitable Shahidian and Soltanizadeh (2017). The main limitation of gradient-based algorithms is the sensibility to local optima; hence, their applicability is restricted to convex functions over the control space. However, the presence of hard constraints might violate this condition; furthermore, the need to analytically evaluate the gradient of the cost function represents a further limitation of gradient-based controllers. To reduce the sensibility to local optima, several strategies can be adopted. In Radak et al. (2017) an UAV is meant to fly a predefined distance towards the peak of the target belief map and, in case of multiple peaks, the heading angle is chosen randomly; the introduction of a stochastic term in the decision process might allow the platform to escape from local optima, at the cost of possible unstable and inefficient trajectories. On the other hand, when the gradient of the cost function can not be analytically computed, or it is computationally expensive to evaluate, it can be estimated through the collected data $\mathbf{z}_{1:t}$ via *gradient-assisted* path-planning algorithms Sun et al.

Table 3.4. Classification of related works based on the control space $\mathcal{A}_m$, the set of constraints $\mathcal{C}$ and the optimization method. Key - MCTS: Monte Carlo Tree Search; BO: Bayesian Optimization

Ref.	Control space			Constraints	Opt. method
	Position	Orientation	Other		
Czarnetzki et al. (2010)		pan-tilt			grid-search
Lathuilière et al. (2019)		pan-tilt		pan speed	Q-learning
Denzler et al. (2003)			zoom		grid-search
Katsilieris et al. (2012)		pan			grid-search
Hoffmann and Tomlin (2009)	2D			collision avoidance	grid-search
Shahidian and Soltanizadeh (2017)	2D			-collision avoidance -connectivity to master	gradient
Sun et al. (2008); Twigg et al. (2012) Fink and Kumar (2010)	2D				gradient
Van Nguyen et al. (2020)	2D			safe distance from the target	grid-search
Koohifar et al. (2016)	2D			turning angle	MPC
Koohifar et al. (2018)	2D				MPC
Liu and Hedrick (2017)	2D			platform dynamics	MPC
Liu et al. (2015)	2D			collision avoidance	MPC
Cliff et al. (2018)	2D				greedy
Hutchinson et al. (2019b); Park and Oh (2020) Radak et al. (2017)	2D				grid-search
Pack et al. (2009)	2D			energy consumption	grid-search
Haubner et al. (2019b)	2D				MCTS
Varotto et al. (2021)	2D	pan-tilt		energy consumption	gradient
Vander Hook et al. (2015)	2D			communication	heuristic
Dogancay (2012)	2D			- collision avoidance - agents connectivity - turn rate	interior point
Ramirez et al. (2014)	3D				grid-search
Meera et al. (2019)	3D			- collision avoidance - time of flight	BO

(2008); Twigg et al. (2012). This kind of strategies are completely model-free (i.e., the observation model is not needed), which makes them highly adaptive. However, they also have many limiting factors: first, they are not effective when the underlying data generative model does not follow a convex function; secondly, they can be applied only when the target is static; finally, in turbulent, cluttered environments, or when sensors are noisy (e.g. in RSSI-based source localization Sun et al. (2008)), the gradient does not always lead directly to the source, with consequent unstable platform behaviors. When the form of the objective function is complex and intractable to analyze, non-convex, nonlinear, high-dimensional and noisy, *heuristic strategies* can be adopted (e.g., ant colony optimization Pérez-Carabaza et al. (2019b), genetic algorithm Pérez-Carabaza et al. (2019a)). Interestingly, heuristic optimization techniques are particularly suitable to handle platform constraints (e.g., communication and collision avoidance) Pérez-Carabaza et al. (2019b), even though they might produce sub-optimal solutions in terms of tracking performance (e.g., localization error) w.r.t. optimal control approaches Koohifar et al. (2016).

Platform orientation When the platform is static (i.e., its position can not be regulated), the APE task can only be accomplished through rotation movements. This is the case of orientable RADAR scanners Katsilieris et al. (2012) and PT cameras, which span the surrounding environment according to discretized action spaces, composed by a fixed number of angular sectors. When the objective is to perform active tracking, but the target dynamics is unknown, Active Disturbance Rejection Controllers (ADRC) Chen, Zhao, and Tan (2014) can be used: the uncertainties on the target dynamics are estimated and compensated through an extended state observer (ESO), and then treated as system disturbances. As an alternative to ADRCs, Model Predictive Controllers (MPCs) can be used Saragih, Kinasih, Machbub, Rusmin, and Rohman (2019b); MPCs generate optimal control signals, have predictive capabilities, embed input saturation constraints in the objective function Liu and Hedrick (2017), and allow to solve non-linear control problems. The optimization problem (3.28) accounts for a one-step look-ahead action; hence, it is referred to as *myopic* path-planning Doğançay et al. (2009). In *receding horizon* path-planning (i.e., MPC) Koohifar et al. (2018, 2016); Liu and Hedrick (2017); Liu et al. (2015) the problem defined in (3.28) is extended to a multiple-step look-ahead horizon. It is important to remark that MPC usually implies high computational costs, it needs an accurate dynamic model of the target, and model uncertainties are very difficult to resolve[12].

High-dimensional control spaces High-dimensional control spaces are common in many robotic systems, like visual servoing architectures and eye-in-hand manipulators Radmard and Croft (2017). Robots with large number of degrees of freedom are proven to be more robust and more efficient than platforms with limited action spaces. As an example, when the purpose is to visually detect a target, acting on the entire camera pose (i.e., position and pan-tilt orientation) is beneficial Varotto et al. (2021): the camera orientation allows to observe the target even when it is not close to the robot; to maximize the target detectability, it is possible to control the robot position, under the assumption that the POD is inversely proportional to the target-platform distance, as in (3.17);

[12]The reader may refer to Bemporad, Morari, and Ricker (2005); Wang (2009) for a comprehensive overview on predictive control

the same formula suggests that, if camera sensors are employed, it is also possible to increase the target detectability by acting on the focal length Pérez-Carabaza et al. (2019a); Shubina and Tsotsos (2010). The main issue of dealing with high-dimensional control spaces is that gradient-based solvers are typically not effective. This is due to the fact that the objective functions are rarely convex with respect to the control space. Hence, heuristic Pérez-Carabaza et al. (2019a) or grid-search Shubina and Tsotsos (2010) optimization strategies are commonly applied to solve (3.28). Usually, to account for joint limits and dynamics constraints, a multiplicative penalty term is introduced to the cost function, or implicitly embedded into the sensor motion model Radmard and Croft (2017).

Constraints Most existing strategies for Active Position Estimation do not include any constraint in the task formulation Hutchinson et al. (2019b); Pack et al. (2009); Park and Oh (2020); Radak et al. (2017), that is, $\mathcal{C}_m = \varnothing$. However, in real-life applications the platform movements are limited by dynamics saturation levels Liu and Hedrick (2017) (e.g., robot speed), while the network communication capabilities are commonly affected by distance-based inter-agent constraints Pérez-Carabaza et al. (2019b); Vander Hook et al. (2015). In addition, time and energy resources play a critical role in APE tasks, usually in the form of (soft or hard) specifications to be satisfied (e.g., in UAV-based Search and Rescue missions) Matthiesen (2010); Nguyen et al. (2013); Noori et al. (2016). Finally, when multiple platforms are employed, or when the environment is cluttered, the risk of collision among robots Hoffmann and Tomlin (2009), or with surrounding obstacles Liu et al. (2015); Meera et al. (2019), should be minimized.

If the constraints can not be violated at any moment during the APE mission (i.e., hard constraints), they must be explicitly considered in the optimization problem (3.28). This is the case of sensing robots with a fixed amount of internal resources, such as battery powered cameras Varotto and Cenedese (2021b), or unmanned vehicles Varotto et al. (2021). An alternative is to incorporate the constraints into the cost function (i.e., soft constraints), obtaining a new unconstrained optimization problem over a functional that aggregates both APE requirements and constraints Meera et al. (2019); Pack et al. (2009); Popovic (2019); Vander Hook et al. (2015). This is particularly useful when robots are allowed to temporarily break their constraints to take better measurements (i.e., move closer to the target, or acquire more informative data).

As suggested by Tab. 3.4, the APE literature offers a wide variety of constrained scenarios; nonetheless, only few of the reviewed works consider energy consumption Varotto et al. (2021), or time of search Miki et al. (2018), as hard specifications. Thus, future research should focus on the impact of hard resource constraints in APE performance Varotto et al. (2021). To this aim, high-dimensional control spaces are found to be useful in balancing resource consumption and APE task completion Pérez-Carabaza et al. (2019a). For instance, (3.17) suggests that, in camera-based APE tasks, the only way to have effective detections is either to zoom-in, or to move closer to the target; however, platform movements are much more energy consuming than the regulation of the camera focal length. Thus, zoom control might have a beneficial impact when dealing with energy-aware SanMiguel and Cavallaro (2014, 2016) and Quality of Sight-aware SanMiguel and Cavallaro (2017) problems, where both detection capabilities and energy preservation matter.

3.4.3 Discussion

As discussed in Sec. 3.4.1, active position estimation is the control of the perception process, which can be formulated as a mathematical optimization program where the cost function depends on the platform state $\mathbf{s}_t$. All the articles reviewed in this Chapter solve the control problem (3.28) assuming perfect knowledge on the platform state. Nonetheless, this is an ideal assumption, since the platform state is often measured from noisy on-board sensors (e.g., IMU units), or it needs to be estimated from other agents Varotto et al. (2019a); in this case, the measurement or estimation inaccuracies might have a significant impact on the active sensing performance Popovic (2019). Therefore, future research should focus on the formulation of APE solutions that mathematically couple both target and platforms uncertainties in a single objective function.

Reinforcement learning (RL) Hernandez-Leal et al. (2019) is a promising research direction within the robotics and control community. The reinforcement learning framework is based on an agent that interacts with the surrounding environment, it observes the consequence of any action taken, and it is capable to measure the success of its actions via a punishment/reward system. In this way, experience is used to learn a policy that, given the current state, chooses the next best action to maximize the expected future reward. Along this line, RL has recently found a fertile ground in the active sensing research field Queralta et al. (2020), since the essence of active sensing is understanding and adapting to changes in the environment, so that future sensing actions can be planned. Interestingly, the mathematical formulation of RL perfectly fits with the APE characterization given in (3.28): the cost function $J(\cdot)$ represents the expected reward, while $\mathbf{u}^*_{t,m}$ is the best action to take. For instance, in camera-based active tracking the reward is typically based on a predefined viewpoint of the target (e.g., a desired viewing angle or a given platform-target distance Luo, Sun, Zhong, Liu, Zhang, and Wang (2018)); in some cases the reward considers the uncertainty associated to the target location estimate Calli et al. (2018). In active search instead, the reward may account for the time of search, the presence of cluttering in the environment, and the probability of detecting the target, which is usually related to the platform-target distance, as in (3.17) Meera et al. (2019); Sandino, Vanegas, Gonzalez, and Maire (2020). The main advantage of RL is that the learning process does not require a dataset of labeled input-output pairs; thus, the robot learns without any human supervision. Usually, the robot starts by performing some random actions and gradually learns to follow a desired search and tracking behavior. Another advantage of RL is that it provides an end-to-end approach from sensing to actuation; therefore, it integrates the perception and control aspects within a single framework Queralta et al. (2020). Finally, the RL framework can also be extended to multi-agent and cooperative scenarios (i.e., multi-agent reinforcement learning, MARL Buşoniu et al. (2010); Hernandez-Leal et al. (2019)), with direct application to distributed and coordinated MSMP-APE problems. Despite the aforementioned benefits, a critical question of reinforcement learning methods is that they require long training sessions, which are not always desiderable in robotics applications Lathuilière et al. (2019). Furthermore, the training process of a RL model involves random exploratory actions, which can be potentially unsafe (especially in human-robot

cooperation) and costly (in terms of hardware damage) Arndt, Hazara, Ghadirzadeh, and Kyrki (2020). These issues are particularly critical in SAR operations, where it is difficult to collect or to create sufficient training data and experiences Queralta et al. (2020). More in general, the dependence on training data may limit the applicability of RL-based solutions in unstructured environments, where very few assumptions can be made on the data generation and acquisition process; hence, the robustness performance obtained during the training phase can not be guaranteed during system deployment. For example, dramatic lighting changing conditions significantly affect the performance of an image recognition algorithm Calli et al. (2018); thus, an unstable system is obtained if during training phase it has not explored and experienced all possible lighting conditions. Similar issues may arise in camera-based search and tracking tasks where the target is unknown or its appearance does not match with the training dataset Calli et al. (2018). Recently, a promising solution to deal with dynamic and unstructured environments lies in transfer learning techniques Arndt et al. (2020); Lathuilière et al. (2019); Luo et al. (2018): ad-hoc simulators are employed to train an inference model; this is then fine-tuned and adapted to the real world environment, since a simulation may not capture all physical phenomena Arndt et al. (2020). As an alternative, it is possible to rely on model-free control strategies (e.g., extremum seeking control Calli et al. (2018)), which do not utilize training data and implement continuous adaptation strategies to respond to the current state of the environment. To sum up, in view of the aforementioned challenges, more research is needed on efficient and robust control techniques in unstructured and cluttered environments Calli et al. (2018); Meera et al. (2019); Sandino et al. (2020).

3.4.4 APE criteria

In Active Sensing applied to target search and localization tasks, the platform state must be regulated in order to satisfy several quality of sensing and quality of localization requirements. For instance, a PTZ camera can be controlled to guarantee minimal detection capabilities or maximum probability of detection; in the former case it is sufficient to move the camera so that the target is inside the FoV, namely $\mathbf{p}_t \in \Phi(\mathbf{s}_{t,m})$[13]; in the latter case, it is necessary to model and optimize the POD function through the platform movements. Another example comes from radio-based tracking systems; these often necessitate of automatic algorithms for antennas alignment to optimize the quality of communication between the platform and an active target with which communication has been established Li, He, Liu, Wan, Gu, Xie, Fu, and Lu (2019).

In this Section, more examples are provided; in particular, APE algorithms are classified as either information-seeking or task-driven, according to the cost function involved in the optimization problem (3.28) (see Tab. 3.5).

Information-seeking control

Information-seeking control aims at choosing the next sensor state such that the expected future measurement optimizes a certain information gathering measure; for this reason, the APE strategies belonging to this control class are biased towards an *explorative*

[13] In this Section the formulas are related to single-platform single-target solutions; hence, to ease the notation, the subscripts used to denote the m-th platform and the ℓ-th target can be omitted.

accomplishment of the positioning task[14]. The most intuitive information-seeking criterion is the minimization of target position *estimate variance* Chung and Furukawa (2009).

Other criteria are based on the optimization of real-valued functions of the *Fisher Information Matrix* (FIM), which is related to the estimation uncertainty ellipsoid Ucinski (2004); Vander Hook et al. (2015). For instance, A-optimality minimizes the trace of the FIM and represents the average variance of the estimate Shahidian and Soltanizadeh (2017); D-optimality instead minimizes the log-determinant of the FIM and it is related to the volume of the uncertainty ellipsoid Dogancay (2012). D-optimality is proven to outperform A-optimality in terms of trace of the error covariance matrix Shahidian and Soltanizadeh (2017). Both criteria are suitable for gradient-based controllers Shahidian and Soltanizadeh (2017) and Model Predictive Controllers Koohifar et al. (2018, 2016). In this latter case, the FIM can be approximated via Monte Carlo numerical methods. In general, the main issue of FIM-based approaches is that they rely on state estimators, like the Kalman Filter and the Extended Kalman Filter, that are ill-suited when target state and sensor models follow severe non-linear/non-Gaussian distributions. In this context, several studies have proven that particle filters Gordon, Ristic, and Arulampalam (2004); Koohifar et al. (2018) appear to be more appropriate, reliable and accurate than conventional Kalman-based methods in many tracking applications.

The information gain obtained with action $\mathbf{u}_t$ can be also quantified by the *mutual information* (MI) between the target position and the sensor observation, namely Hoffmann and Tomlin (2009)

$$J(\mathbf{u}_t|\mathbf{z}_{1:t}, \mathbf{s}_t) = I\left(\mathbf{z}_{t+1}(\mathbf{s}_{t+1}); \mathbf{p}_{t+1}\right) = H(\mathbf{p}_{t+1}) - H\left(\mathbf{p}_{t+1}|\mathbf{z}_{t+1}(\mathbf{s}_{t+1})\right). \tag{3.31}$$

$H(\mathbf{p}_{t+1})$ is the entropy of the random vector $\mathbf{p}_{t+1}$, while $H\left(\mathbf{p}_{t+1}|\mathbf{z}_{t+1}(\mathbf{s}_{t+1})\right)$ represents

Table 3.5. Classification of related works based on the APE criteria. Key – FIM: Fisher Information Matrix; MI: Mutual Information; H: Entropy; D_{KL}: Kullback-Liebler divergence; POD: Probability of Detection

Ref.	APE criterion	
Lathuilière et al. (2019)	-source detection -people in FoV	task-driven
Fink and Kumar (2010); Pack et al. (2009); Radak et al. (2017); Ramirez et al. (2014) Liu and Hedrick (2017); Van Nguyen et al. (2020); Varotto et al. (2021)	distance to target estimate	task-driven
Sun et al. (2008); Twigg et al. (2012)	highest RSSI value	task-driven
Katsilieris et al. (2012)	expected POD	task-driven
Bourgault et al. (2003)	cumulative POD	task-driven
Liu et al. (2015)	-distance to target estimate -POD	task-driven
Vander Hook et al. (2015)	mission time	task-driven
Meera et al. (2019)	target detectability	task-driven
Shahidian and Soltanizadeh (2017)	D/A-optimality	information-seeking
Dogancay (2012); Koohifar et al. (2018, 2016)	D-optimality	information-seeking
Vander Hook et al. (2015)	FIM spectrum	information-seeking
Cliff et al. (2018); Hoffmann and Tomlin (2009)	MI	information-seeking
Hutchinson et al. (2019b); Katsilieris et al. (2012)	expected D_{KL}	information-seeking
Park and Oh (2020)	infotaxis	information-seeking
Czarnetzki et al. (2010); Denzler et al. (2003); Haubner et al. (2019b)	expected H	information-seeking
Fink and Kumar (2010); Meera et al. (2019)	predictive variance	information-seeking
Liu and Hedrick (2017)	covariance trace	information-seeking
Ramirez et al. (2014)	conditional H	information-seeking

[14]The reader may refer to Gupta, Smith, and Shalley (2006) and Xu and Zhang (2014) for an overview on the concepts of exploration and exploitation.

the conditional entropy, that is the expected entropy of the posterior belief, computed as

$$\begin{aligned} H\left(\mathbf{p}_{t+1}|\mathbf{z}_{t+1}(\mathbf{s}_{t+1})\right) &= \mathbb{E}_{\mathbf{z}_{t+1}(\mathbf{s}_{t+1})}\left[H(\mathbf{p}_{t+1}|\mathbf{z}_{t+1}=\mathbf{z}(\mathbf{s_{t+1}}))\right] \\ &= \int_{\mathbf{z}(\mathbf{s}_{t+1})\in\mathcal{Z}_{t+1}} p(\mathbf{z}(\mathbf{s}_{t+1}))H(\mathbf{p}_{t+1}|\mathbf{z}_{t+1}=\mathbf{z}(\mathbf{s}_{t+1}))d\mathbf{z}, \end{aligned} \tag{3.32}$$

where, in general, $\mathbb{E}_{\mathbf{x}}[\cdot]$ denotes the expectation w.r.t. the distribution of the x; $\mathcal{Z}_{t+1}$ is the set of all possible measurements that can be collected at time $t+1$, typically parametrized by the platform state and the control input, from (3.26) and (3.13). In fact, a new measurement is influenced by the next platform state (i.e., $\mathbf{s}_{t+1}$) and, in turn, by the current platform state (i.e., $\mathbf{s}_t$) and the actual control input (i.e., $\mathbf{u}_t$), through (3.26). In this way, the optimization problem (3.28) is well defined. As (3.31) suggests, the cost function depends also on the past observations (i.e., $\mathbf{z}_{1:t}$). From (3.13), each measurement $\mathbf{z}_k$ is function of the target position; therefore, it can be used to recursively update the target belief map, which is then employed to estimate the target position, or to approximate any function of the target position (e.g., the conditional entropy) Boers et al. (2010).

If the value of $\mathbf{z}_{t+1}=\mathbf{z}$ is known in advance Ramirez et al. (2014), then the conditional entropy would be $H(\mathbf{p}_{t+1}|\mathbf{z}_{t+1}=\mathbf{z})$, which can be computed without taking the expectation. Interestingly, mutual information (3.31) is a submodular function Cliff et al. (2018); thus, solving the optimization problem (3.28) via a greedy algorithm provides an optimality (lower) bound of 63% Krause and Guestrin (2011). Formally,

$$J(\tilde{\mathbf{u}}_t \mid \mathbf{z}_{1:t}, \mathbf{s}_t) \geq \left(1-\frac{1}{e}\right) J(\mathbf{u}_t^* \mid \mathbf{z}_{1:t}, \mathbf{s}_t), \tag{3.33}$$

where $\tilde{\mathbf{u}}_t$ is the solution of (3.28) obtained with a greedy algorithm.

Information-seeking capabilities can be achieved also by maximizing the expected *Kullback-Liebler* (KL) divergence before and after a new measurement $\mathbf{z}_{t+1}$ is collected, namely Hutchinson et al. (2019b)

$$J(\mathbf{u}_t|\mathbf{z}_{1:t}, \mathbf{s}_t) = \mathbb{E}_{\mathbf{z}_{t+1}}\left[D_{KL}\left(\mathbf{p}_{t+1}|\mathbf{z}_{1:t}, \mathbf{z}_{t+1}(\mathbf{s}_{t+1}) \;||\; \mathbf{p}_{t+1}|\mathbf{z}_{1:t}\right)\right] \tag{3.34}$$

where $D_{KL}(\mathbf{x} \;||\; \mathbf{y})$ is the KL divergence between the random vectors $\mathbf{x}$ and $\mathbf{y}$ Cover (1999). The generalization of the Kullback-Liebler divergence is the Rényi divergence Aoki et al. (2011a); this is parametrized by a parameter α, used to give emphasis on specific parts of the distribution. Empirical results Aughenbaugh and LaCour (2011); Kreucher, Kastella, and Hero Iii (2005c) show that the Rényi divergence, for some values of $\alpha \neq 1$, provide superior performance than the KL divergence. It is also possible to prove that maximizing (3.34) and (3.31) leads to the same control actions Aoki et al. (2011a); in turn, this is equivalent to minimizing the conditional entropy $H(\mathbf{p}_{t+1}|\mathbf{z}_{t+1})$, since $H(\mathbf{p}_{t+1})$ does not depend on $\mathbf{u}_t$ Cliff et al. (2018); Haubner et al. (2019b); Ryan and Hedrick (2010).

The information-theoretic criteria defined so far induce pure explorative behaviors. *Infotaxis* Vergassola, Villermaux, and Shraiman (2007) is instead used to attain a certain level of exploration and exploitation trade-off and it is defined according to the following

cost function

$$J(\mathbf{u}_t \mid \mathbf{z}_{1:t}, \mathbf{s}_t) = p(\mathbf{p}_{t+1})H(\mathbf{p}_{t+1}) + [1 - p(\mathbf{p}_{t+1})]\, I(\mathbf{z}_{t+1}(\mathbf{s}_{t+1}); \mathbf{p}_{t+1}). \tag{3.35}$$

Infotaxis is a gradient-less measure of information gain and it is particularly useful in STE tasks Park and Oh (2020), due to its capability to work under sparse datapoints.

Task-driven control

In literature several definitions of task-driven control have been suggested. A possible task-driven sensor management criterion is based on the minimization of the *expected estimation error*; in this case, the cost function employed in the optimization problem (3.28) is Aoki et al. (2011a)

$$J(\mathbf{u}_t \mid \mathbf{z}_{1:t}, \mathbf{s}_t) = -\mathbb{E}_{\mathbf{z}_{t+1}(\mathbf{s}_{t+1})}\left[e(\mathbf{p}_{t+1}, \hat{\mathbf{p}}_{t+1})\right], \tag{3.36}$$

The estimated target position is denoted as $\hat{\mathbf{p}}$, while $e(\cdot,\cdot)$ is an estimation performance metric (e.g., Euclidean distance). The definition (3.36) is often impractical, since it may be difficult to find a meaningful metric $e(\cdot,\cdot)$ when there are multiple goals with different importance. The same reason holds when the goals are defined in different dimensions or state spaces Aoki, Bagchi, Mandal, and Boers (2011b). In addition, computing (3.36) incurs a significant computation cost, since it requires to integrate over all possible target positions and all possible sensor measurements Aoki et al. (2011b). For these reasons, definition (3.36) can be reformulated as Kreucher et al. (2005a)

$$J(\mathbf{u}_t \mid \mathbf{z}_{1:t}, \mathbf{s}_t) = \mathbb{E}_{\mathbf{z}_{t+1}(\mathbf{s}_{t+1})}\left[\max_{\mathbf{p}_{t+1}\in\Pi} p(\mathbf{p}_{t+1}|\mathbf{z}_{1:t}, \mathbf{z}_{t+1}(\mathbf{s}_{t+1}))\right], \tag{3.37}$$

which leads the platform to take the action that leads to the largest Maximum A-Posteriori (*MAP*) estimate of target location, that is the action that results in maximal certainty about target position.

In target search applications Liu et al. (2015), a commonly used task-driven criterion is the maximization of the *expected POD* over a given time horizon ν, that is Liu et al. (2015)

$$\begin{aligned} J(\mathbf{u}_t \mid \mathbf{z}_{1:t}, \mathbf{s}_t) &= \mathbb{E}_{\mathbf{p}_{t+1}}\left[p\left(\bigcup_{k=t+1}^{t+\nu} D_k(\mathbf{s}_{t+1})\,|D_{1:t} = \varnothing, \mathbf{z}_{1:t}\right)\right] \\ &= -\mathbb{E}_{\mathbf{p}_{t+1}}\left[\prod_{k=t+1}^{t+\nu} p\left(\bar{D}_k(\mathbf{s}_{t+1})\,|D_{1:t} = \varnothing, \mathbf{z}_{1:t}\right)\right]. \end{aligned} \tag{3.38}$$

In (3.38), $\bar{D}_t$ denotes a non-detection event at time t and $p\left(D_t|\bar{D}_{1:t-1}\right)$ is the probability of detecting the target at instant t, given the non-detection in previous observations. As can be seen from the second equation, each detection event is considered independent to the others. Traditionally, the calculation of (3.38) incurs high computation burden because no analytical form exists and numerical integration over large state space is required. Thus, ad-hoc approximation techniques have been proposed to efficiently solve POD-based search tasks also in real-time applications. For instance, by approximating the target belief map with a Gaussian Mixture Model (GMM) Liu et al. (2015) it is

possible to get a closed-form reformulation of (3.38). Alternatively, it is possible to assume the detection probability as constant during the time horizon $[t+1, t+\nu]$; in this way, the target search problem can be mathematically modeled as a Binomial experiment see Chapt. 5).

Minimum Time Search (MTS) Lanillos, Besada-Portas, Pajares, and Ruz (2012); Pérez-Carabaza et al. (2019a,b) is a task-driven criterion closely related to the POD function. In fact, it aims at minimizing the expected detection time, where the time of first detection (T_D) is modeled as a random variable; thus, the cost function in this case is Lanillos et al. (2012); Perez-Carabaza, Besada-Portas, Lopez-Orozco, and Jesus (2018); Pérez-Carabaza et al. (2019a)

$$\begin{aligned} J(\mathbf{u}_t \mid \mathbf{z}_{1:t}, \mathbf{s}_t) &= -\mathbb{E}\left[T_D\right] \\ &= -\sum_{k=t+1}^{\infty} \left[1 - p(T_D \leq k)\right] \\ &= -\sum_{k=t+1}^{\infty} k\, p\left(D_k(\mathbf{s}_{t+1}) | \bar{D}_{1:k-1}, \mathbf{z}_{1:t}\right) \end{aligned} \tag{3.39}$$

The definition in (3.39) is intractable; for this reason, a truncated time horizon is often considered in real-life applications Lanillos et al. (2012); Lanillos, Yañez-Zuluaga, Ruz, and Besada-Portas (2013); Perez-Carabaza et al. (2018). The main difference between the POD criterion (3.38) and the MTS one (3.39) stems on the fact that the former maximizes the chances of finding the target, but it does not guarantee to minimize the time to first detection; therefore, the POD criterion induces a slightly more explorative behavior during the search mission Lanillos et al. (2012).

Source seeking is another example of task-driven control, often accomplished via probabilistic approaches. In particular, a target-representative setpoint (also known as reference point) is extracted from the belief map; then, the platform is driven towards the setpoint. The MAP Radak et al. (2017) or the MMSE[15] Hasanzade et al. (2018) target position estimates are typical setpoint choices. In source seeking, the cost function is the distance between the current platform position[16] and the setpoint itself, namely

$$J(\mathbf{u}_t \mid \mathbf{z}_{1:t}, \mathbf{s}_t) = -\|\mathbf{s}_t - \hat{\mathbf{p}}_t\|_2. \tag{3.40}$$

This strategy is motivated by the assumption that moving closer to the target reduces the uncertainty of measurements Van Nguyen et al. (2020); however, it becomes inefficient if the belief map is characterized by multiple peaks Radmard and Croft (2017). Finally, when time of search matters, source seeking can be combined with MTS Liu et al. (2015).

To sum up, there is a certain lack of consistency, among the reviewed works, regarding the definition of task-driven control; however, all the proposed aforementioned options are valuable. Therefore, it is possible to conclude that a task-driven APE criterion can be simply defined as the opposite of an information-seeking criterion. In other words, a control law is defined as task-driven when the cost function is not information-theoretic

[15] MMSE: Minimum Mean Square Error.

[16] When the platform position is fixed and only rotation movements are allowed (e.g., pan cameras on a static robot), the distance is computed between the setpoint and the centre of the FoV projection onto the groundplane Π Chen et al. (2014)

and drives the platform closer to its task completion in an exploitative manner (e.g., by reaching a goal position within certain time/energy tolerances, minimizing the estimation error, maximizing the probability of target detection).

Hybrid control

Many practical situations require a trade-off between exploration and exploitation Gupta et al. (2006); Meera et al. (2019); Xu and Zhang (2014), as most APE tasks are designed to accurately reconstruct information about the target location, but not to produce overly time-consuming trajectories during the mission Vander Hook et al. (2015).

For this reason, some works propose hybrid control techniques, where information-seeking and task completion are balanced together. To do this, the most intuitive way is to cast the problem (3.28) into a *multi-objective optimization* framework Marler and Arora (2004), where the cost function is a linear combination between an information gathering measure and a task execution metric Liu and Hedrick (2017); Ramirez et al. (2014).

Alternative approaches cast the APE problem into a *Bayesian Optimization* framework Frazier (2018); Meera et al. (2019); Snoek et al. (2012), which is a powerful strategy for finding the extrema of an objective function in a gradient-free manner and when the cost function is expensive to evaluate, noisy, or non-convex Brochu et al. (2010). In Bayesian optimization, the objective function is approximated by a surrogate function (e.g., Gaussian Process Rasmussen (2003)); then, an acquisition function (e.g., Gaussian Confidence Upper Bound Bai et al. (2016), Expected Improvement Carpin et al. (2015)) chooses the next sensing action, usually managing the trade-off between exploration and exploitation Meera et al. (2019). The acquisition function can also be a linear combination between a source seeking term and an information gathering one Ghassemi and Chowdhury (2020); the former might account for the distance to the maximum expected target position, while the latter should depend on the target belief uncertainty. For a deeper overview on Gaussian Processes and Bayesian Optimization, see Appendix. A.1 and A.2.

Instead of combining exploration and exploitation in a unique cost function, it is also possible to use *switching controllers* to alternate between an information-seeking contribute (e.g., predictive variance) and a source-seeking one (e.g., distance to the Maximum Likelihood target estimate) Fink and Kumar (2010). In any case, switching controllers are liable to unstable behaviors; hence, when they are used, strict convergence guarantees should be provided.

Discussion

According to some previous works Kreucher et al. (2005a), task-driven control slightly outperforms information-seeking in terms of tracking error; in fact, information-based approaches sacrifice a small amount of tracking performance (i.e., estimation bias) to yield smaller estimation uncertainty. In general, explorative strategies are used when the priority is on uncertainty reduction, or if the APE task includes multiple objectives. On the contrary, exploitative strategies are specifically designed for specialized tasks (e.g., source seeking, detection time minimization). This usually requires ad-hoc utility

functions, which is the main limitation of task-driven controllers, as wrong choices may lead to local minima, as well as to unstable and ineffective platform behaviors Radmard and Croft (2017).

Challenges in task-driven control To increase the flexibility of pure-exploitative strategies, hybrid solutions can be adopted, even though in hybrid control the performance strongly depends on accurate exploitation-exploration balancing methods Ghassemi and Chowdhury (2020): if the controller is biased towards the source-seeking term, the estimation process might get stuck on local optima and the target may never be found; on the other hand, higher importance on the exploration term, leads to superior capabilities to explore the environment and, therefore, to find the target. However this may come at the cost of higher completion times, sometimes not compatible with the mission requirements Ghassemi and Chowdhury (2020). In conclusion, hyperparameter tuning is not only cumbersome and time-consuming, but also critical for the performance of hybrid APE solutions. For this reason, automatic hyperparameter tuning algorithms are currently an active research field in machine learning and robotics communities.

Challenges in information-seeking control While pure-exploitative strategies lack in flexibility, one of the main pitfalls of information-seeking approaches is the poor interpretability Aoki et al. (2011b), together with a limited capability in considering time and energy preservation policies. Furthermore, unlike task-based algorithms, information-seeking path planning is not adequate to real-time APE missions, where platforms have limited energy and computation power Van Nguyen et al. (2020). This is due to the intractability of the numerical integration methods involved in the calculation of most information-theoretic cost functions (e.g., Kullback-Liebler divergence, entropy, information gain) Hoffmann and Tomlin (2009). To reduce the computational burden, parametric Liu et al. (2015) and non-parametric Aughenbaugh and LaCour (2011); Czarnetzki et al. (2010); Ryan and Hedrick (2010) approximations are often applied; these, however, have an impact on the quality of the overall solution Hoffmann and Tomlin (2009) and restrict the range of techniques that can be used to solve the optimization problem (3.28). In particular, the lack of any analytical closed-form solution leads to the use of grid-search Ryan and Hedrick (2010), greedy Cliff et al. (2018) and meta-heuristic Pérez-Carabaza et al. (2019b) techniques, which are not always efficient and accurate Liu et al. (2015); Ramirez et al. (2014).

Relations between task-driven and information-seeking strategies The *near universal proxy* theory directly links any arbitrary task-driven function to the optimization of a criterion based on the Rény divergence Hero, Kreucher, and Blatt (2008). In particular, the near universal proxy theory claims that the expected value of any task-driven function is sandwiched between two marginalized Rènyi divergences; consequently, task-driven sensor management criteria could be replaced by a Rènyi divergence-based criterion without altering the performance guarantees. Even though the near universal proxy argument is the strongest theoretical justification until so far for the use of information-driven sensor management, it has been formally rebutted Aoki et al. (2011b); hence, there is still lack of a solid proof regarding the relationship between task-driven and information-driven sensor management.

On the opposite side of the near universal proxy theory, some works Katsilieris et al. (2012); Ryan (2008) have proven that it is possible to reformulate an information-theoretic

optimization problem into an equivalent task-driven one. This, in turn, is less costly and can sometimes be calculated more accurately; hence, reformulating information-seeking tasks into task-driven ones can solve the computational issues discussed in the previous paragraph. The main argument against these results is that they usually hold under specific conditions. For instance, the minimization of the expected entropy is equivalent to the maximization of the POD, if the POD is sufficiently small Ryan (2008). If instead the POD is constant, the optimization problem is equivalent to the maximization of the expected KL divergence (in the sense that both generate the same control input) Katsilieris et al. (2012). Note however that the POD is a joint function of the target state and the control input, as outlined in Sec. 3.1; hence, supposing this quantity to be constant over the entire $\Pi \times \mathcal{S}_m$ space is an ideal assumption, difficult to be met in real-life scenarios (see Fig. 3.2).

Future research lines In spite of the presence of some comparative studies Aoki et al. (2011a); Katsilieris et al. (2012); Kreucher et al. (2005a); Van Nguyen et al. (2020), the relation between task-driven and information-seeking approaches is still an open debate in APE-related research fields. Future research should aim to develop a generalized and unified guideline for the choice of the APE objective function to be used in (3.28). This should also account for the environment conditions, the number of cooperating platforms and their available resources (e.g., on-board computational and power capabilities), the platforms constraints, and for any specific task requirement (e.g., minimum time search). Another possible future trend in APE literature is the design of *Minimum Energy Search* (MES) strategies, which is still an unexplored line of research. Albeit MES might share some similarities with MTS, it is important not to confuse the two approaches: as discussed in Sec. 3.4.4, MTS generates control inputs that minimize the expected detection time (3.39), but no guarantees are provided on the platform energy-consumption, which is an important limitation in real-life applications.

3.5 Conclusion

In this Chapter, the main works in the area of Active Position Estimation have been organized, critically discussed and compared, with a focus on perception, control and multi-agent cooperation. The review identifies the main research challenges and future directions for the APE community. To sum up, multi-modal approaches are considered as one of the most promising future directions, since the majority of the existing works use uni-modal sensing strategies. In particular, it is important to study how different sensing modalities impact on search and tracking tasks. Future research should also be devoted to the comparison among different APE criteria, as the near-universal proxy argument is still an open discussion; furthermore, the existing literature does not clarify how to balance exploration with exploitation, especially under critical mission constraints. To this aim, hybrid control is indicated as the most promising research direction.

4

PAS for Localization: an application to Transmitter Discovery

This Chapter introduces a bi-Radio-Visual PAS scheme (Ra^2 ViPAS) to solve the transmitter discovery problem. Specifically, the correlation between radio and visual measurements is first exploited to learn a target detection model in a self-supervised manner. Then, the model is combined with antenna radiation anisotropies into a Bayesian Optimization framework that controls the platform. The proposed algorithm attains an accuracy of 92%*, overcoming two other probabilistic active sensing baselines.*

Contents

The contents of this chapter are partly available in: Varotto and Cenedese (2021c)
L. Varotto and A. Cenedese, "*Transmitter Discovery through Radio-Visual Probabilistic Active Sensing*". IEEE 25th International Conference on Methods and Models in Automation and Robotics (MMAR), 2021 [accepted]

4.1 Introduction

In the context of Multi-Modal PAS, *Radio-Visual PAS* (RaViPAS) approaches have been recently designed to mitigate the inadequacies of camera and radio-only strategies, especially in localization tasks de San Bernabé, Martinez-de Dios, and Ollero (2017). As discussed in Chapt. 3, the information richness of visual sensors may be impaired by occlusions and Field of View (FoV) directionality de San Bernabé et al. (2017). Radio frequency (RF) signals, instead, have wider ranges Zanella (2016) and enable energy efficient localization from the extraction of the Received Signal Strength Indicator (RSSI) of standard packet traffic Zanella (2016). Furthermore, RF communication has low hardware requirements and comes as parasitic in many real-life scenarios, since most portable devices are WiFi or Bluetooth enabled. Nonetheless, environmental interference often limits the RF-based localization accuracy Zanella (2016).

RaViPAS aims to combine the complementary benefits of RF signals with visual cues. The literature addressing radio-visual sensor fusion is still sparse and the RaViPAS framework is an open research field with methodological challenges and application opportunities: for example, energy-aware strategies are required to alternate the accurate but energy-harvesting camera measurements, with the rough but lighter radio observations (see Chapt. 6). Furthermore, traditional RF-based solutions involve tiring human-labeled calibration processes Zanella (2016).

Contribution This Chapter proposes a RaViPAS scheme to solve the *transmitter discovery* problem hereafter described (Fig. 4.1), which, despite its apparent simplicity, remains a canonical problem in many application domains. A static radio-visual sensing platform is surrounded by targets whose number and location is initially unknown. Only one of them establishes communication with the platform through a radio transmitter (Tx) and the aim is to identify which target is the Tx. The proposed scheme is based on the design of a Bayesian Optimization (BO) controller, leveraging the non-isotropic antenna radiation pattern at the receiver side. The exploration-exploitation trade off capabilities of BO are crucial for robust and efficient target localization Meera et al. (2019). In addition, RSSI-based localization techniques hinging on BO Carpin et al. (2015), do not require any observation model. Hence, the tiring human-labeled calibration processes involved in traditional RF-based solutions can be avoided. To increase the robustness of the localization process, a target visual detectability model is included into the BO framework. To this aim, it is possible to devise a self-supervised training procedure, exploiting the correlation between radio-visual inputs. In this way, the dataset acquisition process is automatized, with consequent minimization of the human intervention. In accordance with the PAS paradigm, the platform is guided by a probabilistic reward function, which is refined as new observations are collected. Numerical results show that the proposed algorithm overcomes a RF-only PAS scheme.

4.2 Problem Statement

Fig. 4.1 shows the main elements of the problem scenario, namely the targets and the sensing platform.

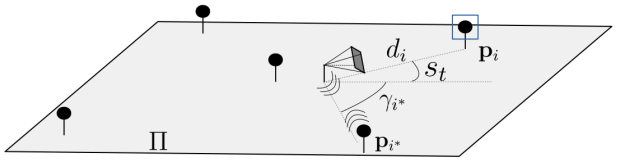

Figure 4.1. Problem scenario. Five targets (black markers) surround the sensing platform over the groundplane Π. The platform communicates through radio signals with one of the targets (Tx), placed at $\mathbf{p}_{i^*}$ and with angular coordinate γ_{i^*}. The platform has orientation s_t and it is detecting a non-Tx target (blue box). The Tx target can be detected only when inside the camera FoV, that is if $s_t \approx \gamma_{i^*}$.

4.2.1 Sensing platform

The sensing platform is a static orientable smart camera, endowed with a real-time target detector Redmon et al. (2016) and two radio receivers[1].

The camera Let $\mathcal{F}_0$ be the global reference frame, whose origin is centered in the platform position[2] and whose axis Z_0 is orthogonal to the groundplane $\Pi \subset \mathbb{R}^2$ where the platform and the targets lie. The platform state is identified with its only degree of freedom, that is the camera pan angle (i.e., the orientation around Z_0)

$$s_t \in [-\pi, \pi] = \mathcal{S}, \quad t = LT_c,\ L \in \mathbb{N}, \tag{4.1}$$

where T_c is the camera frame rate. The camera pan angle defines the FoV, $\Phi(s_t) \subset \Pi$, which is regulated through the control input u_t according to a deterministic Markovian transition model[3]

$$s_{t+T_c} = q(s_t, u_t) = s_t + u_t, \quad u_t \in [-\pi, \pi] = \mathcal{A}. \tag{4.2}$$

The state transitions occur at multiple of $T > T_c$, namely

$$u_t = 0, \quad t \neq HT,\ H \in \mathbb{N}. \tag{4.3}$$

The receivers The platform is endowed with two radio channels. The former, referred to as $\text{Rx}^{(iso)}$, uses an omnidirectional antenna, while the latter, $\text{Rx}^{(dir)}$, has a directional antenna (see Fig. 4.2) Gómez-Tornero, Cañete-Rebenaque, López-Pastor, and Martínez-Sala (2018). The dominating lobe of $\text{Rx}^{(dir)}$ is supposed to be aligned with the camera

[1] https://www.nordicsemi.com/-/media/Software-and-other-downloads/Product-Briefs/nRF52832-product-brief.pdf

[2] Formally, the platform position in Π is $\mathbf{c} = \begin{bmatrix} 0 & 0 \end{bmatrix}^\top$.

[3] s_{t+T_c} is forced to stay within $\mathcal{S}$ by adding (resp., subtracting) 2π if $s_{t+T_c} < -\pi$ (resp, $s_{t+T_c} > \pi$).

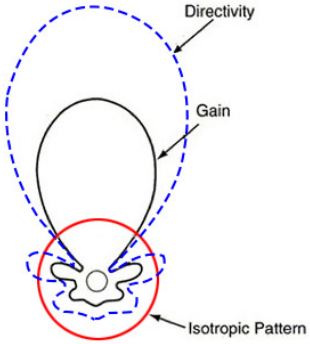

Figure 4.2. Example of radiation pattern of an isotropic (red) and a directional antenna (blue dashed).

optical axis, but the knowledge on the overall radiation pattern is inaccurate. It holds

$$T = T_{\mathrm{RF}} = \nu T_c,\ \nu \in \mathbb{N}, \tag{4.4}$$

where T_{RF} is the sampling interval of both receivers. With the first equality ($T = T_{\mathrm{RF}}$), a control input is computed as a new RSSI sample is collected. The second equality ($T_{\mathrm{RF}} = \nu T_c$) is justified by the fact that radio reception has typically longer sample rates than cameras' acquisition Vollmer and Möllmann (2011).

4.2.2 Targets

The platform is surrounded by N targets, which can be thought as people or objects, in civil or industrial settings. The targets are in Line-Of-Sight (LOS) w.r.t. the platform. In polar coordinates, the position of the i-th target w.r.t. $\mathcal{F}_0$ is denoted as

$$\mathbf{p}_i = \begin{bmatrix} d_i \cos\gamma_i & d_i \sin\gamma_i \end{bmatrix}^\top \in \Pi, \tag{4.5}$$

where $d_i = \|\mathbf{p}_i\|_2$ is the distance of the i-th target to the platform (i.e., Euclidean norm of $\mathbf{p}_i$), and γ_i is the target angular coordinate w.r.t $\mathcal{F}_0$, that is

$$\gamma_i = \arctan\left(\frac{p_{i,Y_0}}{p_{i,X_0}}\right) \tag{4.6}$$

with p_{i,X_0} and p_{i,Y_0} projections of $\mathbf{p}_i$ along X_0 and Y_0, respectively.

The number and the locations of the targets are unknown to the platform, but all N targets can be recognized through the camera object detector; moreover, one of the N targets establishes communication with the platform through an omnidirectional radio Tx. The index of such target is denoted with $i^* \in \{1,\ldots,N\}$.An industrial example of such a scenario comes from logistics and smart warehouse systems (see Fig. 4.3): a static pan camera monitors the surrounding racks where there are N marked boxes, such that they can be recognized by the camera. At each box location there is a radio Tx installed, but only one is active and can communicate with the camera; the camera should identify and localize, among the N locations, which one corresponds to the Tx; in this way, it can

focus on the Tx to perform a specific visual inspection.

Figure 4.3. Example of radio-enabled smart warehouse environment.

4.2.3 Perception modeling

Radio perception Both platform receivers (i.e., $\text{Rx}^{(dir)}$ and $\text{Rx}^{(iso)}$) extract from received data packets the Received Signal Strength Indicator (RSSI), $r \in \mathbb{R}$. As mentioned in Chapt. 3, the RSSI is a measure of the received radio signal power and, therefore, it quantifies the quality of communication between two radio-enabled devices; the RSSI is theoretically related to the Tx-Rx distance d_{i^*} Zanella (2016); thus, for an omnidirectional receiver the following deterministic relation holds

$$r^{(iso)} = g(d_{i^*}), \tag{4.7}$$

where $r^{(iso)}$ denotes the RSSI collected at $\text{Rx}^{(iso)}$. The relation (4.7) is not sufficient to capture the quality of communication decay due to the non-isotropic pattern of $\text{Rx}^{(dir)}$; hence, (4.7) should be modified as

$$r^{(dir)} = g(d_{i^*})\varrho(s_t, \gamma_{i^*}),\ \ \varrho(s_t, \gamma_{i^*}) \in [0, 1]. \tag{4.8}$$

The function $\varrho(\cdot)$ is an attenuation factor that models the effects of the radiation pattern anisotropy in the RSSI signal and it depends on the misalignment between Tx and $\text{Rx}^{(dir)}$: the attenuation is minimal if the domninating lobe points towards $\mathbf{p}_{i^*}$, namely $\varrho(s_t, \gamma_{i^*}) \approx 1$ (no attenuation), if $s_t \approx \gamma_{i^*}$; the attenuation increases as the misalignment between Tx and $\text{Rx}^{(dir)}$ grows and it becomes maximal (i.e., $\varrho(s_t, \gamma_{i^*}) \approx 0$) when the dominating lobe points in the opposite direction w.r.t. $\mathbf{p}_{i^*}$ Gómez-Tornero et al. (2018). Finally, it is worth noting that both $r^{(iso)}$ and $r^{(dir)}$ are functions of $\mathbf{p}_{i^*}$, the former through d_{i^*}, and the latter also through γ_{i^*}. For this reason, a generic RSSI value can be expressed as $r(\mathbf{p}_{i^*})$. Moreover, $r^{(dir)}$ is also function of s_t through the attenuation function $\varrho(\cdot)$.

Motivated by the energy efficiency and the low-cost hardware requirements, RSSI-based localization systems have been widely studied in literature de San Bernabé et al. (2017). Despite this, RSSI localization suffers from several drawbacks.

1. At first, the functions $g(\cdot)$ and $\varrho(\cdot)$ are usually unknown, are hardware-dependent, and need to be estimated through extensive and time-consuming calibration pro-

cedures Zanella (2016). Hence, in this Chapter a self-supervised methodology is suggested; it requires minimal human intervention and does not need to estimate neither $g(\cdot)$, nor $\varrho(\cdot)$ (see Sec. 4.3).

2. A further limitation of RF-only localization systems resides on the high sensitivity to environmental interference (e.g., cluttering and multi-path distortions) Zanella (2016). The target position is often hidden in extremely noisy receiver measurements, according to the following *RSSI observation model*

$$z_{\mathrm{RF},t} = \begin{cases} h(\mathbf{p}_{i^*}, s_t, v_{\mathrm{RF},t}) = r(\mathbf{p}_{i^*}) + v_{\mathrm{RF},t}, & t = MT_{\mathrm{RF}},\ M \in \mathbb{N} \\ \varnothing, & \text{otherwise} \end{cases} \tag{4.9}$$

where $\varnothing$ is an empty observation (i.e., measurements without target information); $z_{\mathrm{RF},t}$ is denoted as either $z^{(iso)}_{\mathrm{RF},t}$ or $z^{(dir)}_{\mathrm{RF},t}$, depending on whether $r(\mathbf{p}_{i^*})$ follows (4.7) or (4.8), respectively. Furthermore, the dependence on the platform state holds only for $\mathrm{Rx}^{(dir)}$. Finally, $v_{\mathrm{RF},t} \sim \mathcal{N}(v|0, \sigma^2_{\mathrm{RF}})$ is the RSSI measurements (zero-mean, Gaussian) noise, with σ_{RF} usually high w.r.t. the range of values spanned by $r(\mathbf{p}_{i^*})$; in real-life applications, the following relation is often satisfied

$$\frac{\sigma_{\mathrm{RF}}}{max\{r(\mathbf{p}_{i^*})\} - min\{r(\mathbf{p}_{i^*})\}} \geq 0.1. \tag{4.10}$$

3. A final issue of RSSI-based localization regards those techniques relying on omnidirectional receivers only Shahidian and Soltanizadeh (2017). As (4.7) highlights, the RSSI values are characterized by inherent ambiguities on the target position (i.e., they do not identify uniquely the position of the radio source). When the receiver is static, these ambiguities induce severe convergence issues in the estimation process (see Chapt. 5 for further details). For this reason, most literature solutions mitigate the ambiguity effects through multiple receivers Cantón Paterna et al. (2017b), multi-modal (e.g., radio-visual) perception systems Varotto et al. (2021), or active sensing schemes Shahidian and Soltanizadeh (2017).

Visual perception As for the camera sensing capability, the visual detection process of target i is modeled as a Bernoulli random variable $D_t \in \{0, 1\}$, with success probability

$$p(D_t = 1|\mathbf{p}_i, s_t) = \begin{cases} P_D(d_i), & \text{if } \mathbf{p}_i \in \Phi(s_t) \\ 0, & \text{otherwise.} \end{cases} \tag{4.11}$$

In line with the discussion of Chapt. 3, the detection success probability is a function of both the camera orientation and the target-platform distance (see Fig. 4.4). In particular, visual depth effects Radmard and Croft (2017) are accounted through $P_D(d_i)$, which is the *Probability of Detection* (POD) when the target is in FoV. Notably, the same POD function can be applied to all targets, which is a reasonable assumption in most practical scenarios where targets are characterized by similar appearance. If $i = i^*$, the dependence on d_i in (4.11) can be equivalently substituted by $r^{(iso)}$, on the basis of (4.7). More specifically,

$$P_D(d_{i^*}) = P_D\left(g^{-1}(r^{(iso)})\right) := p_D\left(r^{(iso)}\right). \tag{4.12}$$

In the following, to ease the notation, $r^{(iso)}$ is confused with r when it is used as argument of the POD function.

Figure 4.4. Example of detection model as function of the distance Radmard and Croft (2017)

.

4.2.4 Problem statement

The Tx discovery problem aims at identifying the Tx among the N targets. The LOS condition implies

$$\nexists(i,j) \text{ s.t. } \gamma_i = \gamma_j, \tag{4.13}$$

that is, each target is uniquely identified by its angular position w.r.t. the platform. Hence, the Tx discovery problem boils down to an *association* task on top of a uni-dimensional *localization* problem within a *noisy* scenario. The quantity to be estimated is the Tx angular position γ_{i^*}, assuming that the dynamics of the targets is slow w.r.t. the localization process. Formally, at time t, the result of the association task is

$$\hat{i}_t = \underset{i \in \{1,\dots,N\}}{\arg\min} |\hat{\gamma}_t - \gamma_i|, \tag{4.14}$$

where $\hat{\gamma}_t$ is the estimate of γ_{i^*} (i.e., the localization outcome), given camera and receivers observations up to time t.

4.3 Methodology

In spite of its simplicity, the Tx discovery problem is ill-posed if tackled with passive strategies, due to the inherent RSSI ambiguities (see Chapt. 5) and the uncertainties on the receivers radiation pattern; thus motivated, an active sensing scheme is here proposed.

Theoretically, the anisotropy of $\text{Rx}^{(dir)}$, coupled with the platform movements, provide sufficient information to reduce the RSSI ambiguities and solve the localization problem. Nonetheless, in practical scenarios the noise in RSSI data affect the stability and reliability of RF-only strategies Dore, Cattoni, and Regazzoni (2007b). Moreover, the lack of knowledge on the antenna radiation pattern often demands extensive calibration procedures Zanella (2016). The localization robustness can be increased by combining radio and visual cues. In particular, the problem is here formulated in a Bayesian probabilistic framework, which accounts for perception uncertainties and increases the adaptivity properties. This leads to the bi-Radio-Visual Probabilistic Active Sensing (Ra^2ViPAS) scheme depicted in Fig. 4.5 and detailed in the following. The Ra^2ViPAS is based on two steps:

1. first, given a set of offline RSSI and camera observations, Gaussian Process Regres-

sion (GPR) is applied to learn, in a self-supervised fashion, the POD as function of the RSSI value, that is $p_D(r^{(iso)})$[4] (Sec. 4.3.1);

2. then, an active localization task is formulated, where the platform control is cast to a black-box optimization problem for which Bayesian Optimization (BO) is employed (Sec. 4.3.2).

4.3.1 POD learning through GPR

This Section is devoted to the first part of Ra^{2}ViPAS, namely the self-supervised POD learning procedure through Gaussian Process Regression.

For a brief overview on Gaussian Processes and Gaussian Process Regression the reader is referred to Appendix A.1. In this work, GPR is employed to learn the POD function $p_D(r)$, on the basis of (4.12); to this aim, the POD is modeled as a GP, namely[5]

$$\widehat{p}_D(r) \sim \mathcal{GP}\left(0, k(r, r')\right). \tag{4.15}$$

The train dataset is

$$\mathcal{D} = \left\{ \left(z^{(iso)}_{\mathrm{RF},j}, \widetilde{p}_{D,j} \right) \right\}_{j=1}^{n_{train}} \tag{4.16}$$

where the inputs $\{z^{(iso)}_{\mathrm{RF},j}\}_{j=1}^{n_{train}}$ are RSSI observations from $\mathrm{Rx}^{(iso)}$, according to the observation model (4.9); the labels $\{\widetilde{p}_{D,j}\}_{j=1}^{n_{train}}$ represent the *empirical POD* between

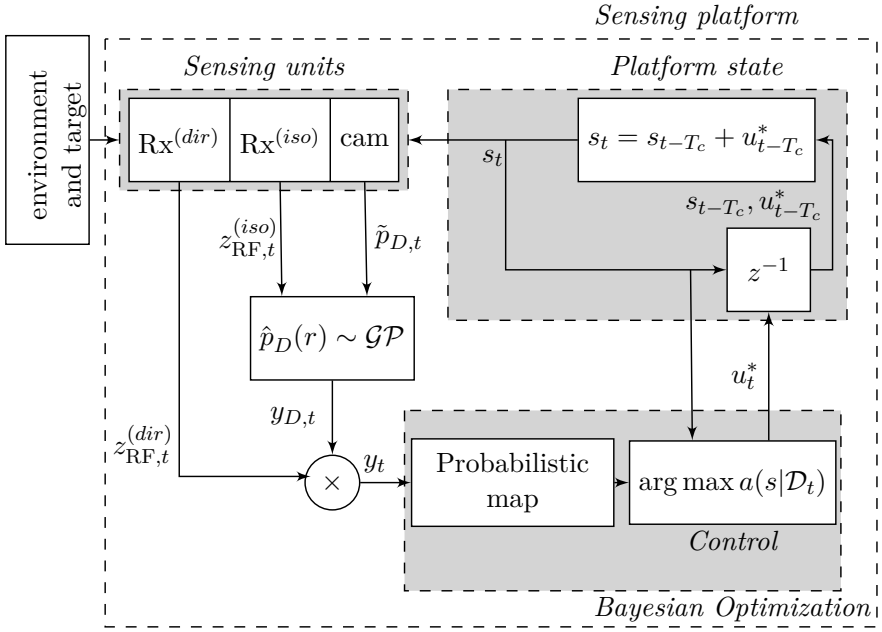

Figure 4.5. Ra^{2}ViPAS scheme for the Tx discovery problem.

[4] To ease notation, $p_D(r^{(iso)})$ is also expressed as $p_D(r)$.
[5] The notation $\widehat{p}_D(r)$ is used to distinguish the GP model with the underlying POD, $p_D(r)$.

two RSSI samples, according to

$$\widetilde{p}_{D,j} = \frac{1}{\nu}\sum_{\ell=1}^{\nu} D_{j,\ell}. \tag{4.17}$$

The equation above exploits the relation (4.4), regarding the ratio between the camera and the receiver sampling periods; $\{D_{j,\ell}\}_{\ell=1}^{\nu}$ are the ν detection outcomes before $z_{\mathrm{RF},j}^{(iso)}$ is collected. Notably, $\mathcal{D}$ is automatically acquired by the platform; in particular, neither sensors calibration nor human labeling are required, which is one of the main contributions of the Ra^{2}ViPAS methodology. What is instead required is human collaboration during the data acquisition process: to generate $\mathcal{D}$, the camera observes a target moving back and forth in front of it; at the same time, the receiver collects the RSSI data from the data packets sent by the Tx-equipped target. Clearly, to facilitate the data association process (i.e., each RSSI sample must be associated with ν camera frames), the camera and the receiver must be synchronized.

It is reasonable to assume that T_{RF} is sufficiently small w.r.t. the target movements, and $g(\cdot)$ has a low sensitivity w.r.t. the distance (i.e., the derivative is bounded). Under these hypotheses, the distance traveled between two consecutive RSSI samples (say, r_j and r_{j+1}) is negligible; hence, from (4.7), the POD remains approximately constant (see Fig. 4.4), that is $p_D(r_j) \approx p_D(r_{j+1})$. By consequence, the sequence of detection events $\{D_{j,\ell}\}_{\ell=1}^{\nu}$ is composed by i.i.d. random variables with distribution $\mathcal{B}\left(p_D(r_j)\right)$, according to (4.7) and (4.11). Then, from the Central Limit Theorem, $\widetilde{p}_{D,j}$ converges in distribution to

$$\mathcal{N}\left(\widetilde{p}_{D,j}\,\middle|\,p_D(r_j), \frac{p_D(r_j)(1-p_D(r_j))}{\nu}\right), \text{ as } \nu \to \infty, \tag{4.18}$$

where condition $\nu \to \infty$ is in practice satisfied for a high frame rate camera Vollmer and Möllmann (2011). From (4.18), the latent function $p_D(r)$ is related to the noisy labels $\widetilde{p}_D$ through the generative model

$$\begin{aligned} \widetilde{p}_{D,j} &= p_D(r_j) + \epsilon_j \\ \epsilon_j &\sim \mathcal{N}\left(\epsilon\,\middle|\,0, \frac{p_D(r_j)(1-p_D(r_j))}{\nu}\right), \end{aligned} \tag{4.19}$$

and the GPR problem is well-posed[6].

Fig. 4.6 shows an example of POD learning through GPR over synthetic data (see Sec. 4.4 for the data generation process). The true POD function is

$$P_D(d) = \left[\left(1+e^{4(d-4.5)}\right)\left(1+e^{-(d-2.5)}\right)\right]^{-1}. \tag{4.20}$$

The GP model uses a zero mean function and a Matern covariance (with length scale hyperparameter set to 1 and smoothness hyperparameter set to 1.5). The regression

[6] As (4.19) highlights, the generative model of the labels is input-dependent (i.e., the GP is heteroscedastic McHutchon and Rasmussen (2011)). Moreover, the GPR works on noisy inputs, according to (4.9). To deal with heteroscedasticity and noisy inputs, several ad-hoc methods have been proposed in literature Kersting, Plagemann, Pfaff, and Burgard (2007); McHutchon and Rasmussen (2011).

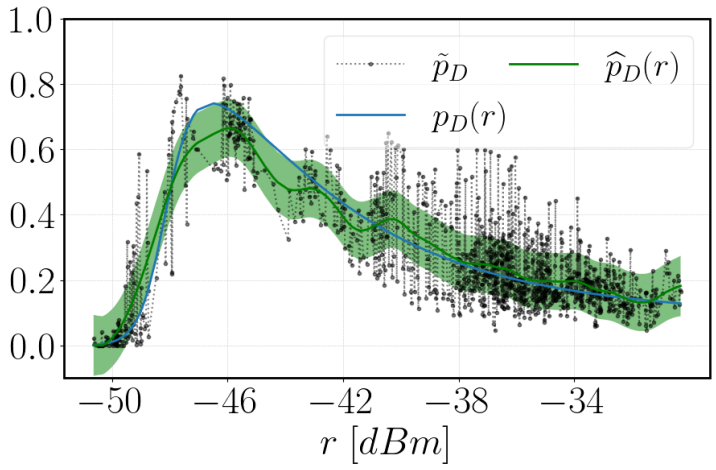

Figure 4.6. Comparison between the latent function $p_D(r)$ (blue line) and the GP model $\widehat{p}_D(r)$ (green, average and 68% confidence interval), learnt on the very noisy synthetic data $\{z^{(iso)}_{\mathrm{RF},j}, \widetilde{p}_{D,j}\}_{j=1}^{n_{train}}$ (black markers).

performance is evaluated through the coefficient of determination $R^2 \in (-\infty, 1]$

$$R^2 = 1 - \frac{\sum_{j=1}^{n_{train}} \left(p_D(r_j) - \widehat{p}_{D,*}(r_j)\right)^2}{\sum_{j=1}^{n_{train}} \left(p_D(r_j) - \frac{1}{n_{train}} \sum_{k=1}^{n_{train}} p_D(r_k)\right)^2} \tag{4.21}$$

where $\widehat{p}_{D,*}(r_j)$ is the POD prediction, according to (A.8). The higher R^2, the better the fit; in the case of Fig. 4.6, $R^2 = 0.94$.

4.3.2 Ra²ViPAS for Transmitter Discovery

In this Section, the second part of Ra²ViPAS is explained, namely the active localization scheme. Fig. 4.5 shows the Ra²ViPAS pipeline: the information gathered by the platform sensing channels is aggregated and injected into a BO scheme; this reconstructs a probabilistic approximation of the localization objective function in (4.14), which is then used to generate platform control inputs, according to a probabilistic controller.

As already discussed in Chapt. 3, Probabilistic Active Localization hinges on a belief map to encode the knowledge on potential target locations and plan future actions accordingly. If the likelihood function is known, Recursive Bayesian Estimation can be applied to keep the map updated as new observations are collected (see Appendix A.3). However, in our framework the likelihood function $p(D_t, z_{\mathrm{RF},t}|\mathbf{p}_{T_{\mathrm{RF}:t}})$ is unknown; hence, RBE techniques can not be applied to compute $\hat{\gamma}_t$. On the contrary, it is possible to cast the localization problem defined in Sec. 4.2 to a suitable active sensing scheme where the platform state is regulated via black-box optimization. To this aim, it is important to notice that the Tx location satisfies the following relation

$$\begin{aligned} \gamma_{i^*} &= \arg\max_{\gamma \in [-\pi,\pi]} \underbrace{J_D(\gamma) J_{\mathrm{RF}}(\gamma)}_{J(\gamma)} \\ J_D(\gamma) &= -|p(D|\mathbf{p}_{i^*}, \gamma_{i^*}) - p(D|\mathbf{p}_{i^*}, \gamma)| \\ J_{\mathrm{RF}}(\gamma) &= \varrho(\gamma, \gamma_{i^*}). \end{aligned} \tag{4.22}$$

Specifically: the detectability term, $J_D(\gamma)$, favors those values of γ where the target has the same (measured) POD of the Tx; on the other side, the RF term, $J_{\text{RF}}(\gamma)$, accounts for the radiation pattern and gives higher rewards when $\text{Rx}^{(dir)}$ is aligned with Tx. The lack of precise knowledge on $\varrho(\cdot)$ makes (4.22) a black-box optimization problem on the cost function $J(\gamma)$; hence, it can be opportunely solved via BO. Recalling that the paltform state is the camera pan angle, it is possible to equivalently optimize over the platform state s_t; thus, at time t the dataset is

$$\mathcal{D}_t = \{(s_j, y_j)\}_{j=T_{\text{RF}}}^{t}, \tag{4.23}$$

with $y_j = y_{D,j} y_{\text{RF},j}$ and

$$\begin{aligned} y_{D,j} &= -|\hat{p}_{D,*}(z_{\text{RF},j}^{(iso)}) - \tilde{p}_{D,j}| \\ y_{\text{RF},j} &= |z_{\text{RF},j}^{(dir)}|/\zeta, \end{aligned} \tag{4.24}$$

where ζ is a user-defined scaling hyper-parameter. Note that, $z_{\text{RF},t}^{(dir)}/\zeta$ is a (noisy) scaled version of $J_{\text{RF}}(\cdot)$, according to (4.9) and (4.8); however, this does not affect the optimization process, since the scale is $g(d_{i^*})$, which does not depend on s_t. Regarding $y_{D,t}$, it is a noisy version of $J_D(\cdot)$, where the noise comes from both the estimation error in $\hat{p}_{D,*}(z_{\text{RF},t}^{(iso)})$ and the measurements error in $\tilde{p}_{D,t}$ (see Sec. 4.3.1).

In this BO implementation, the surrogate model is a Gaussian Process and, according to the usual sequential BO procedure, $J(\gamma)$ is approximated via GPR with the dataset (4.23). The GP model that approximates $J(\gamma)$ represents the probabilistic map from which the next control input is computed (i.e., *probabilistic control*)

$$u_t^* = \begin{cases} \left[\arg\max_{s\in[-\pi,\pi]} a(s|\mathcal{D}_t)\right] - s_t, & t = HT_{\text{RF}} \\ 0, & \text{otherwise}, \end{cases} \tag{4.25}$$

where condition (4.3), with $T = T_{\text{RF}}$, is taken into account; moreover, the control input in (4.25) is the displacement between the optimal and the current camera orientation, that is $\arg\max_{s\in[-\pi,\pi]} a(s|\mathcal{D}_t)$ and s_t respectively. Finally, Ra^{2}ViPAS uses the GP Upper Confidence Bound (UCB) as acquisition function, namely

$$a(\mathbf{x}|\mathcal{D}_t) = \mu_{*,\mathcal{D}_t} + \sqrt{\beta}\sigma_{*,\mathcal{D}_t} \tag{4.26}$$

where (A.8) is applied, $\sigma_{*,\mathcal{D}_t}$ is $\mathbf{\Sigma}_{*,\mathcal{D}_t}$ with $n_{test} = 1$, and β is an exploration-exploitation tuning hyperparameter.

The estimate of γ_{i^*} at time t is the maximum of the GP predictive posterior mean (A.8), namely

$$\hat{\gamma}_t = \arg\max_{\gamma\in[-\pi,\pi]} \mu_{*|\mathcal{D}_t}(\gamma) \tag{4.27}$$

4.4 Numerical Results

The proposed approach has been evaluated through a Python-based synthetic environment[7]. Sec. 4.4.1 describes the main setup parameters, as well as the synthetic data generation process. Sec. 4.4.2 defines the metrics used for performance assessment and the baselines considered for comparison. The numerical simulation results are discussed in Sec. 4.4.3.

4.4.1 Setup parameters

To carry out realistic simulations, most parameters reflect real device characteristics. Tx and $\text{Rx}^{(iso)}$ are supposed to be equipped with Nordic nRF52832 SoCs[8]; $\text{Rx}^{(dir)}$ differs from $\text{Rx}^{(iso)}$ in that the attenuation gain follows the function in Tab. 4.1, inspired by literature Gómez-Tornero et al. (2018). The function $g(\cdot)$ follows the log-distance Path Loss Model (PLM) with a noise level typical of industrial scenarios, that is $\sigma_{\text{RF}} = 3\,\text{dBm}$ Zanella (2016). The radio sampling time is set to $T_{\text{RF}} = 0.1\,\text{s}$, and the value of ν is set to 10, as a reasonable trade-off between typical real-life values and the ideal condition stated in (4.18). The underlying POD function is (4.20), designed according to real-life experiments Varotto et al. (2021). During the POD learning phase, the target is supposed to move according to the stochastic linear model reported in Tab. 4.1. Finally, the synthetic environment simulates a crowded environment with $N = 20$ targets.

Table 4.1. Setup parameters for the MC experiment.

Parameter	Value
T_{RF}	$0.1\,\text{s}$
ν	10
n_{train}	900
n_{test}	120
N_{tests}	50
N	20
$f(d_t, \eta_t)$	$d_t + \eta_t \quad \eta_t \sim \mathcal{N}(0, 0.04)$
$g(d_{i^*})$	$\kappa - 10n\log_{10}(d_t/\delta) \quad \kappa = -30\,\text{dBm},\ n = 2,\ \delta = 1\,\text{m}$
$\varrho(s_t, \gamma_{i^*})$	$1 - 0.5(s_t - \gamma_{i^*})^2$
σ_{RF}	$3\,\text{dBm}$
$P_D(d)$	$\left[\left(1 + e^{4(d-4.5)}\right)\left(1 + e^{-(d-2.5)}\right)\right]^{-1}$

4.4.2 Performance assessment

To capture the performance variability, numerical evaluation is performed through a Monte Carlo (MC) experiment, composed by $N_{tests} = 50$ tests of duration $T_W = n_{test}T_{\text{RF}}$ each, where $n_{test} = 120$ is the number of RSSI collected. At each MC test the position of

[7] https://github.com/luca-varotto/Tx-discovery

[8] https://www.nordicsemi.com/-/media/Software-and-other-downloads/Product-Briefs/nRF52832-product-brief.pdf.

the N targets is randomly chosen.

Performance metric

The performance of the Tx discovery task is evaluated through the *Discovery Rate*

$$DR_t = \frac{1}{N_{tests}} \sum_{j=1}^{N_{tests}} \mathbb{1}_{\hat{i}_t = i^*}, \quad \mathbb{1}_{\hat{i}_t = i^*} = \begin{cases} 1, & \text{if } \hat{i}_t = i^* \\ 0, & \text{otherwise,} \end{cases} \tag{4.28}$$

where $\hat{i}_t$ follows (4.14). From a probabilistic perspective, DR_t is the empirical probability of identifying Tx at time t over the MC tests.

Baselines

The following original baselines are considered, all leveraging the main PAS approach proposed here:

- *RaPAS*: same pipeline of Ra^{2}ViPAS, but only Rx$^{(dir)}$ is used, namely $y_t = y_{\text{RF},t}$, as depicted in Fig. 4.7(a).
- *RaViPAS*: same pipeline of Ra^{2}ViPAS, but Rx$^{(dir)}$ is not used namely $y_t = y_{D,t}$, as depicted in Fig. 4.7(b).

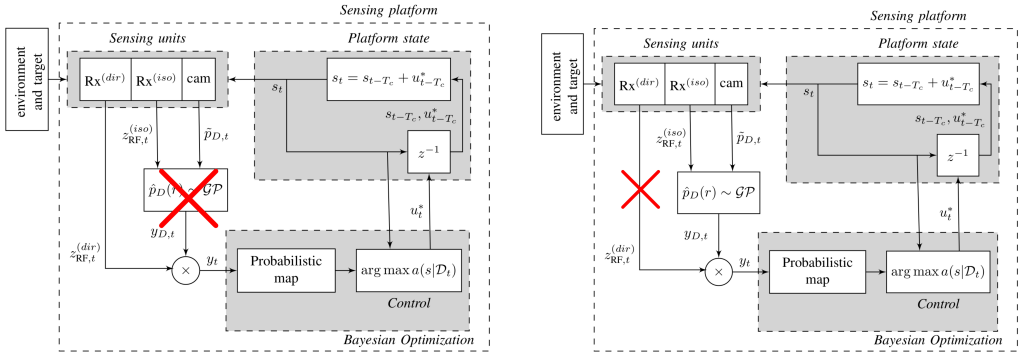

Figure 4.7. Comparison of the Ra^{2}ViPAS scheme with the two baselines used for performance assessment: (a) RaPAS; (b) RaViPAS. The red cross identifies the part of the Ra^{2}ViPAS scheme not used by that specific baseline.

4.4.3 Discussion

Fig. 4.8 depicts the discovery rate of Ra^{2}ViPAS, RaViPAS and RaPAS. The control law (4.25) induces the same tree-steps qualitative behavior in all three algorithms, even though with different final results.

1. Initially, the sparse sampling of the search domain $[-\pi, \pi]$ produces wrong localization solutions; therefore, DR is low.

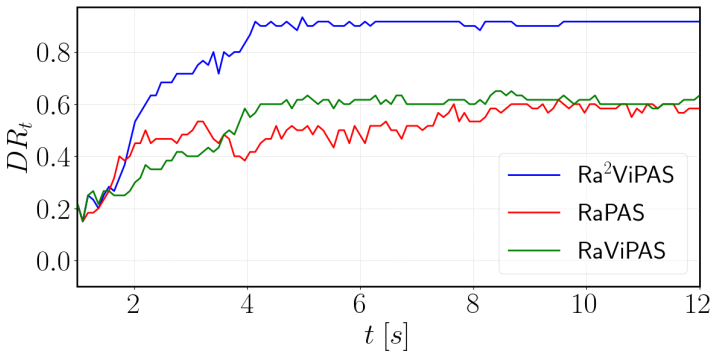

Figure 4.8. Discovery Rate of Ra^{2}ViPAS, RaPAS and RaViPAS over a MC experiment.

2. The sparse domain coverage implies large uncertainties in the surrogate model; this induces an explorative platform behavior, according to (4.26). The exploration process allows the platform to focus on regions that are more likely to optimize the objective function in (4.22); hence, DR increases remarkably in this phase.

3. When enough information is collected, exploitation overcomes exploration, according to the acquisition function (4.26); DR meets a converge.

Fig. 4.8 highlights the fact that the exploration phase of Ra^{2}ViPAS and RaViPAS lasts approximately half with respect to RaPAS ($\approx 4\,$s against $\approx 8\,$s); this is due to the smaller amount of information employed in RaPAS, and it is consistent with the advantages of multi-modal AS solutions discussed in Chapt. 3.

Another interesting feature of the results reported in Fig. 4.8 is that RaPAS and RaViPAS converge to the same value ($\approx 60\%$), while Ra^{2}ViPAS has the highest ($\approx 92\%$) and smoothest convergence behavior. This means that the information fusion process employed in Ra^{2}ViPAS produces a fast and efficient information gain, which leads to high estimation accuracy, stability and robustness.

In conclusion, the performance of Ra^{2}ViPAS justifies the higher hardware requirements w.r.t. RaPAS and RaViPAS.

4.5 Conclusion

This work proposes a probabilistic bi-radio-visual active sensing framework (Ra^{2}ViPAS), applied to the transmitter discovery problem. Target visual detectability and radio signal strength are combined into a Bayesian Optimization framework, responsible for the generation of the platform control movements. The detectability model is learnt in a self-supervised fashion, so that human intervention is minimized. The approach can be extended to various application domains and in the specifically considered scenario the suggested strategy attains a 92% accuracy level, 30% higher than the two baselines under comparison.

5

PAS for Tracking: an application to Target Visual Detection

In this Chapter the focus is on the visual detection problem of a radio-emitting target with an aerial robot, equipped with a radio receiver and a camera. A Recursive Bayesian Estimation scheme uses camera observations to refine radio measurements and generate a probabilistic map; on top of this, an optimal controller regulates the camera pose. Theoretical results support the proposed algorithm, while numerical analyses show higher robustness and efficiency with respect to visual and radio-only baselines. In particular, visual-based tracking often requires long search times before detecting the target, due to the directionality of the sensing domain; conversely, radio signals have larger coverage, but lower tracking accuracy.

Contents

The contents of this chapter are partly available in: Varotto et al. (2021)
L. Varotto, A. Cenedese and A. Cavallaro, "*Probabilistic Radio-Visual Active Sensing for Search and Tracking*". European Control Conference (ECC), 2021 [accepted]

5.1 Introduction

As discussed in Chapt. 3, typical modalities for active sensing include vision, audio and radio. Visual-based tracking provides high accuracy Aghajan and Cavallaro (2009) and does not require the target to use any emitting device. Nonetheless, occlusions and FoV directionality Mavrinac and Chen (2013) limit the range, applicability and success of camera-only platforms Robin and Lacroix (2016), especially for applications where time of detection is critical (e.g., search and rescue missions Murphy et al. (2019)). To collect measurements on wider ranges, and reduce the search phase, other technologies can be used, such as acoustic Haubner et al. (2019b) or radio-frequency (RF) Shahidian and Soltanizadeh (2017) signals. Despite the high localization accuracy of acoustic signals Haubner et al. (2019b), sound pollution and extra hardware requirements (e.g., microphone arrays) are critical pitfalls of this technology Zafari et al. (2019). Conversely, RF signals are energy efficient, have large reception ranges (~ 100 $[m]$) Siekkinen et al. (2012), and low hardware requirements, since platforms are not needed to be equipped with multiple receivers; moreover, the Received Signal Strength Indicator (RSSI) is extracted from standard data packet traffic (see Chapt. 3). For these reasons, RSSI-based localization systems widely appear in literature and in commercial applications, despite environmental interference (e.g., cluttering and multi-path distortions) often limits their accuracy Zanella (2016). In this context, multi-modal sensor fusion techniques have been shown to overcome the inadequacies of uni-modal approaches, being more robust and reliable Lathuilière et al. (2019).

Contribution The algorithm proposed in this Chapter exploits the complementary benefits of radio and visual cues for visually detecting a radio-emitting target with an aerial robot, equipped with a radio receiver and a Pan-Tilt (PT) camera. The control problem is formulated within a probabilistic active sensing framework, where camera measurements refine radio ones within a RBE scheme, used to keep the map updated. The fusion of RF and camera sensor data for target search and tracking is an open problem and the literature addressing this task is sparse. Indeed, this is the first attempt to combine radio and visual measurements within a probabilistic active sensing framework. Furthermore, unlike existing solutions operating on limited control spaces (e.g., platform position Shahidian and Soltanizadeh (2017) or camera orientation Lathuilière et al. (2019)), it is here proposed a gradient-based optimal control, defined on a continuous space comprising both platform position and camera orientation. Theoretical and numerical analyses are provided to validate the effectiveness of the proposed algorithm. What emerges is that bi-modality is proven to increase the target localization accuracy; this, together with the availability of an integrated high-dimensional control space, leads to higher detection success rates, as well as superior time and energy efficiency with respect to radio and vision-only counterparts.

5.2 Problem statement

Fig. 5.1 shows the main elements of the problem scenario, namely the target and the sensing platform[1].

5.2.1 Target

The radio-emitting target moves on a planar environment $\Pi \subset \mathbb{R}^2$, according to a (possibly) non-linear stochastic Markovian state transition model

$$\mathbf{p}_{t+1} = f(\mathbf{p}_t, \boldsymbol{\eta}_t);\ t \in \mathbb{N}. \tag{5.1}$$

where $\mathbf{p}_t \in \Pi$ is the target position at time t, referred to the global 3D reference frame $\mathcal{F}_0$; when expressed in $\mathbb{R}^3$, it is referred as $\mathbf{p}_t^+ = [\,\mathbf{p}_t^\top \quad 0\,]^\top$. The uncertainty on the underlying target movements are captured by the distribution of the stochastic process noise $\boldsymbol{\eta}_t$. The probabilistic form of (5.1), namely $p(\mathbf{p}_{t+1}|\mathbf{p}_t)$, is known as process model.

5.2.2 Sensing platform

The sensing platform is an aerial vehicle (UAV), equipped with an omnidirectional (isotropic) radio receiver and a PT camera on gimbal, and endowed with processing capabilities and a real-time target detector Redmon et al. (2016). The state of the platform is the camera pose, namely

$$\mathbf{s}_t = \begin{bmatrix} \mathbf{c}_t^\top & \boldsymbol{\psi}_t^\top \end{bmatrix}^\top, \tag{5.2}$$

with

$$\mathbf{c}_t \in \mathbb{R}^3; \quad \boldsymbol{\psi}_t = \begin{bmatrix} \alpha_t & \beta_t \end{bmatrix}^\top \in [-\pi/2 + \theta, \pi/2 - \theta]^2 \tag{5.3}$$

and where

$$[-\pi/2 + \theta, \pi/2 - \theta]^2 = [-\pi/2 + \theta, \pi/2 - \theta] \times [-\pi/2 + \theta, \pi/2 - \theta]. \tag{5.4}$$

The UAV position $\mathbf{c}_t$ is referred to $\mathcal{F}_0$, it is supposed to coincide with the camera focal point and its altitude is fixed (i.e., non-controllable); α_t (resp. β_t) is the pan (resp. tilt) angle w.r.t. the camera inertial reference frame $\mathcal{F}_u$ (obtained by a $\mathbf{c}_t$ translation of $\mathcal{F}_0$); θ is the half-angle of view. Similarly to Chapt. 4, the camera state follows a linear

[1] The Euclidean distance between vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$ is denoted as

$$d(\mathbf{a}, \mathbf{b}) = \|\mathbf{a} - \mathbf{b}\|_2 = \left(\sum_{i=1}^{n} (a(i) - b(i))^2 \right)^{1/2}$$

where $a(i)$ is the i-th component of $\mathbf{a}$.
Given a plane Π and a vector $\mathbf{a} \in \mathbb{R}^n$, the orthogonal projection of $\mathbf{a}$ onto Π is denoted as $\mathbf{a}_\Pi$.

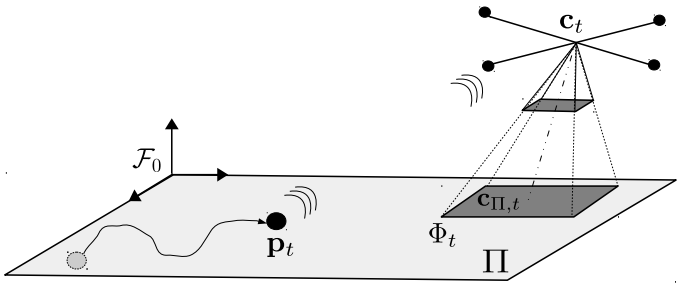

Figure 5.1. Problem scenario. A target moves in a planar environment and establishes a radio signal communication with a camera-embedded UAV. The objective is to control the camera pose so that the target is visually detected.

deterministic Markovian transition model

$$\begin{aligned} &\mathbf{s}_{t+1} = \mathbf{s}_t + \mathbf{u}_t; \\ &\mathbf{u}_t = \begin{bmatrix} \mathbf{u}_{\mathbf{c},t}^\top & \mathbf{u}_{\psi,t}^\top \end{bmatrix}^\top \in \mathcal{A} \end{aligned} \tag{5.5}$$

where $\mathcal{A}$ is the control space (see Chapt. 3). In particular, the control input $\mathbf{u}_{\psi,t} \in \mathbb{S}^2$ acts on the camera PT configuration, while $\mathbf{u}_{\mathbf{c},t} \in \mathbb{R}^3$ acts on the drone position in $\mathbb{R}^3$. However, being the UAV altitude fixed, the focus is only on the planar projection of $\mathbf{c}_t$ onto Π, namely $\mathbf{c}_{\Pi,t}$; therefore, as formally described in Sec. 5.3, $\mathbf{u}_{\mathbf{c},t}$ is composed by a planar control contribute (i.e., $\mathbf{u}_{\mathbf{c}_\Pi,t}$) in the first two entries, and a null contribute in the last entry. As (5.5) suggests, the state transitions occur at multiple of $T = 1$, where T is supposed to coincide with the camera frame rate T_c; moreover, as in Chapt. 4, the camera frame rate T_c is smaller w.r.t. the receiver sampling interval T_{RF}, namely

$$T_{\text{RF}} = \nu T_c, \ \nu \in \mathbb{N}. \tag{5.6}$$

Putting everything together, the chain-relation of the sampling rates introduced in (4.4), here becomes

$$T_{\text{RF}} = \nu T_c = \nu T, \ \nu \in \mathbb{N}. \tag{5.7}$$

Without any loss of generality, later on a normalized frame rate will be considered (i.e., $T_c = 1$). Notably, the controller works at the same frequency of the camera, and not of the receiver as in Chapt. 4 (i.e., the controller works at a higher frequency). As described in Sec. 5.3.1, this implies that some control inputs are computed according to a belief map that has been updated with visual information only.

Inspired by real-life scenarios, the UAV movements are considered energy-consuming with a linear dependence on the flying distance Liu, Chen, Tang, Xu, and Piao (2018), that is

$$\Delta E_t = d(\mathbf{c}_{t+1}, \mathbf{c}_t), \tag{5.8}$$

where ΔE_t is the energy used to move the platform from $\mathbf{c}_t$ to $\mathbf{c}_{t+1}$. The total available energy is denoted as E_{tot}.

5.2.3 Perception modeling

Radio perception Motivated by the long reception ranges of radio signals, the target is supposed to be always within the range of the platform receiver; moreover, the platform can extract the RSSI value, $r_t \in \mathbb{R}$, from received data packets. As already mentioned in Chapt. 3 and 5, the RSSI is related to the platform-target distance $d(\mathbf{c}_t, \mathbf{p}_t^+)$. In particular, the ground-to-air communication should minimize distortions and cluttering phenomena Koohifar et al. (2016); hence, in this Chapter, the unknown function $g(\cdot)$ of (4.7) is supposed to be known and equal to the log-distance path loss model Goldsmith (2005)

$$r_t = \kappa - 10n \log_{10} \left(d(\mathbf{c}_t, \mathbf{p}_t^+) \right). \tag{5.9}$$

The parameters κ and n are estimated via offline calibration procedures Zanella (2016), where the reference distance is $\delta = 1$ $[m]$. Thus, similarly to (4.9) in Chapt. 4, the radio observation model is

$$z_{\mathrm{RF},t} = \begin{cases} h_{\mathrm{RF}}(\mathbf{p}_t, \mathbf{s}_t, v_{\mathrm{RF},t}) = r_t + v_{\mathrm{RF},t}, & t = MT_{\mathrm{RF}},\ M \in \mathbb{N} \\ \varnothing, & \text{otherwise} \end{cases} \tag{5.10}$$

where T_{RF} is the receiver sampling interval and the dependence on the platform state and target position follows from (5.9); $v_{\mathrm{RF},t} \sim \mathcal{N}\left(v|0, \sigma_{\mathrm{RF}}^2\right)$ is the RSSI measurements noise.

Visual perception Camera measurements are modeled through the projection perspective geometry Aghajan and Cavallaro (2009)

$$\mathbf{z}_{c,t} = \begin{cases} h_c(\mathbf{p}_t, \mathbf{s}_t, \mathbf{v}_{c,t}) = \mathbf{P}(\mathbf{s}_t)\tilde{\mathbf{p}}_t + \mathbf{v}_{c,t}, & D_t = 1 \text{ and } t = NT_c,\ N \in \mathbb{N} \\ \varnothing, & \text{otherwise} \end{cases} \tag{5.11}$$

where $\mathbf{v}_t \sim \mathcal{N}\left(\mathbf{v}|\mathbf{0}_2, \mathbf{\Sigma}_c\right)$ is the noise of camera observations; $\tilde{\mathbf{p}}_t$ is the homogeneous representation of $\mathbf{p}_t^+$, namely

$$\tilde{\mathbf{p}}_t = \begin{bmatrix} \mathbf{p}_t^\top & 1 \end{bmatrix}^\top; \tag{5.12}$$

and

$$\mathbf{P}(\mathbf{s}_t) = \begin{bmatrix} \mathbf{I}_2 & \mathbf{0}_2 \end{bmatrix} \mathbf{K} \begin{bmatrix} \mathbf{R}(\boldsymbol{\psi}_t) & \mathbf{c}_t \end{bmatrix} \in \mathbb{R}^{2\times 4} \tag{5.13}$$

is the projection matrix that maps $\tilde{\mathbf{p}}_t$ onto the camera image plane. It depends on $\mathbf{K} \in \mathbb{R}^{3\times 3}$, the matrix of intrinsic parameters, and $\mathbf{R}(\boldsymbol{\psi}_t)$, the camera rotation matrix w.r.t. $\mathcal{F}_0$.

Finally, the target visual detection process is modeled as a Bernoulli random variable $D_t \in \{0, 1\}$, with success probability

$$p(D_t = 1|\mathbf{p}_t, \mathbf{s}_t) = \begin{cases} \Upsilon(\mathbf{p}_t, \mathbf{s}_t), & \text{if } \mathbf{p}_t \in \Phi(\mathbf{s}_t) \\ 0, & \text{otherwise} \end{cases} \tag{5.14}$$

with $\Phi(\mathbf{s}_t)$ camera FoV onto Π (see Fig. 5.1), and

$$\Upsilon(\mathbf{p}_t, \mathbf{s}_t) = \left[1 + e^{\gamma(d(\mathbf{c}_t, \mathbf{p}_t^+)/f - \epsilon)}\right]^{-1} \left(1 + e^{-\gamma\epsilon}\right). \tag{5.15}$$

Thus, from (5.14), the target can be detected only if inside the camera FoV, while from (5.15), the detection probability is proportional to the quantity $d(\mathbf{c}_t, \mathbf{p}_t^+)/f$; this is the resolution at which the target is observed, where f is the camera focal length. In (5.14), the success probability depends also on some constant values, such as the resolution at which the target is no longer well detectable (i.e., $\epsilon > 0$), and the spatial rate of decrease of the target detectability (i.e., $\gamma > 0$). Notably, $\Upsilon(\cdot)$ is a realistic POD model, as proven by literature Radmard and Croft (2017) and by field results (see Chapt. 6); moreover, this model is consistent with the definition provided in Chapt. 3 and it is very similar to the model used to generate the synthetic data in Sec. 4.4 in Chapt. 4.

5.2.4 Problem statement

With the formalism introduced in the previous Sections, the *visual target detection problem* can be formulated as the control of the platform state $\mathbf{s}_t$ (through $\mathbf{u}_t$) to realize event $D_t = 1$, given camera and receiver observations $\mathbf{z}_{c,T_c:t}$ and $z_{\mathrm{RF},T_c:t}$.

5.3 Methodology

To solve the problem defined in Sec. 5.2.4, a probabilistic bi-modal active sensing approach is proposed. As shown in Fig. 5.2, radio-visual measurements are aggregated to form a bi-modal likelihood function; this is used to keep the target belief map updated through an RBE scheme. Finally, an optimal controller is fed with the probabilistic map and generates the platform control input.

5.3.1 Bi-modal likelihood and probabilistic map

In active sensing frameworks, the configuration of the sensing robot is not constant; hence, the likelihood function $p(\mathbf{z}_t|\mathbf{p}_t)$ should account also for the platform state, that is $p(\mathbf{z}_t|\mathbf{p}_t, \mathbf{s}_t)$. Given this premise, and recalling the stochastic characterization of the radio observation model (5.10), the *RF likelihood* is

$$p(z_{\mathrm{RF},t}|\mathbf{p}_t, \mathbf{s}_t) = \begin{cases} \mathcal{N}\left(z|r_t, \sigma_{\mathrm{RF}}^2\right), & t = MT_{\mathrm{RF}} \\ 1, & \text{otherwise.} \end{cases} \tag{5.16}$$

Note that $z_{\mathrm{RF},t}$ updates the belief map only when it carries information on the target position (i.e., $z_{\mathrm{RF},t} \neq \varnothing$, at $t = MT_{\mathrm{RF}}$).

To define the *visual likelihood*, the observation model (5.11) is combined with the POD

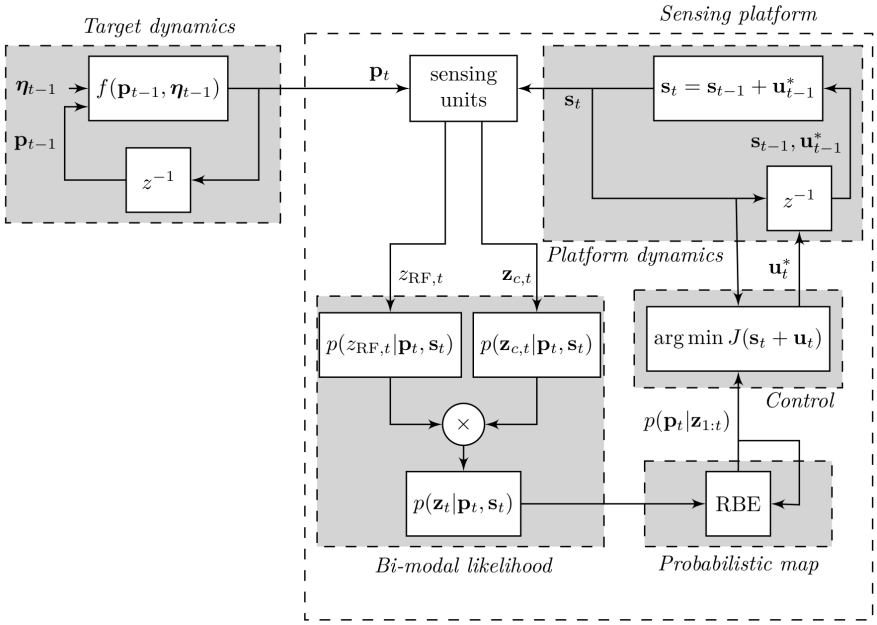

Figure 5.2. Scheme of the proposed Probabilistic Radio-Visual Active Sensing algorithm.

model (5.14)

$$p(\mathbf{z}_{c,t}|\mathbf{p}_t, \mathbf{s}_t) = p(\mathbf{z}_{c,t}|\mathbf{p}_t, \mathbf{s}_t, D_t)p(D_t|\mathbf{p}_t, \mathbf{s}_t) = \begin{cases} \mathcal{N}(\mathbf{z}|\mathbf{P}(\mathbf{s}_t)\tilde{\mathbf{p}}_t, \boldsymbol{\Sigma}_c)\Upsilon(\mathbf{p}_t, \mathbf{s}_t), & D_t = 1,\ \mathbf{p}_t \in \Phi(\mathbf{s}_t) \\ 0, & D_t = 1,\ \mathbf{p}_t \notin \Phi(\mathbf{s}_t) \\ 1 - \Upsilon(\mathbf{p}_t, \mathbf{s}_t), & D_t = 0,\ \mathbf{p}_t \in \Phi(\mathbf{s}_t) \\ 1, & D_t = 0,\ \mathbf{p}_t \notin \Phi(\mathbf{s}_t). \end{cases} \tag{5.17}$$

By aggregating radio and visual likelihoods, the following *bi-modal likelihood* is obtained

$$p(\mathbf{z}_t|\mathbf{p}_t, \mathbf{s}_t) = p(z_{\mathrm{RF}.t}|\mathbf{p}_t, \mathbf{s}_t)p(\mathbf{z}_{c,t}|\mathbf{p}_t, \mathbf{s}_t). \tag{5.18}$$

Then, (5.18) is applied to the update stage of the particle filter that implements the RBE scheme (see Appendix A.3).

5.3.2 Optimal controller

The platform control input is computed by solving the following optimization program

$$\begin{aligned} \mathcal{C}: \ \mathbf{u}^*_t = & \arg\min_{\mathbf{u} \in \mathcal{A}} J(\mathbf{u}|z_{\mathrm{RF},1:t}, \mathbf{z}_{c,1:t}, \mathbf{s}_t) \\ & \text{s.t. } d\left(\mathbf{c}_{t+1}, \mathbf{c}_t\right) \leq E_t \end{aligned} \tag{5.19}$$

where E_t is the residual energy at time t, computed as

$$E_t = E_{tot} - \sum_{k=0}^{t-1} \Delta E_k = E_{tot} - \sum_{k=0}^{t-1} d(\mathbf{c}_{k+1}, \mathbf{c}_k). \tag{5.20}$$

The cost function is chosen as

$$J(\mathbf{u}|z_{\mathrm{RF},1:t}, \mathbf{z}_{c,1:t}, \mathbf{s}_t) = \frac{1}{2} d\left(\mathbf{c}_{\Pi,t+1}, \hat{\mathbf{p}}_t\right)^2, \tag{5.21}$$

where $\mathbf{c}_{\Pi,t+1}$ is the projection of $\mathbf{c}_{t+1}$ onto Π, and $\hat{\mathbf{p}}_t$ is the MAP estimate[2] of the target position; formally

$$\hat{\mathbf{p}}_t = \mathbf{p}_t^{(i^*)};\ i^* = \underset{i\in[1,N_s]}{\arg\max}\, w_t^{(i)}. \tag{5.22}$$

Note that $J(\cdot)$ extracts information from the belief map, according to the probabilistic active sensing approach (Fig. 5.2). In fact, it is function of $z_{\mathrm{RF},1:t}$ and $\mathbf{z}_{c,1:t}$ through $\hat{\mathbf{p}}_t$ (see Appendix A.3). It is also function of $\mathbf{s}_t$ and $\mathbf{u}$ through $\mathbf{c}_{\Pi,t+1}$, which depends on the next camera pose (i.e., $\mathbf{s}_{t+1} = \mathbf{s}_t + \mathbf{u}_t$), according to the inverse perspective geometry (with known and fixed UAV altitude w.r.t. Π) Aghajan and Cavallaro (2009).

The convexity of $J(\cdot)$ w.r.t. $\mathbf{c}_{\Pi,t+1}$ allows to solve (5.19) with the following gradient-based control law

$$\begin{cases} \mathbf{s}_{\tau+1} = \mathbf{s}_\tau + \mathbf{u}_\tau,\ \tau \in [0, \tau_{max}] \\ \mathbf{u}_\tau = \begin{bmatrix} \mathbf{u}_{\mathbf{c},t} \\ \hline \mathbf{u}_{\psi,t} \end{bmatrix} = -\left[\begin{array}{c|c} \mathbf{G_c} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{G}_\psi \end{array}\right] \begin{bmatrix} \frac{\partial J(\mathbf{c}_{\Pi,\tau+1})}{\partial \mathbf{c}_\Pi} \\ \hline 0 \\ \hline \frac{\partial J(\mathbf{c}_{\Pi,\tau+1})}{\partial \psi} \end{bmatrix} \end{cases} \tag{5.23}$$

where τ_{max} accounts for the maximum number of iterations in order to accommodate the next incoming measurement at $t+1$. $\mathbf{G_c} \in \mathbb{R}^{3\times3}$ and $\mathbf{G}_\psi \in \mathbb{R}^{2\times2}$ are suitable control gain matrices. By choosing $\mathbf{G_c}$ entries small, energy is preserved, since (5.23) commands short UAV movements. Conversely, larger $\mathbf{G_c}$ and $\mathbf{G}_\psi$ lead to a more reactive system, capable of getting close to the setpoint $\hat{\mathbf{p}}_t$ within a maximum number of iterations τ_{max}. Consequently, if the localization procedure is accurate (i.e., $\hat{\mathbf{p}}_t \approx \mathbf{p}_t$), the condition $\mathbf{p}_t \in \Phi(\mathbf{s}_t)$ is likely to be satisfied and $d(\mathbf{c}_t, \mathbf{p}_t^+)$ is small, which is necessary to have high detection probabilities, according to (5.14). Finally, it is important to remark that $J(\cdot)$ is *purely-exploitative* and *not energy-aware*: in the controller design problem $\mathcal{C}$, energy appears only in the constraint and no energy preservation Liu et al. (2018), nor information-seeking (explorative) Radmard and Croft (2017) criteria are included.

5.4 Theoretical results

This Section formally motivates the use of an action space involving the entire camera pose, as in (5.5), and supports the choice of a combined radio-visual perception system. We observe that the particle weight distribution of a particle filter-based probabilistic

[2] As discussed in Appendix A.3, an alternative would be the MMSE estimate Hasanzade et al. (2018), which leads to smoother trajectories, but it is also proven to make $\mathcal{C}$ more prone to local minima Radmard and Croft (2017).

map[3] is an indicator of the target localizability: highly-weighted regions allow to focus the position estimate, while uniform weight patterns suggest ambiguity in the target localization. In this respect, it is possible to show that radio-only solutions need the sensing platform to move in order to solve localization ambiguity (Ths. 1-2 that follow), which can be conversely attained through a radio-visual approach also with a static platform (Th. 3).

Theorem 1. *Let the following hypotheses hold*

1. *the target moves according to an unbiased random walk, i.e.,*

$$p(\mathbf{p}_{t+1}|\mathbf{p}_t) = \mathcal{N}(\mathbf{p}|\mathbf{p}_t, \sigma^2\mathbf{I}_2); \tag{5.24}$$

2. *the platform is static, i.e.* $\mathbf{c}_t = \mathbf{c}$; $\forall t$;

3. *the RBE scheme updates through* (A.15b) *exploiting only the RF likelihood* (5.16).

Then,

$$\begin{aligned} &\mathbb{E}\left[\omega_t^{(i)}|z_{RF,1:t}\right] = \mathbb{E}\left[\omega_t^{(j)}|z_{RF,1:t}\right],\ t \geq 0 \\ &\forall i,j \in [1, N_s]\ s.t.\ d(\mathbf{p}_0^{(i)}, \mathbf{c}_\Pi) = d(\mathbf{p}_0^{(j)}, \mathbf{c}_\Pi). \end{aligned} \tag{5.25}$$

Proof. The dynamic model associated to the target unbiased random walk is

$$\mathbf{p}_{t+1} = \mathbf{p}_t + \boldsymbol{\eta}_t, \quad \boldsymbol{\eta}_t \sim \mathcal{N}(\boldsymbol{\eta}|\mathbf{0}_2, \sigma^2\mathbf{I}_2) \tag{5.26}$$

Equivalently,

$$\mathbf{p}_t = \mathbf{p}_0 + \sum_{k=0}^{t-1} \boldsymbol{\eta}_k,\ t > 0. \tag{5.27}$$

Given $\boldsymbol{\eta}_t$, $\boldsymbol{\eta}_s$ i.i.d. for any $s \neq t$, it follows

$$\bar{\boldsymbol{\eta}}_{t-1} := \sum_{k=0}^{t-1} \boldsymbol{\eta}_k \sim \mathcal{N}\left(\boldsymbol{\eta}|\mathbf{0}_2, t\sigma^2\mathbf{I}_2\right). \tag{5.28}$$

Then, the squared distance $d_t^{(i),2} := d(\mathbf{p}_t^{(i)}, \mathbf{c}_\Pi)^2$ is

$$\begin{aligned} d_t^{(i),2} &= \|\mathbf{p}_t^{(i)} - \mathbf{c}_\Pi\|_2^2 = \|\mathbf{p}_0^{(i)} + \bar{\boldsymbol{\eta}}_{t-1} - \mathbf{c}_\Pi\|_2^2 \\ &= \sum_{\ell=1}^{2} \left(p_0^{(i)}(\ell) - c_\Pi(\ell)\right)^2 + \sum_{\ell=1}^{2} \bar{\eta}_{t-1}(\ell)^2 \\ &+ 2\sum_{\ell=1}^{2} \left(p_0^{(i)}(\ell) - c_\Pi(\ell)\right) \bar{\eta}_{t-1}(\ell),\ t > 0. \end{aligned} \tag{5.29}$$

It holds,

$$2\left(p_0^{(i)}(\ell) - c_\Pi(\ell)\right)\bar{\eta}_{t-1}(\ell) \sim \mathcal{N}\left(\eta|0, 4t\sigma^2\left(p_0^{(i)}(\ell) - c_\Pi(\ell)\right)^2\right) \tag{5.30}$$

[3]The symbols and notation used in the Theorems (and in their proofs) of this Section refer to the particle filter introduction provided in Appendix A.3.

and, since the components of $\bar{\boldsymbol{\eta}}_{t-1}$ are i.i.d. with distribution $\mathcal{N}(0, t\sigma^2)$,

$$\sum_{\ell=1}^{2} \bar{\eta}_{t-1}(\ell)^2 \sim t\sigma^2 \chi^2(2). \tag{5.31}$$

Finally, recalling that $\sum_{\ell=1}^{2}(p_0^{(i)}(\ell) - c_\Pi(\ell))^2 = d_0^{(i)}$,

$$d_t^{(i),2} \sim \mathcal{N}(d^2|\mu_d, \sigma_d^2) + t\sigma^2\chi^2(2),\ t \geq 0 \tag{5.32}$$

with

$$\mu_d = d_0^{(i),2}, \quad \sigma_d^2 = 4t\sigma^2 d_0^{(i),2}. \tag{5.33}$$

From (5.32)-(5.33), $d_t^{(i),2}$ is identically distributed for all particles with the same initial distance from the platform. Using likelihood (5.16) and with $z_{\mathrm{RF},1:t}$ given, $\omega_t^{(i)}$ is a (non-linear) function of $d_t^{(i),2}$, through (5.9). Then, condition (5.25) follows. ■

Theorem 2. *Let the following hypotheses hold*

1. *the target moves according to an unbiased random walk, i.e.,*

$$p(\mathbf{p}_{t+1}|\mathbf{p}_t) = \mathcal{N}(\mathbf{p}|\mathbf{p}_t, \sigma^2\mathbf{I}_2); \tag{5.34}$$

2. *the platform moves at a fixed altitude according to a deterministic linear Markovian motion, i.e.,* $\mathbf{c}_{\Pi,t+1} = \mathbf{c}_{\Pi,t} + \mathbf{u}_{\mathbf{c}_\Pi,t}$;

3. *the RBE scheme updates through (A.15b) exploiting only the RF likelihood (5.16).*

Then,

$$\begin{aligned} &\exists (i,j) \in [1, N_s] \ s.t.\ d(\mathbf{p}_0^{(i)}, \mathbf{c}_{\Pi,0}) = d(\mathbf{p}_0^{(j)}, \mathbf{c}_{\Pi,0}) \ \textit{and} \\ &\mathbb{E}\left[\omega_t^{(i)}|z_{RF,1:t}, \mathbf{u}_{\mathbf{c}_\Pi,1:t}\right] \neq \mathbb{E}\left[\omega_t^{(j)}|z_{RF,1:t}, \mathbf{u}_{\mathbf{c}_\Pi,1:t}\right],\ t > 0. \end{aligned} \tag{5.35}$$

Proof. The platform planar dynamics can be equivalently written as

$$\mathbf{c}_{\Pi,t} = \mathbf{c}_{\Pi,0} + \bar{\mathbf{u}}_{t-1}, \quad \bar{\mathbf{u}}_{t-1} := \sum_{k=0}^{t-1} \mathbf{u}_{\mathbf{c}_\Pi,k},\ t > 0. \tag{5.36}$$

With similar computations involved in Th. 1, it is possible to show that the squared distance $d_t^{(i),2} := d(\mathbf{p}_t^{(i)}, \mathbf{c}_{\Pi,t})^2$ is distributed as

$$d_t^{(i),2} \sim \mathcal{N}(d^2|\mu_d, \sigma_d^2) + t\sigma^2\chi^2(2, \lambda),\ t > 0 \tag{5.37}$$

with

$$\begin{aligned} &\mu_d = d_0^{(i),2} - 2(\mathbf{p}_0^{(i)} - \mathbf{c}_{\Pi,0})^\top \bar{\mathbf{u}}_{t-1}, \quad \sigma_d^2 = 4t\sigma^2 d_0^{(i),2} \\ &\lambda = -\sum_{\ell=1}^{2} \bar{u}_{\mathbf{c}_\Pi,t-1}(\ell)^2 \end{aligned} \tag{5.38}$$

where $\chi^2(2, \lambda)$ is a non-central chi-squared distribution, with non-centrality parameter λ. From (5.37)-(5.38), the distribution of $d_t^{(i),2}$ depends on both $d_0^{(i),2}$ and $(\mathbf{p}_0^{(i)} - \mathbf{c}_{\Pi,0})^\top \bar{\mathbf{u}}_{t-1}$.

It is always possible to find $i \in [1, N_s]$ and $j \in [1, N_s]$ $(i \neq j)$, such that $d_0^{(i)} = d_0^{(j)}$ and

$$(\mathbf{p}_0^{(i)} - \mathbf{c}_{\Pi,0})^\top \bar{\mathbf{u}}_{t-1} \neq (\mathbf{p}_0^{(j)} - \mathbf{c}_{\Pi,0})^\top \bar{\mathbf{u}}_{t-1}.$$

Then, in general,

$$\mathbb{E}\left[\omega_t^{(i)} | z_{\mathrm{RF},1:t}, \mathbf{u}_{\mathbf{c}_\Pi,1:t}\right] \neq \mathbb{E}\left[\omega_t^{(j)} | z_{\mathrm{RF},1:t}, \mathbf{u}_{\mathbf{c}_\Pi,1:t}\right]. \tag{5.39}$$

■

Theorem 3. *Let the following hypotheses hold*

1. *the target moves according to an unbiased random walk, i.e.,*

$$p(\mathbf{p}_{t+1}|\mathbf{p}_t) = \mathcal{N}(\mathbf{p}|\mathbf{p}_t, \sigma^2 \mathbf{I}_2); \tag{5.40}$$

2. *the platform is static, i.e.* $\mathbf{c}_t = \mathbf{c}$; $\forall t$;

3. *the RBE scheme updates through* (A.15b) *exploiting the radio-visual likelihood* (5.18).

Then,

$$\begin{aligned} &\exists (i,j) \in [1, N_s] \ s.t. \ d(\mathbf{p}_0^{(i)}, \mathbf{c}_\Pi) = d(\mathbf{p}_0^{(j)}, \mathbf{c}_\Pi) \ and \\ &\mathbb{E}\left[\omega_t^{(i)} | z_{RF,1:t}, \mathbf{z}_{c,1:t}\right] \neq \mathbb{E}\left[\omega_t^{(j)} | z_{RF,1:t}, \mathbf{z}_{c,1:t}\right], \ t > 0 \end{aligned} \tag{5.41}$$

Proof. Suppose (w.l.o.g.) that $\mathbf{p}_0^{(i)}$ and $\mathbf{p}_0^{(j)}$ satisfy $d(\mathbf{p}_0^{(i)}, \mathbf{c}_\Pi) = d(\mathbf{p}_0^{(j)}, \mathbf{c}_\Pi)$; hence, $d_t^{(i),2}$ and $d_t^{(j),2}$ are equally distributed, from Th. 1. Suppose now that $\mathbf{p}_t^{(i)} \in \Phi(\mathbf{s}_t)$, but $\mathbf{p}_t^{(j)} \notin \Phi(\mathbf{s}_t)$. From (5.17)-(5.18), $\omega_t^{(i)}$ and $\omega_t^{(j)}$ are different functions of two equally distributed random variables (once the observations $z_{\mathrm{RF},1:t}$ and $\mathbf{z}_{c,1:t}$ are fixed).

Thus, in general,

$$\mathbb{E}\left[\omega_t^{(i)} | z_{\mathrm{RF},1:t}, \mathbf{z}_{c,1:t}\right] \neq \mathbb{E}\left[\omega_t^{(j)} | z_{\mathrm{RF},1:t}, \mathbf{z}_{c,1:t}\right]. \tag{5.42}$$

■

5.4.1 Discussion

Th. 1 statistically characterizes the *axis-symmetric ambiguity*, one of the main issues of RSSI-based localization systems. As (5.9) suggests, RSSI values, collected by isotropic receivers, bring information only on the target-receiver distance; consequently, the RF likelihood (5.16) is axis-symmetric w.r.t. the receiver. Thus, in a particle-based RBE scheme, particles at the same distance from the receiver are expected to equally weighted (on average). When the sensing platform is static, this phenomenon generates severe convergence issues, since the belief map is toroidal with a non-unique MAP estimate, as depicted in Fig. 5.3(a). For this reason, most literature solutions propose to use multiple receivers Cantón Paterna et al. (2017b) in order to reduce the axis-symmetric ambiguity and find a unique solution in localization problems.

Th. 2 guarantees that, when multi-receiver platforms are not feasible, the axis-symmetric effect can be mitigated by using a single moving receiver. Therefore, including the platform position in the control space, as in (5.5), is beneficial for the convergence of RSSI-based probabilistic active tracking strategies.
Th. 3 says that another viable solution is the adoption of a bi-modal radio-viual perception system, since the visual likelihood (5.17) reduces the ambiguity in the particles weighting process. This represents one of the main advantages of radio-visual approaches w.r.t. radio-only counterparts. To completely exploit the disambiguation effect of visual observations, camera orientation should be included in the control space, as in (5.5). In particular, with the control law defined in Sec. 5.3, the platform uses camera measurements to discard wrong MAP candidates. In this way, most of the weight will concentrate around the true target position over time, as Fig. 5.3(b) shows.

In conclusion, by combining radio-visual cues in the RBE scheme and by comprising the entire camera pose in $\mathcal{A}$, the estimation ambiguities can be reduced. This makes the localization procedure faster and more accurate (see Sec. 5.4.2). Accordingly, the overall target visual detection task becomes more robust and more efficient (see Sec. 5.5).

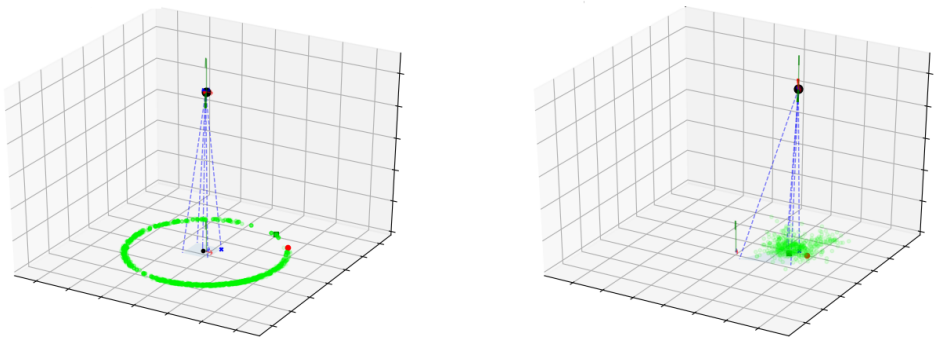

(a) (b)

Figure 5.3. Probabilistic map obtained with the uni-modal RF likelihood (a) and the bi-modal radio-visual likelihood (b).

5.4.2 Proof of concept - target localization

To support the theoretical results, the control law (5.19) is used to localize a static target (i.e., $\mathbf{p}_t = \mathbf{p}$; $\forall t$) in a Python-based synthetic environment (see Fig. 5.3)[4]. For consistency with the hypotheses of Th. 1-3, the process model is an unbiased random walk. The receiver is kept static (i.e., $\mathbf{G_c} = \mathbf{0}$), to stress the effects of the axis-symmetric ambiguity, as well as the improvements introduced by visual information. RSSI data are collected with sampling rate $T_{\text{RF}} = 2$ ($T_c = 1$) and noise level $\sigma_{\text{RF}} = 2$ $[dBm]$; the target is placed at distance $d(\mathbf{p}^+, \mathbf{c}) = 10$ $[m]$ from the platform. The proposed bi-modal radio-visual approach (RF+V) is compared with a visual-only (V) and a radio-only (RF) one. To capture the performance variability, the evaluation is assessed through a Monte Carlo (MC) simulation with $N_{tests} = 150$ tests of $T_W = 100$ iterations each. The

[4] https://github.com/luca-varotto/RF-Vision

performance index is the localization error $e_t = d(\mathbf{p}, \hat{\mathbf{p}}_t)$.

Fig. 5.4 depicts the average localization error (μ_{e_t}) and its 68% confidence interval over the MC tests (with standard deviation σ_{e_t}). Only RF+V converges to the groundtruth; namely, the estimate is unbiased and with zero variability over the MC tests

$$\begin{aligned} \mu_{e_{T_W}} &:= \mathbb{E}[e_{T_W}] = 0 \ [m] \\ \sigma_{e_{T_W}} &:= \sqrt{\mathbb{E}[(e_{T_W} - \mu_{e_{T_W}})^2]} = 0 \ [m]. \end{aligned} \tag{5.43}$$

The detrimental effect of the axis-symmetric ambiguity can be appreciated in the results of the RF approach: as proven in Th. 1, the MAP estimate is non-unique; hence, at any time, the estimate $\hat{\mathbf{p}}_t$ may be not representative of the true target position. This leads RF to converge, on average, to a biased estimate ($\mu_{e_{T_W}} \approx 1 \ [m]$), with remarkable inter-test variability ($\sigma_{e_{T_W}} \approx 3 \ [m]$).

The oscillations of e_t are even higher for the visual-only approach. The main issue here is related to the limited sensing domain of the camera: as (5.17) suggests, in case of non-detection events ($D_t = 0$), only particles inside the camera FoV are updated, which might be a small subset of the entire particles set. Consequently, many particles share the same weight for a long time, leading to a stronger ambiguity than the axis-symmetry (especially in the first iterations). Therefore, the localization procedure experiences a very slow and erratic converge behavior.

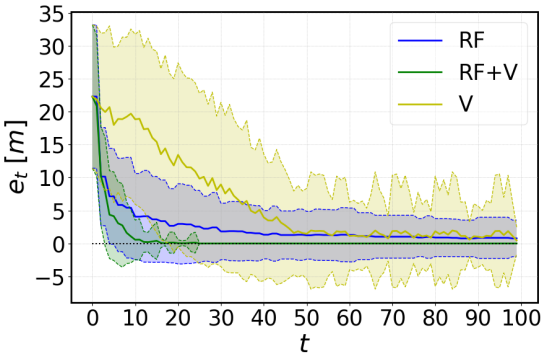

Figure 5.4. Localization error (average and 68% confidence interval over a MC simulation) with radio-only (RF), visual-only (V) and radio-visual (RF+V) likelihood.

5.5 Numerical results

In this Section the localization problem tackled in Sec. 5.4.2 is extended to the active tracking of a moving target, as stated in Sec. 5.2. At first, the main setup parameters are introduced (Sec. 5.5.1), as well as the performance metrics used to validate the proposed approach (Sec. 5.5.2). Then, Sec. 5.5.3 numerically motivates the choice of the entire camera pose as control space, in accordance with the theoretical results of Sec. 5.4. Sec. 5.5.4 instead shows that bi-modality allows to achieve higher robustness and time-efficiency with respect to several uni-modal baselines. Finally, in Sec. 5.5.5 the proposed algorithm is proven to be even more energy-efficient than a radio-only

counterpart, albeit the control law (5.19) does not account for any energy-preserving term.

5.5.1 Setup parameters

As in Sec. 5.4.2, RF+V, RF and V are compared within the same Python-based synthetic environment and over a MC simulation with $N_{tests} = 150$ tests of $T_W = 100$ iterations each. The belief map is approximated by $N_s = 500$ particles, initialized according to a uniform distribution over Π. The underlying target motion is simulated using a linear stochastic Markovian transition planar model, namely

$$\mathbf{p}_{t+1} = \mathbf{p}_t + \mathbf{q}_t; \quad \mathbf{q}_t \sim \mathcal{N}(\mathbf{q}|\boldsymbol{\mu}_\mathbf{q}, \boldsymbol{\Sigma}_\mathbf{q}), \ \boldsymbol{\mu}_\mathbf{q} \neq \mathbf{0}_2. \tag{5.44}$$

The initial condition $\mathbf{p}_0$ is randomly changed at each MC test, as well as the platform initial planar position $\mathbf{c}_{\Pi,0}$. The process model is an unbiased random walk, the receiver sampling rate is $T_{\text{RF}} = 10$ ($T_c = 1$), the total available energy is fixed to $E_{tot} = 300$, and the RSSI noise level is $\sigma_{\text{RF}} = 3.5$ $[dBm]$.

5.5.2 Performance assessment

The following performance indexes are computed from the MC simulation, recalling the definition of Empirical Cumulative Distribution Function (ECDF) for a generic variable q:

$$\begin{aligned} p(q \leq Q) &= \frac{1}{N_{tests}} \sum_{j=1}^{N_{tests}} \mathbb{1}_{q_j \leq Q} \\ \mathbb{1}_{q_j \leq Q} &= \begin{cases} 1, & \text{if } q_j \leq Q \\ 0, & \text{otherwise} \end{cases} \end{aligned} \tag{5.45}$$

where q_j is the value of the j-th MC test and Q is the considered threshold.

1. **Detection success rate**: *robustness index* that counts the number of successful target visual detections over the MC tests; formally,

$$\begin{aligned} \varpi &= \frac{1}{N_{tests}} \sum_{j=1}^{N_{tests}} \mathbb{1}_{D,j} \\ \mathbb{1}_{D,j} &= \begin{cases} 1, & \text{if } \exists t_D \in [1, T_W] \text{ s.t. } D_{t_D} = 1 \\ 0, & \text{otherwise} \end{cases} \end{aligned} \tag{5.46}$$

   where $\mathbb{1}_{D,j}$ is 1 only if the target is detected ($D_{t_D} = 1$) at any instant during the j-th test ($\exists t_D \in [1, T_W]$); t_D is the time of first detection, namely

$$t_D = \min\{t \in [1, T_W] : D_t = 1\}. \tag{5.47}$$

2. **Detection time ECDF**: *time-efficiency index* that accounts for the time required to visually detect the target over the MC tests. The ECDF is obtained applying

(5.45) with $q = t_D$ and $Q = t$.

3. **Energy consumption ECDF**: *energy-efficiency index* that accounts for the amount of available energy $E_{T_W} \in [0, E_{tot}]$ at the end of the task. The ECDF is obtained applying (5.45) with $q = E_{T_W}$ and $Q = E$.

5.5.3 Impact of the control space

The controller $\mathcal{C}$, described in Sec. 5.3, acts on a high-dimensional space comprising the full camera pose (namely, its position and orientation). This choice has been theoretically motivated in Ths. 2-3 of Sec. 5.4 and it is here supported by numerical results.

Specifically, the proposed control law (i.e., RF+V) is compared with two variants.

1. The first baseline does not allow pan-tilt movements (i.e., $\mathbf{G_\Psi} = \mathbf{0}$) and the camera is fixed and facing downwards (namely, $\alpha_t = \beta_t = 0$, $\forall t$). Without pan-tilt actuation, the target can be detected only when the UAV gets very close to it; more precisely, $\mathbf{p}_t \in \mathbf{\Phi}(\mathbf{s}_t)$ (necessary condition to have $D_t = 1$) is satisfied only when the platform is approximately above the target (i.e., $\mathbf{c}_{\Pi,t} \approx \mathbf{p}_t$)[5]. However, this condition is difficult to be met for every initial condition $\mathbf{p}_0$ and $\mathbf{c}_{\Pi,0}$, due to the localization errors and to the bias in the target motion. More specifically, localization errors imply that $\hat{\mathbf{p}}_t \approx \mathbf{p}_t$ does not hold; thus, $\mathbf{c}_{\Pi,t} \approx \mathbf{p}_t$ may not be met (hence, $\mathbf{p}_t \in \mathbf{\Phi}(\mathbf{s}_t)$ neither), since $\mathcal{C}$ drives the platform towards $\hat{\mathbf{p}}_t$. At the same time, the target motion bias ($\boldsymbol{\mu}_\mathbf{q}$) increases the difficulty of the task: consider that, at time t, the localization error is small enough to drive the platform close to the target, so that $\mathbf{p}_t \in \mathbf{\Phi}(\mathbf{s}_t)$; at time $t+1$ the target may move outside the camera FoV and its position must be accurately estimated again to keep $\mathbf{p}_{t+1} \in \mathbf{\Phi}(\mathbf{s}_{t+1})$. In conclusion, if the control space is restricted to the UAV position only, the success of the tracking task hinges on the accuracy and reactivity of the localization process. In this regard, the full potentialities of the proposed radio-visual approach are poorly exploited, since the disambiguation effects of camera observations (see Th.3) are applied only locally (i.e., in a neighbourhood of the UAV position). In this way, the ambiguities introduced by RSSI data, and trasnsmitted to the particles, are propagated in time, with consequent detrimental effects on the overall localization accuracy.

2. The second baseline considers a static platform (i.e., $\mathbf{G_c} = \mathbf{0}$). In this case, the UAV is not capable to get closer to the target; hence, many detection failures may occur, even when the position accuracy is sufficiently good to have the target inside the camera FoV. In fact, the detection capability of the camera depends on the resolution at which the target is observed and, according to (5.14)-(5.15), the POD is very low when the platform-target distance is large. Interestingly, choosing $\mathbf{G_c} = \mathbf{0}$ can also lead to the worst case scenario in which the tracking task is unfeasible; this happens when $d(\mathbf{p}_0^+, \mathbf{c}_0)$ is large enough to have almost zero POD

[5] Recall that another necessary condition is to have an UAV altitude small enough to guarantee non-zero POD when $\mathbf{p}_t \in \mathbf{\Phi}(\mathbf{s}_t)$, according to (5.15). Otherwise, the tracking problem is ill-defined.

Table 5.1. Detection success rate of the compared approaches.

	RF+V	RF+V: $\mathbf{G}_{\boldsymbol{\Psi}} = \mathbf{0}$	RF+V: $\mathbf{G}_{\mathbf{c}} = \mathbf{0}$
ϖ	77%	39%	29%
	RF	V	2RF
ϖ	55%	26%	64%

and the target moves away from the platform. Formally, from-(5.14)-(5.15)

$$\frac{d(\mathbf{p}_t^+, \mathbf{c}_t)}{f} \gg \epsilon, \ \forall t \geq 0. \tag{5.48}$$

The considerations above are supported by the numerical results reported in Tab. 5.1 (top row): the detection success rate of RF+V is 77%, against 39% of the case $\mathbf{G}_{\boldsymbol{\Psi}} = \mathbf{0}$ and 29% of the case $\mathbf{G}_{\mathbf{c}} = \mathbf{0}$. Moreover, Fig. 5.5 shows a remarkable difference in the detection time ECDF between the proposed solution and the other two versions with reduced control space.

In conclusion, considering the entire camera pose as control space induces more robustness and more time-efficiency in the target visual detection task.

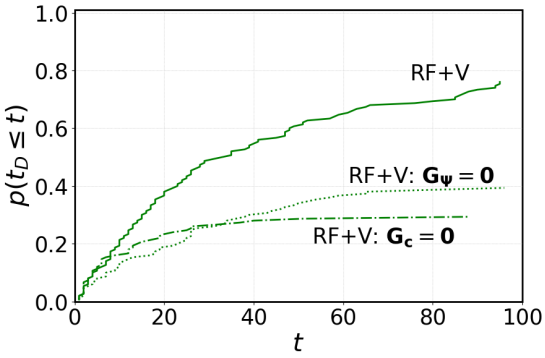

Figure 5.5. MC simulation results. Detection time ECDF with different control spaces: full camera pose (RF+V), no pan-tilt movements (RF+V: $\mathbf{G}_{\boldsymbol{\Psi}} = \mathbf{0}$), static platform (RF+V: $\mathbf{G}_{\mathbf{c}} = \mathbf{0}$).

5.5.4 Impact of the sensing modalities

Sec. 5.4.2 shows that aggregating radio-visual cues allow to achieve a more accurate localization than radio-only or visual-only approaches. Hence, control law $\mathcal{C}$, combined with the localization accuracy of a radio-visual approach, should regulate the platform position and the camera orientation in such a way that the UAV is often close to the target and this, in turn, is often inside the FoV. These two conditions should result into high detection probabilities, according to (5.14)-(5.15). As a consequence, radio-visual active tracking should be more robust and more time-efficient than radio and visual-only counterparts. Indeed, Tab. 5.1 (bottom row) highlights the fact that RF+V overcomes both RF and V in terms of detection success rate; moreover, Fig. 5.6 shows the higher time-efficiency of the bi-modal strategy with respect to the two uni-modal baselines.

One may argue that the superior performance of the bi-modal strategy is due to the larger amount of information involved. To rebut this assumption, RF+V is also compared with 2RF: instead of combining radio-visual data (heterogeneous radio-visual sensor fusion), the observations coming from two receivers are aggregated (homogeneous radio-only sensor fusion). Thus, the measurements of 2RF are

$$\mathbf{z}_{2\mathrm{RF},t} = \begin{bmatrix} z_{\mathrm{RF}^{(1)},t} & z_{\mathrm{RF}^{(2)},t} \end{bmatrix}^{\top} \tag{5.49}$$

where $z_{\mathrm{RF}^{(\ell)},t}$ identifies the RSSI sample collected by the ℓ-th receiver at time t, according to the model (5.10). Under the assumption of independence between the observations of the two receivers, the likelihood of 2RF becomes

$$p(\mathbf{z}_{2\mathrm{RF},t}|\mathbf{p}_t, \mathbf{s}_t) = \prod_{\ell=1}^{2} p(z_{\mathrm{RF}^{(\ell)},t}|\mathbf{p}_t, \mathbf{s}_t), \tag{5.50}$$

with $p(z_{\mathrm{RF}^{(\ell)},t}|\mathbf{p}_t, \mathbf{s}_t)$ as in (5.16). Tab. 5.1 and Fig. 5.6 show that 2RF gains some robustness and time-efficiency with respect to RF, but not enough to overcome RF+V. This means that the tracking performance is not only affected by the *quantity*, but also by the *quality*, of the aggregated information. In particular, exploiting complementary cues (e.g. radio and visual ones) is more advantageous than combining homogeneous data extracted from different sources; in fact, one can prove that the axis-symmetric ambiguity is not reduced when applying multiple receivers placed at the same location.

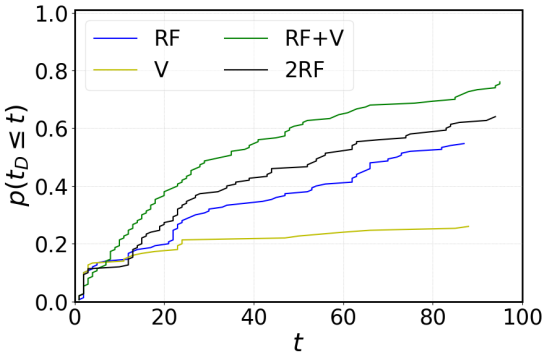

Figure 5.6. MC simulation results. Detection time ECDF with different type and amount of information employed for the tracking task: radio-visual (RF+V), visual-only (V), radio-only (RF), and radio-only with two receivers (2RF).

5.5.5 Energy-efficiency

As mentioned in Sec. 5.3, the proposed control law (5.19) does not include any explicit energy-aware term; hence, $\mathcal{C}$ is not expected to generate energy-preserving platform movements (unless $\mathbf{G_c}$ is kept small, but this strategy has already been proven to be ineffective). Therefore, when the total available energy runs out, that is

$$\exists t \in [1, T_W] \text{ s.t. } E_t = 0, \tag{5.51}$$

the platform becomes static and the detection capabilities dramatically decrease (see Sec. 5.5.3). In the absence of an energy-aware control technique, the only way to avoid energy run-out is by producing accurate target position estimates. In fact, if $\hat{\mathbf{p}}_t \approx \mathbf{p}_t$, the controller $\mathcal{C}$ drives the platform towards the target through a smooth and direct trajectory, which is more energy-efficient than irregular patterns, according to the energy model (5.8).

Given this premise, an considering the localization superiority of RF+V against RF (Sec. 5.3 and 5.5.4), the expectation is to have more energy preservation in RF+V than in RF. Indeed, the two energy consumption ECDFs showed in Fig. 5.7 satisfy

$$p(E_{T_W}^{\mathrm{RF+V}} \leq E) < p(E_{T_W}^{\mathrm{RF}} \leq E);\ \forall E \in [0, E_{tot}]. \tag{5.52}$$

This means that RF drains the initial available energy faster than RF+V. In particular, the out-of-energy condition (5.51) is higher in RF (76%) than in RF+V (68%).

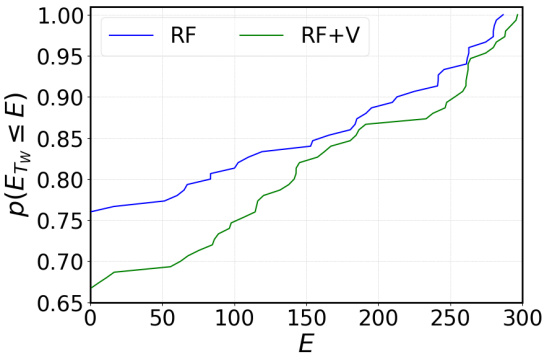

Figure 5.7. MC simulation results. Energy consumption ECDF between RF+V and RF.

5.6 Conclusion

This work discusses the problem of visually detecting a RF-emitting target under energy constraint. To this aim, a probabilistic active sensing approach, leveraging radio-visual information, is proposed. A theoretical discussion highlights the main limitations of radio-only localization, as well as the benefits introduced by the bi-modal strategy. These findings are then validated through numerical results, from which bi-modality is shown to be more robust and more time and energy-efficient than uni-modal camera and radio-only counterparts (in terms of detection success rate, time to first detection, and energy consumption). Bi-modality is proven to be even more effective than homogeneous sensor fusion strategies, since the combination of two receivers does not get comparable performance of the suggested radio-visual integration.

6

PAS for Energy-constrained Video Surveillance: an application to Camera Wake-Up

The application described in this Chapter considers a wireless smart camera, equipped with a low-energy radio receiver, and used to visually detect a moving radio-emitting target. To preserve the camera lifetime, probabilistic energy-aware controller is designed to switch on and off the camera. The decision-making module relies on the radio signal strength measured at the receiver side and used to estimate the target detectability, via Recursive Bayesian Estimation combined with self-supervised Gaussian Process Regression. Both numerical and experimental results validate the proposed approach in terms of target detection accuracy and energy consumption.

Contents

The contents of this chapter are partly available in: Varotto and Cenedese (2021b)
L. Varotto and A. Cenedese, "*Probabilistic RF-Assisted Camera Wake-Up through Self-Supervised Gaussian Process Regression*". IEEE 29th Mediterranean Conference on Control and Automation (MED), 2021 [accepted]

6.1 Introduction

Wireless sensors are battery-powered devices that do not require any wiring infrastructure to work Jelicic, Magno, Brunelli, Bilas, and Benini (2014). Fostered by the emerging Internet of Things (IoT) paradigm Zanella, Bui, Castellani, Vangelista, and Zorzi (2014) and the evolution in embedded systems, and thanks to their high scalability and to their fast and economical deployment, wireless sensors have become a propulsive technology in various applications, from home and industrial automation, to environmental monitoring Kandris, Nakas, Vomvas, and Koulouras (2020). However, the finite battery capacity often affects the sensors reliability, with consequent degradation of the sensing performance and with high maintenance costs for battery replacement. This becomes critical in applications like environment monitoring Lopes and Ruiz (2008) and video surveillance Jelicic et al. (2014), where energy-hungry sensors (e.g., cameras SanMiguel and Cavallaro (2016)) monitor events that occur sporadically and unpredictably Aghdasi, Nasseri, and Abbaspour (2013) (e.g., anomalous events Peri, Khorramshahi, Saketh Rambhatla, Shenoy, Rawat, Chen, and Chellappa (2020)). Thus motivated, the design of efficient energy management schemes has recently become a major theme in research Engmann, Katsriku, Abdulai, Adu-Manu, and Banaseka (2018). In this regard, modern computer vision systems have been designed to include attentional strategies that limit processing in an attempt to both mimic human visual behavior, as well as to economize on processing in the face of limited resources Bajcsy et al. (2018). More in general, several researchers within the wireless sensor networks community have proposed *multi-tier wake-up strategies* Jelicic et al. (2014) to preserve sensors lifetime and functionality: the first sensor layer is composed by low-power scalar nodes (e.g., Pyroelectric Infrared sensors, PIRs Jelicic et al. (2014)), which continuously monitor the environment; if any event of interest is detected, the power-consuming sensors (second layer) are activated. In this way, the energy-hungry sensors can be switched off most of the time and triggered only if needed.

In this Chapter, the focus is on wireless smart cameras which are nowadays ubiquitous in both industrial and civil contexts. Their on-board sensing, processing and communication capabilities Rinner and Wolf (2008) enable complex vision-based applications, even if at the cost of high-power consumption SanMiguel and Cavallaro (2016). For this reason, a number of energy preservation strategies for camera networks have been proposed. These are often based on multi-layer architectures Aghdasi et al. (2013), where cameras are coupled with scalar nodes, such as PIRs Magno et al. (2013) or Radio Frequency Identification (RFID) tags Alemdar, Durmus, and Ersoy (2010). PIR nodes are widely used, being low-cost, low-power and reliable presence detectors. RFIDs are limited by the possibility of detecting radio-emitting objects only. This, however, allows to reduce false alarms and wrong target identifications, thanks to the ID information

carried by radio-frequency (RF) signals. Moreover, some RFIDs are passive (i.e., they do not have an internal power source); hence, they do not impact on the system energy consumption. Despite this, passive RFIDs are limited by short sensing ranges. As an example, authors of Lopes and Ruiz (2008) propose a three-tier architecture: the first sensor layer is composed of passive infrared (PIR) scalar nodes, which monitor the environment continuously. When a heat-emitting event is detected by PIR nodes, the second-tier, made up of camera nodes with limited capabilities, is activated. At this step, a visual identification process occurs and, if the target is of interest, more powerful cameras are used to track the target. This kind of multi-tier solutions typically require extra hardware costs, but they significantly reduce the energy consumption, leading to maintenance costs savings, while providing high sensing performance Jelicic et al. (2014).
Contribution - The setup considered in this Chapter comprises a wireless smart camera, equipped with a low-energy radio receiver. The purpose is to visually detect a RF-emitting target and to minimize the overall energy consumption of the platform. To this aim, a multi-tier wake-up system activates the camera only when the target detectability is sufficiently high. Similarly to Chapt. 4, Gaussian Process Regression is used to learn, in a self-supervised fashion and exploiting the correlation between radio-visual inputs, the map between the Received Signal Strength Indicator (RSSI) at the receiver Zanella (2016) and the target probability of detection (POD). In this way, the same advantages discussed in Chapt. 4 are obtained, namely the human intervention is minimized and it is possible to avoid the tiring human-labeled calibration processes involved in traditional RF-based solutions Zanella (2016). On the other side, o account for the uncertainties on the target dynamics and on the perception process, the problem is formulated within a Bayesian probabilistic framework Smith (2013), similarly to Chapt. 5. This is the first attempt to design a RSSI-based probabilistic camera wake-up strategy. Furthermore, the suggested algorithm is suitable for mobile camera applications, as opposed to most existing works, relying on large static sensor networks Jelicic et al. (2014); Magno et al. (2013). Finally, numerical and experimental results validate the effectiveness of the proposed algorithm: it maintains high detection rates and introduces rema rkable energy savings compared to non-wake-up and random wake-up policies.

6.2 Problem statement

Fig. 6.1 shows the main elements of the problem scenario, namely the target and the sensing platform

6.2.1 Sensing platform

The sensing platform is a smart camera, endowed with a real-time target detector Redmon et al. (2016); Viola and Jones (2001) and a low-energy radio receiver. The camera works at a frame rate $1/T_c$ and can be either in *active* or in *sleep* mode: the former allows to detect the target, but with a high energy cost; the latter, instead, is a power saving mode and makes the camera inoperative SanMiguel and Cavallaro (2016). Therefore, the camera operational mode allows to define the platform state as a boolean variable,

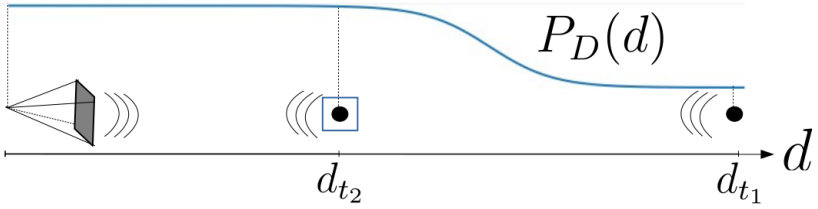

Figure 6.1. Problem scenario. A target (black marker) moves towards a camera, which has to visually recognize the target, supported by radio communication. At d_{t_1} the target is far from the camera and the detection module is ineffective, while at d_{t_2} the target is likely to be visually detected (blue box). In a resource-constrained framework, the camera should activate only when the detection probability $P_D(d)$ is sufficiently high.

namely

$$s_t \in \{0,1\}, \quad t = LT_c,\ L \in \mathbb{N} \tag{6.1}$$

where $s_t = 0$ and $s_t = 1$ denote the sleep and active mode, respectively. The platform state is regulated through a *switching-state control input* u_t; hence, s_t follows the deterministic Markovian transition model

$$s_{t+T_c} = s_t + u_t \tag{6.2}$$

$$u_t \in \mathcal{A}(s_t) = \begin{cases} \{0,1\}, & \text{if } s_t = 0 \\ \{0,-1\}, & \text{if } s_t = 1. \end{cases} \tag{6.3}$$

Similarly to the setup defined in Chapt. 4, and differently from that of Chapt. 5, the controller is here supposed to perform a state transition at multiples of $T > T_c$, where T is set equal to the receiver sampling interval T_{RF}. In this way, a new control input is computed as a new RSSI sample is collected; hence, between two consecutive RSSI samples the camera operational mode remains unchanged, namely

$$u_t = 0, \quad t \neq HT,\ H \in \mathbb{N}. \tag{6.4}$$

Finally, as in the previous Chapters, the following relation is supposed to hold regarding the sampling rate of camera and receiver

$$T_{\text{RF}} = \nu T_c,\ \nu \in \mathbb{N}. \tag{6.5}$$

Putting everything together, the same chain-relation of (4.4) holds:

$$T = T_{\text{RF}} = \nu T_c,\ \nu \in \mathbb{N}. \tag{6.6}$$

The camera energy consumption follows the model

$$\begin{aligned} E_c(s_{HT+T_c}, s_{HT}) =& T\left[(P_A + P_{OD})s_{HT+T_c} + P_S(1 - s_{HT+T_c})\right] \\ &+ E_{trans}|s_{HT+T_c} - s_{HT}|. \end{aligned} \tag{6.7}$$

If the new state is $s_{HT+T_c} = 1$, the power consumption is P_A with an additional cost P_{OD}, due to the object detector SanMiguel and Cavallaro (2016). Otherwise, a

reduced power $P_S < P_A$ is employed. If the camera operational mode changes (i.e., $s_{HT+T_c} \neq s_{HT}$), it requires a power of P_{trans} for a transition time T_{trans} (i.e., a total energy of $E_{trans} = P_{trans}T_{trans}$) SanMiguel and Cavallaro (2016). The power consumption of the receiver, P_{Rx}, is constant.

Overall, the accumulated platform energy consumption at a generic time instant $t \in]HT, (H+1)T]$ is

$$\begin{aligned} E_t = & \sum_{h=1}^{H-1} E_c(s_{hT+T_c}, s_{hT}) + E_{trans}|s_{HT+T_c} - s_{HT}| \\ & + (t - HT)\left[(P_A + P_{OD})s_{HT+T_c} + P_S(1 - s_{HT+T_c})\right] + tP_{\text{Rx}} \end{aligned} \tag{6.8}$$

The summation accounts for the energy consumption caused by the previous $H-1$ transitions; tP_{Rx} is the accumulated receiver energy consumption; the remaining part is the accumulated energy consumption between HT (current transition moment) and t (current instant of time).

6.2.2 Target

The radio-emitting target is supposed to move in front of the camera; hence, its state is the distance from the camera focal point, namely $d_t \in \mathbb{R}_+$. The *process model* is a (possibly) non-linear stochastic Markovian transition model

$$d_{t+T_c} = f(d_t, \eta_t). \tag{6.9}$$

As usual, the uncertainty on the underlying target movements are captured by the distribution of the stochastic process noise η_t.

6.2.3 Perception modeling

Radio perception Similarly to Chapt. 4 and 5, the target is supposed always inside the receiver sensing range and the RSSI value, $r_t \in \mathbb{R}$, is extracted from received data packets. The map between the target-platform distance, d_t, and the RSSI value,

$$r_t = g(d_t), \tag{6.10}$$

As in Chapt. 4, the function $g(\cdot)$ is unknown. By combining the process model (6.9) with (6.10), the (noise-free) RSSI value r_t can be equivalently considered as the target state. Therefore, the target dynamics becomes

$$r_{t+T_c} = f(g^{-1}(r_t), \eta_t) := f'(r_t, \eta'_t), \tag{6.11}$$

where the new process noise, η'_t, accounts for uncertainties in both target dynamics and in the underlying RSSI-distance relation.

Given the process model reformulation in (6.11), the target state is *hidden* in the

receiver measurements, according to the following *RSSI observation model*

$$z_{\mathrm{RF},t} = \begin{cases} h_{\mathrm{RF}}(\mathbf{p}_t, v_{\mathrm{RF},t}) = r_t + v_{\mathrm{RF},t}, & t = MT_{\mathrm{RF}},\ M \in \mathbb{N} \\ \varnothing, & \text{otherwise.} \end{cases} \tag{6.12}$$

The lack of samples (when $t \neq MT_{\mathrm{RF}}$) is modeled with empty observations $\varnothing$; this will turn useful for the formulation of a probabilistic controller as in Sec. 6.3.1. As can be noticed, the main difference w.r.t. the previous Chapters is that $h_{\mathrm{RF}}(\cdot)$ does not depend on the platform state, since the receiver is static and always working (the wake-up control is applied only to the camera).

Visual perception Differently from the assumptions made in the previous Chapters, the target is always inside the camera FoV. The target visual detection process is modeled as a Bernoulli random variable $D_t \in \{0, 1\}$, with success probability

$$p(D_t = 1 | d_t, s_t) = \begin{cases} P_D(d_t), & \text{if } s_t = 1 \\ 0, & \text{if } s_t = 0. \end{cases} \tag{6.13}$$

As (6.13) highlights, the detection success probability is a function of both the camera operational mode (i.e., the state s_t) and the target-platform distance, through $P_D(d_t)$. This is the POD when the camera is active and accounts for the visual depth effects (see Fig. 6.1). In the following, the dependence on d_t is equivalently substituted by r_t, on the basis of (6.10); formally

$$P_D(d_t) = P_D\left(g^{-1}(r_t)\right) := p_D\left(r_t\right). \tag{6.14}$$

6.2.4 Problem statement

With the formalism introduced in the previous Sections, the *RF-assisted camera wake-up* problem can be formulated as the control of the platform state s_t (through u_t), given camera and receiver observations $\mathbf{z}_{c,T_c:t}$ and $z_{\mathrm{RF},T_c:t}$ (i.e., all platform measurements). The purpose is the realization of the event $D_t = 1$, with a a minimal amount of energy E_t spent.

6.3 Methodology

To solve the RF-assisted camera wake-up problem, defined in Sec. 6.2, a two-step probabilistic (energy-aware) bi-modal active sensing approach is proposed.

1. First, given a set of offline RSSI and camera observations, GPR is applied to learn, in a self-supervised fashion, the POD as function of the RSSI value, that is $p_D(r)$. The method is the same as that described in Sec. 4.3.1; hence, the reader is addressed to Chapt. 4 for a detailed description of the first step of the proposed pipeline.

2. As shown in Fig. 6.2, the GPR result is incorporated into a RBE framework to

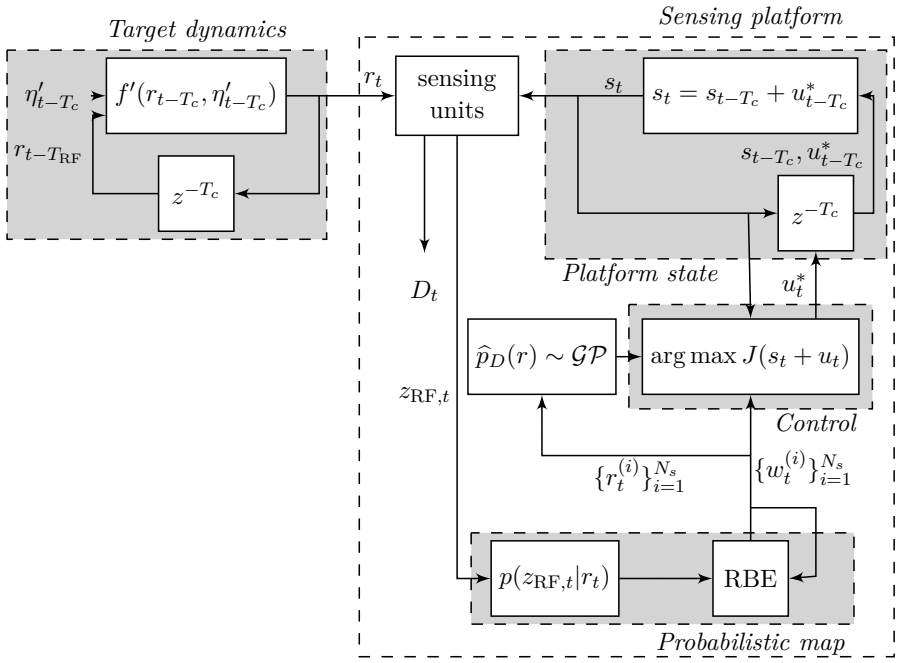

Figure 6.2. Scheme of the RF-assisted camera wake-up algorithm.

predict the target detectability and, therefore, to switch the camera operational mode. More specifically, radio measurements are used to keep the target belief map updated through an RBE scheme. The map feeds the pre-trained Gaussian Process module, which estimates the expected target POD over a time horizon. This estimate and the updated RBE map are injected into a probabilistic energy-aware controller, which is responsible for generating the platform control input.

6.3.1 Probabilistic RF-assisted camera wake-up

Probabilistic map The RBE scheme involved in Fig. 6.2 is implemented through a particle filter (see Appendix A.3). Given the target state r_t, the process model (6.11), the observations $z_{\mathrm{RF},T_c:t}$ and the observation model (6.12), the RBE map estimates the posterior distribution $p(r_t|z_{\mathrm{RF},T_c:t})$ via

$$p(r_t|z_{\mathrm{RF},T_c:t}) \approx \sum_{i=1}^{N_s} w_t^{(i)} \delta\left(r_t - r_t^{(i)}\right). \tag{6.15}$$

The particles and their weights are computed as

$$r_t^{(i)} = f'\left(r_{t-T_c}^{(i)}, \eta'_{t-T_c}\right) \tag{6.16a}$$

$$w_t^{(i)} \propto w_{t-T_c}^{(i)} p\left(z_{\mathrm{RF},t}|r_t^{(i)}\right). \tag{6.16b}$$

In our framework, the radio observation model (6.12) leads to the following *RF*

likelihood

$$p(z_{\mathrm{RF},t}|r_t) = \begin{cases} \mathcal{N}\left(z|r_t, \sigma^2_{\mathrm{RF}}\right), & t = MT_{\mathrm{RF}} \\ 1, & \text{otherwise.} \end{cases} \tag{6.17}$$

The belief map is updated only when $z_{\mathrm{RF},t}$ carries information on the target state (i.e., $z_{\mathrm{RF},t} \neq \varnothing$, at $t = MT_{\mathrm{RF}}$).

Probabilistic Energy-aware Controller The platform control input is computed by solving the following optimization program

$$\mathcal{C} : u_t^* = \begin{cases} \arg\max_{u \in \mathcal{A}(s_t)} J(u|z_{\mathrm{RF},T_c:t}, s_t), & t = HT_{\mathrm{RF}} \\ 0, & \text{otherwise,} \end{cases} \tag{6.18}$$

where condition (6.4), with $T = T_{\mathrm{RF}}$, is taken into account. The cost function is a combination of a *detection term* $J_D(\cdot)$ and an *energy-aware term* $J_E(\cdot)$, both related to u_t through (6.2):

$$J(u|z_{\mathrm{RF},T_c:t}, s_t) = J_D(u|z_{\mathrm{RF},T_c:t}, s_t) + \gamma J_E(u|s_t) \tag{6.19}$$

The detection term is

$$J_D(u|z_{\mathrm{RF},T_c:t}, s_t) = \mathbb{E}\left[p\left(\sum_{\ell=1}^{\nu} D_{t+T_{\mathrm{RF}},\ell} \geq 1 \middle| u \right) \middle| z_{\mathrm{RF},T_c:t}, s_t \right], \tag{6.20}$$

where $\{D_{t+T_{\mathrm{RF}},\ell}\}_{\ell=1}^{\nu}$ are the ν detection outcomes before the next RSSI sample, $z_{\mathrm{RF},t+T_{\mathrm{RF}}}$, is collected; hence, $J_D(\cdot)$ represents the expected probability of detecting the target at least once before a new RSSI measurement comes. Clearly, it is a function of the next platform state (i.e., s_{t+T_c}), from which the dependence on u_t and s_t follows, according to (6.2). Consequently, the detection cost function can be re-written as $J_D(s_{t+T_c}|z_{\mathrm{RF},T_c:t})$. From (6.13), $J_D(s_{t+T_c}|\cdot) = 0$ if $s_{t+T_c} = 0$; otherwise, $\{D_{t+T_{\mathrm{RF}},\ell}\}_{\ell=1}^{\nu}$ is a sequence of i.i.d. random variables with distribution $\mathcal{B}\left(p_D(r_{t+T_{\mathrm{RF}}})\right)$ (see Chapt. 4); thus

$$\sum_{\ell=1}^{\nu} D_{t+T_{\mathrm{RF}},\ell} \sim \mathcal{B}in\left(\nu, p_D(r_{t+T_{\mathrm{RF}}})\right) \tag{6.21}$$

and, therefore,

$$p\left(\sum_{\ell=1}^{\nu} D_{t+T_{\mathrm{RF}},\ell} \geq 1 \middle| u \right) = 1 - \left[1 - p_D(r_{t+\mathrm{RF}})\right]^{\nu}. \tag{6.22}$$

Finally, we get

$$\begin{aligned} J_D(s_{t+T_c} = 1|z_{\mathrm{RF},T_c:t}) &= 1 - \int \left[1 - p_D(r_{t+T_{\mathrm{RF}}}|z_{\mathrm{RF},T_c:t})\right]^{\nu} p(r_{t+T_{\mathrm{RF}}}|z_{\mathrm{RF},T_c:t}) dr_{t+T_{\mathrm{RF}}} \\ &\approx 1 - \sum_{i=1}^{N_s} w_t^{(i)} \left[1 - \widehat{p}_{D,*}(r_{t+T_{\mathrm{RF}}}^{(i)})\right]^{\nu}, \end{aligned} \tag{6.23}$$

where the RBE prediction step is applied to compute the particles $\{r_{t+T_{\mathrm{RF}}}^{(i)}\}_{i=1}^{N_s}$ from those at time t; at the same time, the weights $\{\omega_t^{(i)}\}_{i=1}^{N_s}$ encrypt the information carried by measurements $z_{\mathrm{RF},T_c:t}$ through the likelihood function (6.17). Finally, the posterior

POD $p(r_{t+T_{\mathrm{RF}}}|z_{\mathrm{RF},T_c:t})$ is estimated through the pre-trained GP-based POD function, fed with the RBE particles, namely $p(r_{t+T_{\mathrm{RF}}}|z_{\mathrm{RF},T_c:t}) \approx \widehat{p}_{D,*}(r^{(i)}_{t+T_{\mathrm{RF}}})$

The energy-aware term is

$$J_E(u|s_t) = \frac{E_c(s_{t+T_c}, s_t)}{E_{max}}, \tag{6.24}$$

which accounts for the energy consumption induced by the new state s_{t+T_c}, according to (6.7). This quantity is then normalized by the maximum energy consumption

$$E_{max} = (P_A + P_{OD})T_{\mathrm{RF}} + E_{trans}. \tag{6.25}$$

The hyperparameter $\gamma \geq 0$ is used to tune the importance of the energy preservation contribution: the higher γ, the higher must be the expected POD to justify larger energy expenditures for the camera activation (with $\gamma = 0$, the camera is always active).

6.4 Numerical and experimental results

The proposed approach is firstly evaluated with a Python-based synthetic environment[1] and then through a real-world experiment. Sec. 6.4.1 describes the main (numerical and experimental) setup parameters, as well as the synthetic data generation process. Sec. 6.4.2 defines the metrics used for performance assessment and the baselines considered for comparison. The numerical simulation results are discussed in Sec. 6.4.3, while the experimental validation is reported in Sec. 6.4.4.

6.4.1 Setup parameters

To carry out realistic simulations, most parameters reflect the experimental setup also in the synthetic experiments (see Tab. 6.1, where blue font highlights the difference between simulative and experimental setups). The platform and the target communicate via Bluetooth protocol, using Nordic nRF52832 SoCs[2]. RSSI observations are collected with a sampling time of $T_{\mathrm{RF}} = 0.1\,\mathrm{s}$. The camera node is the integrated webcam of a Dell Inspiron (Intel Core i7-4500U, 3.00 GHz processor), which collects images at resolution 640×480 pixels and with a frame rate of 20 Hz; hence, its transition energy, as well as the energy consumption in active and sleep modes, can be computed accordingly SanMiguel and Cavallaro (2016). For the target detection, the camera is endowed with a real-time face detector Viola and Jones (2001), whose energy consumption is studied in LiKamWa, Priyantha, Philipose, Zhong, and Bahl (2013). To simulate detection events in numerical tests, the POD in (4.20) is considered[3]. The value of ν with the experimental frame rates is 2; nonetheless, a value of $\nu = 10$ is used in the numerical tests, as a reasonable trade-off with the ideal condition stated in (4.18). The target is supposed to move according

[1] https://github.com/luca-varotto/RSSI-camera-wakeUp

[2] https://www.nordicsemi.com/-/media/Software-and-other-downloads/Product-Briefs/nRF52832-product-brief.pdf.

[3] $P_D(d) = \left[\left(1 + e^{4(d-4.5)}\right)\left(1 + e^{-(d-2.5)}\right)\right]^{-1}$.

Table 6.1. Setup parameters for simulations and experiment.

Parameter	Simulation	Experiment
T_{RF}	0.1 s	0.1 s
ν	10	2
P_{Rx}	20 mW	20 mW
P_A	100 mW	100 mW
P_S	1 mW	1 mW
P_{trans}	50 mW	50 mW
T_{trans}	3 ms	3 ms
P_{OD}	100 mW	100 mW
n_{train}	900	1073
n_{test}	500	936
N_s	100	100
N_{tests}	50	1
γ	1	1
$f(d_t, \eta_t)$	$d_t + \eta_t$ $\eta_t \sim \mathcal{N}(0, 0.04)$	unknown
σ_{RF}	3	unknown
$g(d_t)$	$\kappa - 10n \log_{10}(d_t/\delta)$ $\kappa = -30$ dBm, $n = 2$, $\delta = 1$ m	unknown

to the stochastic linear model reported in Tab. 6.1; the initial condition d_0 is chosen according to a uniform random variable in $[0.5, 6]$m. The log-distance path loss model (PLM) Zanella (2016) is used to generate synthetic RSSI data and to design the RF likelihood, with noise level $\sigma_{\mathrm{RF}} = 3$. By combining the target dynamics with the PLM, the RSSI process model (6.11) is obtained.

6.4.2 Performance assessment

To capture the performance variability, numerical evaluation is performed through MC simulations with $N_{tests} = 50$ tests of duration $T_W = n_{test} T_{\mathrm{RF}}$ each, where n_{test} is the number of RSSI collected.

Performance metrics The following performance indexes are computed from the MC simulation, recalling the definition of Empirical Cumulative Distribution Function (ECDF) defined in Sec. 5.5.2 of Chapt. 5.

- *Accuracy ECDF*: the main objective of any wake-up system is to save energy without reducing the effectiveness of the perception system. A bias in the wake-up decision process may induce imbalances between sensing performance and energy efficiency. More specifically, the camera wake-up problem can be seen as a binary decision task (camera is either active or not) based on the binary classification about the target detectability (target is either detectable or not). Therefore, the higher the classification accuracy, the higher the capability to activate the camera only when needed: leaving the camera active when the target is not detectable leads to (unjustified) energy waste; on the contrary, a bias towards the sleep mode increases the risk of losing the target. Moreover, if the wake-up system is unstable

(i.e., active and sleep states alternate faster than target dynamics) the transition costs may reduce the energy efficiency. The system accuracy is computed as

$$acc = \frac{1}{T_W} \sum_{t=T_c}^{T_W} \left[\underbrace{D_t s_t}_{TP} + \underbrace{(1 - D_t)(1 - s_t)}_{TN} \right], \tag{6.26}$$

which is the empirical probability of having either a true positive (TP) or a true negative (TN). At time t, a TP is realized if the camera is active and the target is detected. Conversely, a TN consists in switching-off the camera and having a non-detection event if the camera would have been active. The accuracy ECDF is then obtained applying (5.45) with $q = acc$ from (6.26) and $Q = a$.

- *Total Energy Consumption ECDF*: clearly, a high accuracy is necessary, but not sufficient to justify the additional power costs introduced by the ancillary equipment of multi-tier architectures (the radio receiver, in our case). More precisely, the use of a multi-layer wake-up system is well-founded only if the accuracy-induced energy savings overcome the costs due to the ancillary equipment. Since accuracy does not provide any information about the power usage, it is possible to evaluate a wake-up strategy also on the total energy consumption E_{T_W}, computed from (6.8), and recalling (5.45) to obtain the corresponding ECDF (with $q = E_{T_W}$ and $Q = E$, expressed in energy units).

- *Final Confusion-Energy ECDF*: accuracy and energy efficiency are sufficient to characterize the performance of a wake-up strategy; despite this, a performance index that aggregates these two metrics is needed, in order to ease the comparative analysis between different techniques. As a matter of fact, a very accurate solution may employ energy-hungry sensors that increase the energy usage; conversely, an energy-preserving algorithm sometimes do not yield high accuracy levels. In conclusion, it is useful to define the *confusion-energy* at $t = T_W$ (i.e., final CE)

$$CE_{T_W} = \frac{1}{E_{T_W}} \sum_{t=2T_c}^{T_W} (E_t - E_{t-T_c})\,(s_t + D_t - 2 s_t D_t)\,. \tag{6.27}$$

The CE is the ratio of energy spent in case of any mis-classification (i.e., $s_t \neq D_t$), over the total amount of energy usage. Thus, it is a measure of power management efficiency. The final CE ECDF is obtained similarly to the previous case, resorting to the general formulation in (5.45) with $q = CE_{T_W}$ and threshold value $Q = E_{norm}$, where energy normalization is applied.

Baselines The following baselines are considered for comparison in the simulative scenario:

- *always*: the camera is always left active (i.e., no camera wake-up strategy is applied) and the receiver is not used;

- *rnd*: the camera is activated randomly, that is $s_t \sim \mathcal{B}(0.5)$, and the receiver is not used;

- *gt*: the camera is activated according to the probabilistic control law $\mathcal{C}$, but assuming perfect knowledge on the underlying POD function $p_D(r)$, on the underlying path loss model, and on the platform-target distance at each time instant. Hence, $p_D(r)$ can be perfectly evaluated and, in turn, the computation of the cost function $J(u|z_{\mathrm{RF},T_c:t}, s_t)$ does not involve any approximation or estimation step. For this reason, *gt* works as comparison groundtruth and represents the performance upper bound.

The proposed approach differs from *gt* in that the underlying POD function $p_D(r)$ is estimated with the self-supervised GPR described in Sec. 4.3.1. Moreover, neither the platform-target distance nor the PLM are known; hence, the target state is predicted through the RBE scheme described in Sec. 6.3.1. For the aforementioned reasons, the proposed solution is referred as S^2GPR.

6.4.3 Numerical results discussion

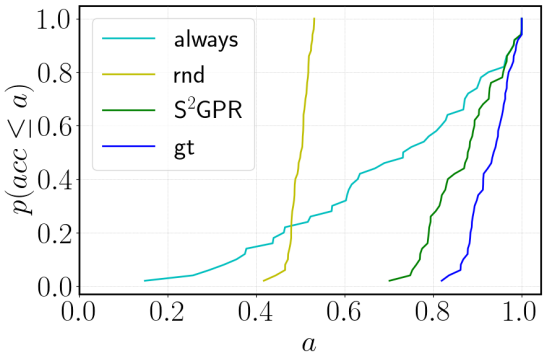

Figure 6.3. Numerical results: accuracy ECDF.

Accuracy Fig. 6.3 shows the accuracy ECDF of the proposed S^2GPR and the three baselines under comparison (*always*, *rnd*, *gt*). The accuracy of *always* experiences the highest variability over the MC simulations, being strongly related to the target behavior at each MC test. In fact, the accuracy depends only on the ratio between true and false positives (since $s_t = 1, \forall t$), which, in turn, depends on the time spent by the target at detectable distances. This result suggests the importance of adopting a camera activation scheme to increase the robustness w.r.t. all possible target behaviors. On the other hand, the accuracy of *rnd* is independent on the target dynamics and it is concentrated around 0.5; this is due to the fact that, with $s_t \sim \mathcal{B}(0.5)$, the probability of a TP or a TN is 0.5, irrespectively to the target detectability. This suggests that, to have better accuracy levels, the camera activation scheme should exploit some kind of target detectability awareness. As Fig. 6.3 shows, the exploitation of a RF-assisted strategy to account for both target movements and its probability of detection, allows to attain remarkable improvements; in particular, *gt* and S^2GPR are better than the non-RF-assisted baselines for every MC test. More specifically, the accuracy of *gt* is never smaller than 0.8, due to the perfect knowledge employed in the target detectability evaluation. Despite this, the probabilistic nature of $\mathcal{C}$ makes it difficult to reach $acc = 1.0$ for every MC test, even in this ideal scenario. The proposed algorithm performs slightly worse than *gt*, since GPR

and RBE introduce some estimation errors that affect the computation of the expected target detection probability. In any case, the accuracy performance is satisfactory, since it is never lower than 0.7 and in the range $[0.7 - 0.8]$ with probability 0.2.

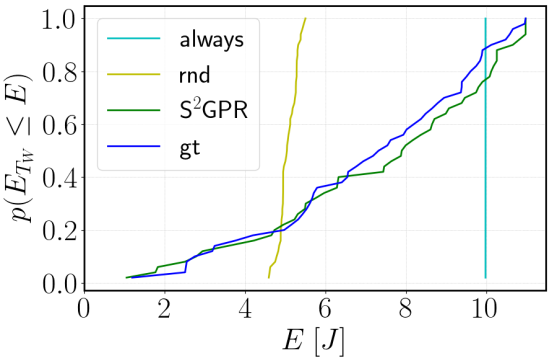

Figure 6.4. Numerical results: total energy consumption ECDF.

Energy consumption Fig. 6.4 depicts the ECDF of the total amount of energy usage, E_{T_W}. The consumption related to the *always* mode is obviously constant and equal to 10 J, since the camera never changes its operational mode. The RF-assisted strategies (*gt* and S^2GPR) have a very similar power usage, with minor preservation improvements coming from the perfect knowledge involved in *gt*. Interestingly, S^2GPR requires less power than *always* for the 80% of the MC simulations. The remaining 20% may be due to target movements lingering around a borderline distance for the detectability (e.g., $P_D(d) \approx 0.5$): in this case, the trade-off between J_D and J_E in (6.18) produces unstable and rapid state changes; therefore, the energy savings induced by camera deactivations are negligible and the frequent transition costs become relevant for the overall energy computation. The least energy-demanding algorithm is *rnd*, with an average energy consumption of 5.1 J. Note, however, that *rnd* is also the least accurate for the 80% of the MC tests, according to Fig. 6.3; for this reason, it is fundamental to to draw the main conclusions by considering a performance index that combines and balances together accuracy and energy consumption (i.e., the confusion energy metric).

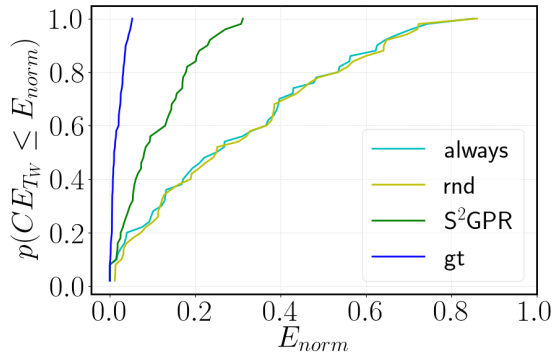

Figure 6.5. Numerical results: final confusion-energy ECDF.

Energy efficiency Fig. 6.5 shows the ECDF of the final CE, CE_{T_W}. The accuracy gap between *gt* and S^2GPR is reflected also on this quantity. In any case, the RF-assisted solutions are always more efficient than the others, despite the extra energy spent for

the receiver. In particular, the low accuracy of *rnd* leads to poor performance in terms of energy efficiency, despite the highest energy preservation levels depicted in Fig. 6.4. These findings prove the impact of accuracy on the energy usage efficiency. Overall, the results shown in Fig. 6.5 justify the adoption of multi-tier wake-up strategies.

6.4.4 Experimental results discussion

In this paragraph, the results of an experimental test are described, where the reader may refer to Sec. 6.4.1 for the setup and to Tab. 6.1 for the main parameters. The groundtruth baseline *gt* is not considered, since no perfect knowledge is available in real scenarios. Moreover, since the RSSI generative model and the noise level are unknown, it is possible to use an educated-guess value ($\widehat{\sigma}_{\mathrm{RF}} = 3$) for the RF-likelihood. Regarding the RSSI process model (6.11), a linear stochastic Markovian dynamics with variance 0.01 is adopted.

Fig. 6.6 depicts the POD function learnt via GPR on the real data. With respect to the synthetic version (Fig. 4.6), output labels are noisier, due to the smaller value of ν and according to (4.19). Moreover, depth effects impact on the target detectability also at smaller distances (larger RSSI values), as expected indeed in real-world computer vision systems Radmard and Croft (2017).

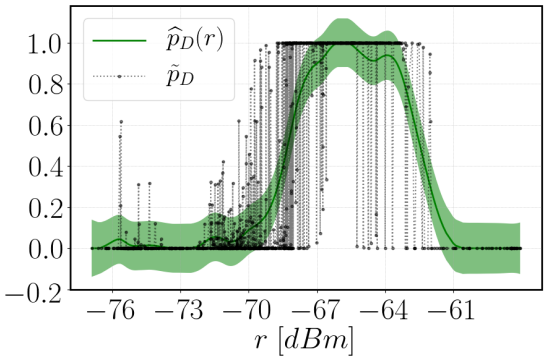

Figure 6.6. Experimental results: GP model $\widehat{p}_D(r)$ (green), learnt on real data $\{z_{\mathrm{RF},j}, \widetilde{p}_{D,j}\}_{j=1}^{n_{train}}$ (black markers).

To exemplify, Figs. 6.7(a) and 6.7(b) show sample images for a TP and a TN case: here, camera active and sleep modes are indicated by a green and red marker, respectively. In summary, the main quantitative results are given in Tab. 6.2: notably, S^2GPR exhibits the highest accuracy and the smallest confusion-energy, and achieves 68% and 37% energy saving w.r.t. *always* and *rnd*, respectively. As a final remark, it is important to note that one major benefit of the proposed approach is that it performs well also with small datasets. A total of $n_{train} = 1073$ training data were collected, at a sampling rate of $T_{\mathrm{RF}} = 0.1\,\mathrm{s}$; hence, the data collection campaign lasted only 107.3 s. A further advantage is the self-supervised procedure, which allows to automatize the dataset acquisition process.

(a) (b)

Figure 6.7. Experimental results. a) True Positive example: the target is detected and the camera is active (green marker on left top). b) True Negative example: the target is not detected and the camera is left silent (red marker on left top).

Table 6.2. Experiment results: accuracy, final CE and total energy consumption (best values in blue font).

		Method		
		always	*rnd*	S^2GPR
Metrics	*acc*	0.31	0.48	0.81
	CE_{T_W}	0.68	0.69	0.21
	E_{T_W}	23.47 J	12.17 J	7.61 J

6.5 Conclusion

This work proposes an energy-preservation scheme for a wireless camera, used to detect a radio-emitting target. A probabilistic energy-aware controller activates the camera based on the predicted target detectability. Predictions are made through Gaussian Process Regression, combined with Recursive Bayesian Estimation. The GP model is learnt in a self-supervised fashion, exploiting the correlation between radio-visual measurements and avoiding manual annotation. Experiments on a real platform show a 68% energy reduction w.r.t. a camera than continuously monitors the environment.

7

PAS for IoT: an application to parking estimation in connected vehicles

In the context of assisted driving and connected vehicle technologies, this Chapter proposes an online and adaptive scheme for parking availability mapping. Specifically, an information-seeking active sensing approach is adopted to select the incoming data, thus preserving the on-board storage and processing resources; then, the parking availability is estimated through Gaussian Process Regression. The proposed algorithm is compared with several baselines, which attain inferior performance in terms of mapping convergence speed and adaptivity capabilities; moreover, the proposed approach comes at the cost of a very small computational demand.

Contents

The contents of this chapter are partly available in: Varotto and Cenedese (2021a)
L. Varotto and A. Cenedese, "*Online and Adaptive Parking Availability Mapping: An Uncertainty-Aware Active Sensing Approach for Connected Vehicles*". IEEE 32nd Intelligent Vehicles (IV2021) - 2nd workshop on Online Map Validation and Road Model Creation (MAVROC), 2021 [accepted]

7.1 Introduction

IoT technology for connected and automated vehicles Research on connected vehicles represents a continuously evolving technological domain, fostered by the emerging Internet of Things (IoT) paradigm and the recent advances in intelligent transportation systems. Nowadays, vehicles are platforms capable of generating, receiving and automatically act based on large amount of data. In the context of assisted driving, connected vehicles technology provides real-time information about the surrounding traffic conditions Čičić, Barreau, and Johansson (2020). Such information is expected to improve drivers' quality of life, for example, through Advanced Driver Assistance Systems (ADASs), or by adopting decision making strategies according to the current parking availability status. Indeed, the research on ADAS is an extremely active and promising field of study across all levels: from the Level 1 of *Driver Assistance*, with possible intelligent intervention to steering, braking, and acceleration, to the futuristic Level 5 of *Full Automation*, where the vehicle can operate all driving conditions SAE On-Road Automated Vehicle Standards Committee and others (2014). Along these avenues, though, it has been recently highlighted how to unleash the full potential of autonomous driving: the smartness of the infrastructure is at least as important as the intelligence inside the vehicle itself Gopalswamy and Rathinam (2018); Joy and Gerla (2017); Mamduhi, Hashemi, Baras, and Johansson (2020). By *infrastructure* it is usually meant the information and communication systems distributed along the roads that will support the autonomous driving capabilities, but also the dynamic information-rich collaborating network that emerges from the many vehicles that are actually circulating Uhlemann (2015, 2018).

Such an information flow allows the definition of real-time maps that are needed for efficient and safe navigation: to this aim, mapping is an all-inclusive procedure that accounts for topology of roads, geometry of the environment, and more volatile data such as traffic levels, congestion or accident occurrences, and parking availability Wong, Gu, and Kamijo (2021). In order to acquire such mixture of permanent and temporary information, active sensing methodologies come into play, where the information acquisition process is controlled and driven by some optimal criteria chosen according to the task or the goal to be accomplished. Specifically, the vehicle can be interpreted as a mobile sensing platform that leverages environmental perception from on-board sensors, information coming from fixed infrastructure, and data exchange with other vehicles to perform scene reconstruction and understanding Siegel, Erb, and Sarma (2017). The combination with active sensing approaches allows to improve the quality of measurements and the efficiency of the information gathering process Varotto et al. (2021). This framework results particularly effective when the information of interest is highly dynamic, such

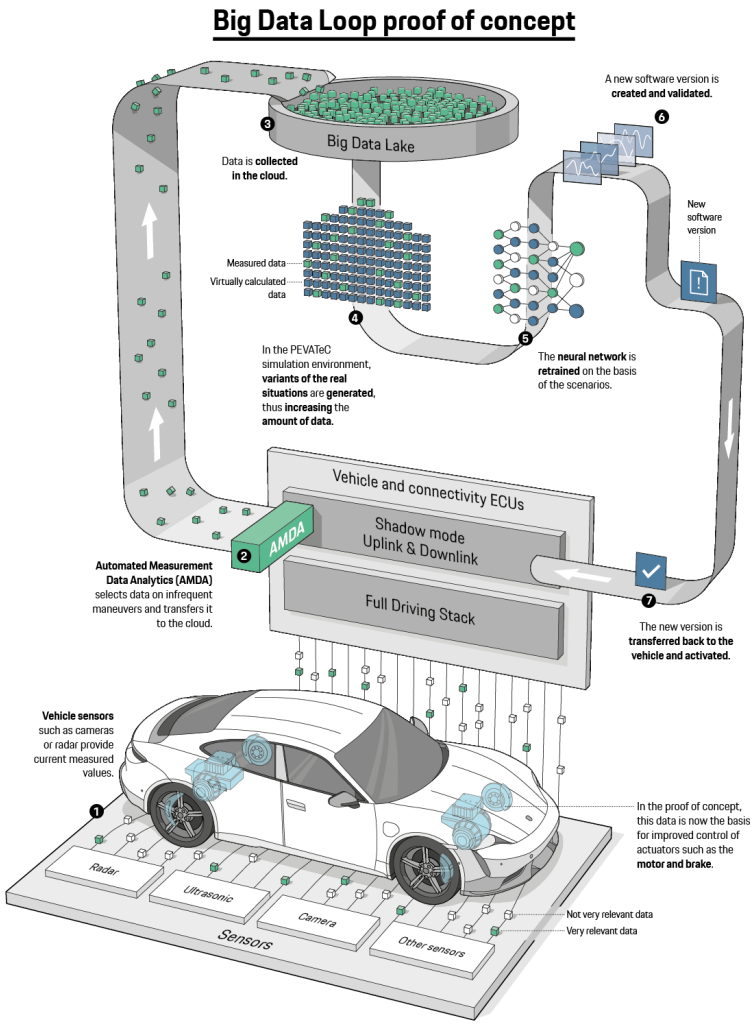

Figure 7.1. The Cariad's self-learning adaptive cruise control system is a prominent example of how IoT can be used within the framework of connect vehicles.

as that related to traffic condition and parking availability. In this case, the data coming from the network of connected vehicles provide high-fidelity real-time information that supports the driver's decision making process and improves the overall system mobility Talebpour and Mahmassani (2016).

As an example, Fig. 7.1[1] depicts the scheme of a self-learning adaptive cruise control system developed at Cariad (a subsidiary of the Volkswagen Group responsible for automotive software). Data from the vehicle is actively selected and transmitted wirelessly to the cloud, where it is used to adapt a neural network model the on the basis of the current road scenario. Afterwards, the improved algorithm is tested and integrated on a new software release, which is fed back again into the vehicle. For what concerns this Thesis, the most interesting module is the AMDA (Automated Measurement Data

[1] https://newsroom.porsche.com/en/2021/innovation/porsche-engineering-big-data-loop-25029.html.

Analytics), which is responsible for the data selection process. As a matter of fact, radar sensors and cameras generate an immense amount of data, most of which is not relevant to the function under consideration. Driving on an empty highway, for example, offers no learning opportunities for a distance controller. Moreover, evaluating or storing all the data would be far too time-consuming and memory demanding.

Parking Availability Mapping With such premises, the parking forecast problem can be seen as a probabilistic map building problem that can be addressed through an incremental information discovery process: namely, the problem is to estimate the parking availability in a certain objective area by gathering information along the path towards the desired place; the result of the estimation process is then used to take actions accordingly, in order to maximize the parking probability itself.

In the last few years, the strong demand for ADAS technologies, and the quest for the autonomous vehicles, have fuelled a growing literature on the design of strategies and models for smart parking systems Bischoff, Maciejewski, Schlenther, and Nagel (2018); Lin, Rivano, and Le Mouël (2017a). Indeed, smart vehicles equipped with assistive systems exploit semantic maps of the environment to navigate and operate. In Westfechtel, Ohno, Mizuno, Hamada, Kojima, and Tadokoro (2018), for example, a map of the environment is build from on-board sensors (basically LiDAR and GPS) and a graph-based approach is proposed to include topological and usage parking information, obtained also from multiple days observation data. This map is then used to evaluate the parking spots occupancy rate and therefore predict their availability. The method proposed in Westfechtel et al. (2018) detects and uses parked vehicles in the surroundings to estimate parking lot topology and infer vacant parking spots via a graph-based approach. The suggested solution is validated both in structured and unstructured parking lots. A collaborative approach is proposed in Li, Yang, Xiao, Valde, Wrenn, and Leflar (2018a) with the specific focus of the management of parking garages, where a global probabilistic occupancy map emerges from individual vehicle perception (e.g. through LiDAR, IMU, GPS) and inter-vehicle information sharing; the final objective is to support autonomous parking: from the available parking spot identification, to the navigation and motion planning, to the precise parking maneuvering. The multi-vehicle data collection is considered also in the ParkNet system Mathur, Jin, Kasturirangan, Chandrasekaran, Xue, Gruteser, and Trappe (2010), where the problem of urban street parking is addressed by means of drive-by-parking monitoring; the collected data from GPS and ultrasonic rangefinders are then fused into a single centralized map, which allows the generation of statistics for optimal management policies.

On the modeling aspects, the state of the art reports several research avenues. In particular, the employment of Gaussian Processes (GPs) to model occupancy maps appears as a convenient method to attain a probabilistic description of the behavior of mobile objects in the environments, given their physical spatial structure and correlation O'Callaghan, Ramos, and Durrant-Whyte (2009). Dynamic GP occupancy maps are further developed in Senanayake, O'Callaghan, and Ramos (2017), where variational inference methods are employed to speed up the estimation process and adapt to LiDAR produced dataflow. A different approach is taken by other models that instead of focusing on the occupancy map, consider the point of view of the vehicles that are moving towards the parking slots. For example, in Rajabioun and Ioannou (2015), a multivariate autoregressive model is

proposed where temporal and spatial correlations are taken into account to compute the probability of parking availability at the vehicle estimated arrival time; this system exploits real-time and historical data, where the latter are used to learn the parameters of the model and the former are used to carry out an accurate up-to-date prediction. Finally, queuing models are a popular and established method to describe the behavior of the vehicles-parking slots interaction and are employed in Tavafoghi, Poolla, and Varaiya (2019) to provide probabilistic estimation of parking availability, leveraging on learning techniques to attain the prediction model parameters.

Contributions This Section considers a smart transportation scenario, populated by smart infrastructures (I), as well as intelligent and connected vehicles (V) Uhlemann (2018). These are able to detect parking slots and to classify them as either free or occupied Westfechtel et al. (2018). They are also endowed with communication capabilities, so that V2V and V2I data exchange are allowed. In this framework, the purpose is the design of an online parking mapping scheme, which leverages on local and remote data to adapt to dynamic traffic conditions. To this aim, the task is formulated as a Gaussian Process Regression problem, coupled with an active sampling module; the latter is based on an information-seeking data selection criterion, which retains only those datapoints that are expected to reduce the uncertainty of the model.

This is the first attempt to design a parking mapping algorithm by leveraging active sensing within the framework of connected vehicles. In particular, with respect to the state of the art, the proposed methodology employs a data-driven approach (i.e., GPR) with the training dataset collected in a multi-source scenario, where V2V and V2I communication guarantees incremental learning and continuous map update. The computational complexity of most data-driven multi-source learning procedures increase without an appropriate control algorithm; therefore, the proposed active sensing module selects the most informative samples, thus minimizing information redundancy and avoiding memory and processing overload. Numerical experiments validate the effectiveness of the proposed algorithm against several baselines; according to them, the combination of GPR with active sensing improves the quality and the efficiency of the estimation process. In particular, the suggested approach attains high performance in terms of mapping convergence speed and adaptivity capabilities, at the cost of small computational demands.

7.2 Problem statement

Fig. 7.2 provides an overview of the problem scenario: a road path is populated by mobile sensing platforms (e.g., autonomous vehicles Bischoff et al. (2018)) with sensing, processing and communication capabilities on-board (i.e., intelligent and connected vehicles Siegel et al. (2017); Uhlemann (2018)). Each platform establishes communication with surrounding smart infrastructures (e.g., smart parking stations) and with the other vehicles, according to the V2I Uhlemann (2015) and V2V Krasniqi and Hajrizi (2016) paradigms. The platform has also access to cloud storage resources, where information on current traffic conditions can be downloaded. The purpose is to estimate the parking availability (PA) over the entire route. To this aim, informative data from on-board

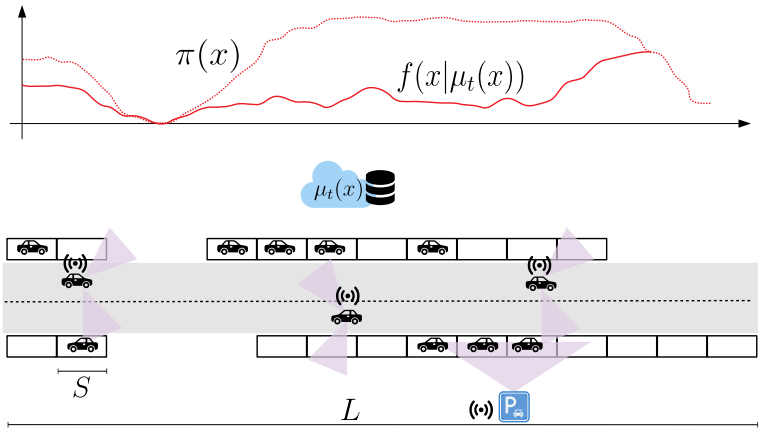

Figure 7.2. Problem scenario. A road path of length L is populated by intelligent and connected vehicles, as well as with smart parking infrastructures. Each vehicle is capable to detect parking slots and to classify them as free or occupied. Traffic conditions $\mu_t(x)$ are obtained from online cloud databases; hence, by establishing communication with external information sources (e.g, connected vehicles, and parking stations), each vehicle can learn the parking availability map $f\,(x|\mu_t(x))$, which differs from the a-priori parking availability $\pi(x)$, over the entire road path.

sensors and external resources need to be gathered, selected and processed.

One of the main objectives of this Chapter is to stress the flexibility and the generalization capabilities of the MS-PAS scheme introduced in Chapt. 3. In Chapt. 3-6 the targets were moving agents (e.g., people), but the MS-PAS scheme has been designed to accommodate for a more general concept of target, namely any entity whose properties are measured and estimated from platform observations. In this respect, the "target" considered in this Chapter is the environment (more specifically, the road), from which the parking availability level is estimated.

7.2.1 Environment

Consider a road path $\mathcal{X} = [0, L]$ over which the following function is defined

$$\pi(x) : \mathcal{X} \mapsto [0, 1]. \tag{7.1}$$

It represents the (on-street Mathur et al. (2010); Rajabioun and Ioannou (2015)) *a-priori parking availability* along the path $\mathcal{X}$. Being related to structural and environmental properties of the street in traffic-free conditions (e.g., presence and number of parking slots), it is a time-invariant quantity. In this work, $\pi(x)$ is the total count of slots over a spatial window W, namely

$$\pi(x) = \frac{1}{\lfloor W/S \rfloor} \sum_{h=0}^{\lfloor W/S \rfloor} \mathbb{1}(x - hS), \quad x \in \mathcal{X} \tag{7.2}$$

where S is the length of a parking slot (homogeneous slots are considered); $\lfloor W/S \rfloor$ is the maximum number of parking slots over the window W and it works as normalization factor; $\mathbb{1}(\cdot)$ is an indicator function, such that $\mathbb{1}(x - hS) = 1$ if there is a parking slot between

$x - hS$ and $x - (h+1)S$; $\mathbb{1}(x - hS) = 0$ otherwise. In practice, $\pi(x)$ is produced once and for all, via parking spot detection and mapping algorithms Westfechtel et al. (2018).

The road path is supposed to be characterized, at each time instant t, by a time-varying *traffic density* function

$$\mu_t(x) : \mathcal{X} \mapsto [0, 1]. \tag{7.3}$$

This function maps each point of $\mathcal{X}$ into a normalized density value, proportional to the traffic congestion level. It is realistic to assume that there exists a *parking availability attenuation function*, which depends on the current traffic density level Aryandoust, van Vliet, and Patt (2019); Bischoff et al. (2018); Cao and Menendez (2015), namely

$$\lambda\left(\mu_t(x)\right) : [0, 1] \mapsto [0, 1]. \tag{7.4}$$

This function defines the parking availability reduction according to the traffic density $\mu_t(x)$ at a certain location x. Hence, the *parking availability map* (PAM) (also referred to as occupancy map Westfechtel et al. (2018)) is computed as

$$f\left(x|\mu_t(x)\right) : \ \mathcal{X} \mapsto [0, 1] \tag{7.5}$$

$$x \to \pi(x)\lambda\left(\mu_t(x)\right), \tag{7.6}$$

where notation $f\left(x|\mu_t(x)\right)$ means that $f(x)$ is computed, given $\mu_t(x)$[2]. By aggregating information on environmental properties through $\pi(x)$, and on current traffic conditions through $\mu_t(x)$, the PAM represents the number of available parking slots between x and $x - W$, normalized by the maximum number of parking slots (i.e., $\lfloor W/S \rfloor$). In fact, it is computed as

$$f\left(x|\mu_t(x)\right) = \frac{1}{\lfloor W/S \rfloor} \sum_{h=0}^{\lfloor W/S \rfloor} \mathbb{1}_{\mu_t}(x - hS), \quad x \in \mathcal{X} \tag{7.7}$$

where $\mathbb{1}_{\mu_t}(x - hS)$ is an indicator function, such that, at time t, $\mathbb{1}_{\mu_t}(x - hS) = 1$ if between $x - hS$ and $x - (h+1)S$ there is a parking slot and it is available; $\mathbb{1}_{\mu_t}(x - hS) = 0$ otherwise.

7.2.2 Sensing platform

This work leverages on the active sensing framework to solve the parking availability estimation problem Mathur et al. (2010); therefore, the sensing platform is one of the main elements of the problem scenario. In particular, the platform interacts with the surrounding environment and, by actively selecting the incoming information, it can reconstruct the PAM.

Platform observation model The sensing platform is mobile and equipped with communication capabilities (i.e., it is a *connected vehicle* Uhlemann (2018)). It also embeds sensing, computational and logical resources on-board, so that it is capable to detect on-street parking slots and to recognize their availability, as in Bischoff et al.

[2]For ease of notation, $f\left(x|\mu_t(x)\right)$ is sometimes referred also as $f\left(x\right)$.

(2018). Thus, a platform measurement is

$$z_t = \frac{1}{\lfloor W/S \rfloor} \sum_{h=0}^{\lfloor W/S \rfloor} \tilde{\mathbb{1}}_{\mu_t}(s_t - hS), \tag{7.8}$$

where $s_t \in \mathcal{X}$ is the platform state at time t, represented by the vehicle location over the road path. The vehicle is supposed to move according to a (possibly) non-linear deterministic Markovian dynamics

$$s_{t+1} = \begin{cases} s_t + v(\mu_t(x))T, & t = LT,\ L \in \mathbb{N} \\ s_t, & \text{otherwise,} \end{cases} \tag{7.9}$$

where T is the transition time of the discretized vehicle model, and $v(\mu_t(x))$ is the car velocity influenced by the current traffic level. In (7.8), the binary function $\tilde{\mathbb{1}}_{\mu_t}(s_t - hS) \in \{0, 1\}$ is the *estimated occupancy*: $\tilde{\mathbb{1}}_{\mu_t}(s_t - hS) = 1$ if the platform classifies the slot between $s_t - hS$ and $s_t - (h+1)S$ as available; $\tilde{\mathbb{1}}_{\mu_t}(s_t - hS) = 0$ otherwise. To account for sensors nuisance and for the estimation uncertainty associated to common parking slot classification algorithms Bischoff et al. (2018), $\tilde{\mathbb{1}}_{\mu_t}(s_t - hS)$ is a noisy version of the true value $\mathbb{1}_{\mu_t}(s_t - hS)$. In particular, we consider the following *PA observation model*

$$\begin{aligned} z_t = & \underbrace{\frac{1}{\lfloor W/S \rfloor} \sum_{h=0}^{\lfloor W/S \rfloor} \mathbb{1}_{\mu_t}(s_t - hS)(1 + \epsilon_1)}_{z_{t,1}} \\ & + \underbrace{\frac{1}{\lfloor W/S \rfloor} \sum_{h=0}^{\lfloor W/S \rfloor} (1 - \mathbb{1}_{\mu_t}(s_t - hS))\, \epsilon_0}_{z_{t,0}} \end{aligned} \tag{7.10}$$

where ϵ_0 and ϵ_1 are noisy terms: $\epsilon_1 \sim \mathcal{B}(p_1)$ with support in $\{-1, 0\}$, while $\epsilon_0 \sim \mathcal{B}(p_0)$ with support in $\{0, 1\}$. In practice, ϵ_1 induces a free slot to be either not detected or detected as occupied; hence, it accounts for mis-detections and mis-classifications. On the other hand, ϵ_0 induces an absent or occupied slot to be seen as a free parking space; thus, it accounts for sensor artifacts and mis-classifications as well. From the Bernoulli distribution of ϵ_0 and ϵ_1, it follows that $z_{t,0}$ and $z_{t,1}$ are Binomial random variables; by exploiting the Gaussian approximation of Binomial random variables, the following approximation holds

$$\begin{aligned} & z_t \approx h(\mu_t, s_t, \epsilon_t) = f(s_t | \mu_t(s_t)) + \epsilon_t \\ & \epsilon_t \sim \mathcal{N}\left(\epsilon \,\middle|\, 0, \sigma^2\right). \end{aligned} \tag{7.11}$$

Note that the observation model (7.11) is linear, with noise component concentrated in the Gaussian term ϵ_t[3].

External information sources In line with the IoT paradigm, the platform is supposed to receive information from *external sources*, such as other connected vehicles, cloud databases, and smart parking stations Krasniqi and Hajrizi (2016); Uhlemann

[3]Later on, the dependence of z_t on the platform state will be made clear through the notation $z_t(s_t)$, while the dependence on the current traffic conditions and on the observation noise is left implicit.

(2018). The set of external sources that establish communication with the platform at time t is denoted as

$$\mathcal{V}_t = \{v_{t,i}\}_{i=1}^{N_t}, \tag{7.12}$$

while the location of the i-th source is indicated with $x_i \in \mathcal{X}$. As (7.12) suggests, the number of communicating agents is a function of time, as in common real-life scenarios. In fact, smart parking stations and connected cars may have a limited communication range; therefore, two communicating (resp. non communicating) platforms at time t may lose (resp. establish) connection at time $t+T$, due to their relative movements. Moreover, the number of platforms travelling along $\mathcal{X}$, and endowed with compatible communication technologies, is non-constant in time (in principle, it should follow the traffic density trends). For the aforementioned reasons, the i-th source at time t may be different from the i-th source at different instants of time[4]. It is also important to remark that the vehicle is capable to receive a limited amount of N_{max} data at each iteration t, that is

$$0 \leq N_t \leq N_{max}. \tag{7.13}$$

Each agent in $\mathcal{V}_t$ sends to the platform the measured PAM at its location, namely $z_t(x_i),\ i = 1, \ldots, N_t$, where each $z_t(x_i)$ follows (7.11).

The platform is also supposed to have access to an external database from which perfect information on the traffic density level over the entire path can be gathered; in other words, at each time instant t, the platform knows $\mu_t(x)$ for each $x \in \mathcal{X}$, thanks to cloud resources[5].

7.2.3 Problem statement

Given the premises provided in this Section, the platform can accumulate large amount of data over time; in particular, if no data selection strategies are applied, at time t the platform stores the following dataset

$$\mathcal{D}_t = (\mathbf{X}_t, \mathbf{y}_t)\,. \tag{7.14}$$

This is incrementally fed with new incoming data, namely

$$\begin{aligned} \mathbf{X}_t &= \begin{bmatrix} \mathbf{X}_{t-T_S} & s_t & x_1 & \ldots & x_{N_t} \end{bmatrix} \\ \mathbf{y}_t &= \begin{bmatrix} \mathbf{y}_{t-T_S} & z_t(s_t) & z_t(x_1) & \ldots & z_t(x_{N_t}) \end{bmatrix}. \end{aligned} \tag{7.15}$$

where T_S is the sampling time, that is, between $t - T_S$ and t the platform collects a new measurement, $z_t(s_t)$, accumulates the data received from the external sources i.e., the set of data

$$\{(x_i, z_t(x_i))\}_{i=1}^{N_t}, \tag{7.16}$$

and updates the knowledge on $\mu_t(x)$ over $\mathcal{X}$.

[4]This explains why the variable x_i has not reference to time: it is not meant to describe the spatio-temporal trajectory of a specific source.

[5]If the platform does not have access to any cloud database, the traffic state can also be reconstructed exploiting the measurements coming from on-street sensors and connected vehicles along the road Liu, Barreau, Čičić, and Johansson (2021a).

Data selection as an active sensing problem According to (7.15), the cardinality of $\mathcal{D}_t$ is

$$\begin{aligned} |\mathcal{D}_t| &= |\mathcal{D}_{t-1}| + (1 + N_t) \\ &= t + \sum_{\tau=1}^{t} N_\tau, \end{aligned} \tag{7.17}$$

which linearly increases as the time grows. Consequently, to operate with the limited on-board computational and storage resources of the platform, efficient data collection, filtering and management schemes are required. In this work filtering is cast to an active sampling problem (see Sec. 7.3), which requires the introduction of the *extended platform state*, namely

$$\bar{\mathbf{s}}_t = \begin{bmatrix} s_t & \mathbf{b}_t \end{bmatrix}^\top. \tag{7.18}$$

The platform position is coupled with the *choice state* $\mathbf{b}_t$: a $(N_{max}+1)$-dimensional boolean vector (i.e., $\mathbf{b}_t \in \{0,1\}^{(N_{max}+1)}$) used to indicate which information sources are considered by the platform at time t; formally, it is defined as

$$\mathbf{b}_t(\ell) = \begin{cases} 1, & \text{if } \ell = 1 \text{ and } z_t(s_t) \text{ is selected} \\ 0, & \text{if } \ell = 1 \text{ and } z_t(s_t) \text{ is not selected} \\ 1, & \text{if } \ell \in]1, N_t + 1] \text{ and } z_t(x_\ell) \text{ is selected} \\ 0, & \text{if } \ell \in]1, N_t + 1] \text{ and } z_t(x_\ell) \text{ is not selected} \\ 0, & \text{if } \ell > N_t + 1. \end{cases} \tag{7.19}$$

The first entry indicates if the platform retains its own measurement; from the second to the $(N_t + 1)$-th entry, $\mathbf{b}_t$ indicates which external source is used to update the dataset; the remaining entries are set to 0, since no more observations are available, From this definition, the selected data at time t are

$$\mathcal{Z}_t(\mathbf{b}_t) = \begin{cases} z_t(s_t) \cup \bigcup_{\ell:\mathbf{b}_t(\ell)=1} z_t(x_\ell), & \text{if } \mathbf{b}_t(1) = 1 \\ \bigcup_{\ell:\mathbf{b}_t(\ell)=1} z_t(x_\ell), & \text{otherwise.} \end{cases} \tag{7.20}$$

Accordingly, the selected sources are

$$\chi_t(\mathbf{b}_t) = \begin{cases} s_t \cup \bigcup_{\ell:\mathbf{b}_t(\ell)=1} x_\ell, & \text{if } \mathbf{b}_t(1) = 1 \\ \bigcup_{\ell:\mathbf{b}_t(\ell)=1} x_\ell, & \text{otherwise.} \end{cases} \tag{7.21}$$

Consequently, (7.15) can be re-formulated using sets $\chi_t(\mathbf{b}_t)$ and $\mathcal{Z}_t(\mathbf{b}_t)$, as

$$\mathcal{D}_t = \mathcal{D}_{t-T_S} \cup (\chi_t(\mathbf{b}_t), \mathcal{Z}_t(\mathbf{b}_t)). \tag{7.22}$$

Finally, it is possible to cast the data selection process into a data acquisition control problem (i.e., active sensing); this is done by associating the following dynamics to the choice state

$$\mathbf{b}_t = \mathbf{b}_{t-T_S} + \bar{\mathbf{u}}_t, \tag{7.23}$$

with the control input $\bar{\mathbf{u}}_t$ disabled between two sampling instants, that is[6]

$$\bar{\mathbf{u}}_t = \mathbf{0}, \quad t \neq HT_S,\ H \in \mathbb{N}. \tag{7.24}$$

The control input $\bar{\mathbf{u}}_t$ is suitably chosen to modify $\mathbf{b}_{t-T_S}$, on the basis of which information sources to select at time t. In this way, it is possible to choose what to sense, according to the active sensing paradigm, as *selection represents an integral process of active sensing*, according to Bajcsy et al. (2018).

The PAME problem With the formalism introduced in this Section, the *parking availability map estimation (PAME)* problem can be formulated as the control of the platform choice state $\mathbf{b}_t$ (through $\bar{\mathbf{u}}_t$), given the dataset $\mathcal{D}_t$. The objective is to reconstruct the latent PAM function $f(x|\mu_t(x))$, over $\mathcal{X}$. Moreover, the problem is solved *online* when the reconstructed function, $\hat{f}(x|\mu_t(x))$, is updated as new samples are collected. Finally, a PAME technique is considered *adaptive* if it can handle the time-varying fluctuations of the traffic density level, namely it is capable to deal with the condition $\mu_t(x) \neq \mu_{t-T_S}(x)$.

7.3 Methodology

The methodology adopted to solve the PAME problem leverages on a GPR scheme that incrementally and adaptively learns the underlying function $f(x|\mu_t(x))$, exploiting the dataset $\mathcal{D}_t$. As GPR strategies do not scale well with the cardinality of the dataset Das, Roy, and Sambasivan (2018); Liu, Ong, Shen, and Cai (2020); Moore, Chua, Berry, and Gair (2016); Shahriari, Swersky, Wang, Adams, and De Freitas (2015), the learning process is combined with an active sensing scheme to select the incoming samples according to an uncertainty-aware acquisition policy.

To sum up, the proposed online and adaptive PAME algorithm is based on the following three steps.

1. The core of the algorithm is the exploitation of $\mathcal{D}_t$ to learn $f(x|\mu_t(x))$ via GPR (the reader may refer to Appendix A.1 for an overview in GPR).

2. To preserve the on-board storage and processing resources an active data selection procedure is applied; in particular, the data acquisition task is treated as an active sensing (active sampling) problem, according to the formulation introduced in Sec. 7.2.3. In this framework, an information-seeking data selection criterion is employed to retain only those datapoints that are expected to reduce the uncertainty of the model.

3. Finally, adaptation to time-varying traffic density is achieved by leveraging on cloud resources; these allow to detected and remove obsolete data in $\mathcal{D}_t$. In this way, it is possible to prevent data contamination issues during the GPR learning process.

Alg. 1 provides a schematic overview of the suggested solution, while the following Sections describe in detail each of the aforementioned steps.

[6] Note that, in general, $T \neq T_S$; hence, the two elements of the extended platform state, namely s_t and $\mathbf{b}_t$, update at different frequencies. This is due to the fact that s_t is related to the dynamics of the vehicle, while $\mathbf{b}_t$ refers to the data collection process.

Algorithm 1: Online and adaptive PAME.

Initialization:

- $s_0 = 0$
- $\mathcal{D}_0 = (\varnothing, \varnothing)$

while $s_t \leq L$ **do**
- update knowledge on $\mu_t(x)$
- compute $\mathcal{D}_{obsolete}$
- **if** $\mathcal{D}_{obsolete} \neq \varnothing$ **then**
  - $\mathcal{D}_{t-T_S} \leftarrow \mathcal{D}_{t-T_S} \setminus \mathcal{D}_{obsolete}$
- **end**
- measure $z_t(s_t)$
- receive $\{(x_i, z_t(x_i))\}_{i=1}^{N_t}$
- solve (7.28) $\rightarrow \bar{\mathbf{u}}_t^*$
- solve (7.23) with $\mathbf{u}_t = \bar{\mathbf{u}}_t^* \rightarrow (\chi_t(\mathbf{b}_t), \mathcal{Z}_t(\mathbf{b}_t))$
- apply (7.22) $\rightarrow \mathcal{D}_t$
- perform GPR on $\mathcal{D}_t \rightarrow \hat{f}(x|\mu_t(x))$
- $s_{t+T_S} = s_t + v(\mu_t(x))T_S$
- $t \leftarrow t + T_S$

end

Output:

- $\hat{f}(x|\mu_t(x)) \sim \mathcal{GP}\left(0, k(x, x')\right)$

7.3.1 GPR-based PAME

Consider the PAME framework described in Sec. 7.2 and suppose a time-invariant traffic density function, i.e.,

$$\mu_{t-T_S}(x) = \mu_t(x), \ \forall x \in \mathcal{X}, \ \forall t. \tag{7.25}$$

At any time instant t, it is possible to employ GPR to learn the PAM function $f(x|\mu_t(x))$. To this aim, $\mathcal{D}_t$ can be used as training dataset, while the underlying PAM can be modelled as a GP with zero mean function and Matern covariance function $k(x, x')$ (with length scale hyperparameter set to 1 and smoothness hyperparameter set to 1.5), that is[7]

$$\widehat{f}(x|\mu_t(x)) \sim \mathcal{GP}\left(0, k(x, x')\right). \tag{7.26}$$

Note that the GPR problem is well-posed Williams and Rasmussen (2006), according to (7.11).

7.3.2 Uncertainty-aware active sampling

In active sensing, attention is a set of mechanisms that tune and control the data acquisition process. One important aspect of this tuning is to select which portion of the input to process at any time Active sensing systems can also select the manner in which this input is processed, suppressing irrelevant computations and enhancing the relevant ones, improving responses to task-relevant characteristics, and sharpening

[7] The notation $\hat{f}(x)$ is used to distinguish the GP model w.r.t. the underlying one, $f(x)$.

decision processes. This is the essence of attentional behavior and active sampling, which represents an important subset of the full range of active sensing techniques.

With reference to the PAME task, it can theoretically be solved with the iterative GPR procedure outlined in Sec. 7.3.1: as a new datapoint $z_t(s_t)$ is collected, it is included in the dataset $\mathcal{D}_{t-T_S}$, together with the samples coming from the external sources, as in (7.15); at this point, a new GPR process is performed on $\mathcal{D}_t$. In this way, as $|\mathcal{D}_t|$ increases, the GPR converges to an optimal reconstruction of the latent function Carron et al. (2015); Shahriari et al. (2015). Nonetheless, one of the main drawbacks of GPR is the poor scalability properties with respect to the dataset cardinality Das et al. (2018); Liu et al. (2020); Moore et al. (2016); Shahriari et al. (2015). In particular, the PAME framework induces a fast intractability in the learning process, since the dataset gains $N_t + 1$ new data at each iteration, as (7.17) shows.

Given these premises, it is here proposed an active sampling scheme to select which incoming data (among the $N_t + 1$) are kept in memory (i.e., included in the dataset) and used in the successive GPR iterations. Motivated by the fast dataset growth and the good estimation properties of GPs with small datasets Williams and Rasmussen (2006), it is sufficient to consider only one sample per iteration, according to the following *uncertainty-aware policy*

$$a(x_*|\mathcal{D}_{t-T_S}) = \sigma_{*,\mathcal{D}_{t-T_S}} \tag{7.27}$$

where (A.8) is applied on the query point x_* and $\sigma_{*,\mathcal{D}_{t-T_S}}$ is $\mathbf{\Sigma}_{*,\mathcal{D}_{t-T_S}}$ with $p = 1$. The predicted variance is a measure of the process entropy Viseras, Wiedemann, Manss, Magel, Mueller, Shutin, and Merino (2016); hence, by choosing the source whose position is characterized by the highest predicted variance, it is possible to sample where the entropy is highest, that is, at the most informative location. This leads the selection process to be achieved through the following control law

$$\bar{\mathbf{u}}_t^* = \left[\underset{\substack{\mathbf{b}\in\{0,1\}^{(N_{max}+1)} \\ \sum_{\ell=1}^{(N_{max}+1)} \mathbf{b}(\ell)=1}}{\arg\max} \; a(\chi_t(\mathbf{b})|\mathcal{D}_{t-T_S}) \right] - \mathbf{b}_{t-T_S}, \quad t = HT_S,\ H \in \mathbb{N}. \tag{7.28}$$

The control input $\bar{\mathbf{u}}_t^*$ is applied to compute the next choose state, $\mathbf{b}_t$, according to (7.23); in turn, $\mathbf{b}_t$ allows to select the next query point $\chi_t(\mathbf{b}_t)$, as well as its correlated measurement, $\mathcal{Z}_t(\mathbf{b}_t)$. Note that $|\chi_t(\mathbf{b}_t)| = |\mathcal{Z}_t(\mathbf{b}_t)| = 1$, due to the condition $\sum_{\ell=1}^{(N_{max}+1)} \mathbf{b}(\ell) = 1$.

In Sec. 7.4 the active sampling module is shown to remarkably decrease the computational load of the map learning process; in particular, the acquisition function (7.27) leads to fast convergence and adaptability.

7.3.3 Adaptation to time-varying traffic density

So far, a time-invariant traffic density function has been considered. Nonetheless, real-life scenarios are characterized by time-varying traffic densities, as stated in Sec. 7.2. Consequently, some data in $\mathcal{D}_t$ become obsolete as $\mu(\cdot)$ changes over time. More formally,

$$\mathcal{D}_{obsolete} = \{(x_j, y_j) \in \mathcal{D}_t : \mu_{t-T_S}(x_j) \neq \mu_t(x_j)\} \subseteq \mathcal{D}_t \tag{7.29}$$

is the sub-dataset of obsolete data at time t, that is, the set of locations (and respective measurements) where the traffic density has changed during the last sampling period.

As mentioned in Sec. 7.2, the platform captures information on the current traffic conditions from cloud databases. Thus, it is capable to compute $\mathcal{D}_{obsolete}$ and, therefore, it can remove obsolete data from the current dataset (see Alg. 1); in this way, the successive GPR iterations are not to contaminated by spurious information. This technique, combined with the multi-source data collection technique described previously, allows to achieve reactive and adaptive mapping performance under dynamic traffic conditions (see Sec. 7.4).

7.4 Numerical results

The proposed approach has been evaluated through a Python-based synthetic environment[8]. Sec. 7.4.1 describes the main setup parameters, as well as the synthetic data generation process. Sec. 7.4.2 defines the metrics used for performance assessment and the baselines considered for comparison. The numerical simulation results are discussed in Sec. 7.4.3.

7.4.1 Setup parameters

In the synthetic environment the platform is required to travel a length of $L = 10\,\text{km}$ with a velocity $v(\mu_t(x))$, function of the local traffic density conditions (see Tab. 7.1). The traffic density is a piece-wise constant function where each constant portion is 1 km-long and with a value randomly chosen between 0 and 1. The time-variance is simulated through a Bernoulli random variable of parameter $p_{change} = 0.2$, which is used to change the value of one of the constant portions (the segment is randomly chosen, as well as its new value). The a-priori parking availability function is generated according to its definition (7.2), where $W = 100\,\text{m}$ and the indicator function $\mathbb{1}(x)$ is a binary piece-wise constant function where each portion is 5 m-long (i.e., $S = 5\,\text{m}$) and whose binary value is chosen randomly. The number of connected external sources at time t is modeled as $N_t \sim \mathcal{U}(0, N_{max})$, with $N_{max} = 10$. Measurements are simulated according to the model (7.11) with $\sigma = 3 \times 10^{-2}$. For all the remaining relevant parameters, refer to Tab. 7.1.

7.4.2 Performance assessment

To capture the simulation and performance variability, numerical evaluation is performed through a Monte Carlo (MC) experiment, composed by $N_{tests} = 10$ tests.

Baselines The following baselines are considered for comparison.

- *NoCom*: the platform is a non-connected vehicle; hence, it relies on its measurements only, i.e.,

$$\mathcal{D}_t = \left(s_{1:t}, \{z_k(s_k)\}_{k=1}^{t}\right), \ \forall t. \tag{7.30}$$

[8] https://github.com/luca-varotto/Pparking

Table 7.1. Setup parameters for the MC experiment.

Parameter	Value
$\lambda(\mu)$	$(1+e^{20(\mu-0.5)})^{-1}$
$v(\mu)$	$90e^{-\mu}$km/h
L	10 km
N_{tests}	10
N_{max}	10
N_t	$\sim \mathcal{U}(0, N_{max})$
p_{change}	0.2
T	10 s
S	5 m
W	100 m
σ	3×10^{-2}

- *NoSel*: the platform is a connected vehicle, but the active sampling module is not applied (i.e., $\mathcal{D}_t$ is updated with all N_t available data).
- *Rnd*: the platform is a connected vehicle and the active sampling module uses a random acquisition function; more formally, the choice state is not computed as in (7.28), but according to

$$\mathbf{b}_t(\ell) = \begin{cases} 1, & \ell = \ell_{Rnd} \\ 0, & \text{otherwise} \end{cases} \tag{7.31}$$

  with $\ell_{Rnd} \sim \mathcal{U}(1, N_t + 1)$.

Fig. 7.3 shows a realization of $\pi(x)$, with the corresponding PAM function to be estimated, according to the synthetic generation procedure detailed in Sec. 7.4.1. As an immediate observation, it can be seen that *noCom* attains good estimation performance only on a small neighbourhood of the current platform position ($s_t = 4.1$ km). In principle, its training dataset (7.30) should induce an accurate GP model estimate for any $x \leq s_t$ (for all past locations). Nevertheless, many of these datapoints become obsolete over time, due to the dynamic traffic conditions; thus, the accuracy remains good only locally (i.e., in a neighbourhood of s_t, where samples are more recent). To sum up, *noCom* has poor adaptivity capabilities, since obsolete data can not be replaced by more up to date samples; for this reason, *noCom* is not suitable to accurately estimate the latent function over the entire domain $\mathcal{X}$.

Performance metrics The following performance indexes are computed from the MC simulation.

- *Learning curve*: the estimation quality at time t is estimated through the Root Mean Squared Error (RMSE) between the underlying function $f(\cdot|\mu_t(\cdot))$ and its estimated version $\hat{f}(\cdot|\mu_t(\cdot))$, i.e.

$$RMSE_t = \sqrt{\frac{1}{L}\int_0^L \left[\hat{f}(x|\mu_t(x)) - f(x|\mu_t(x))\right]^2 dx}. \tag{7.32}$$

  The learning curve is defined as $RMSE_t/RMSE_0$, where $RMSE_0$ is computed before any measurement is collected (i.e., during model initialization) and it is

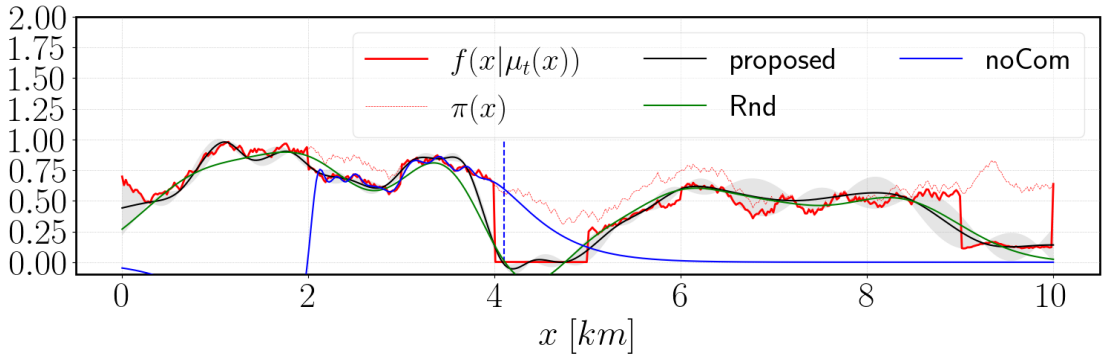

Figure 7.3. Example of PA curves. Synthetic PAM $f(x|\mu_t(x))$ (red) with its a-priori availability function $\pi(x)$ (red dashed), plotted against the estimated $\hat{f}(\cdot|\mu_t(\cdot))$ (with values not constrained in the interval $[0, 1]$): the proposed approach (black) is compared with *Rnd* (green) and *noCom* (blue). The gray shaded area is the uncertainty associated to the GP model obtained with Alg. 1. The current platform position (vertical and dashed blue line) is $s_t = 4.1\,\text{km}$ over the $L = 10\,\text{km}$ path.

equal for all baselines, since the initial model parameters are fixed.

- *Processing time ratio*: it is computed as

$$\tau_t = \frac{T_{C,t}^{(proposed)}}{T_{C,t}^{(noSel)}} \tag{7.33}$$

where $T_{C,t}^{(proposed)}$ and $T_{C,t}^{(noSel)}$ are the computation time for the proposed algorithm and *noSel*, respectively; both are referred to an entire GPR iteration: model training and prediction.

7.4.3 Discussion

Time-invariant traffic This paragraph considers the simplified scenario in which the traffic density is time-invariant; this allows to analyze the asymptotic estimation properties of the methods under comparison. As depicted in Fig. 7.4, all the methods

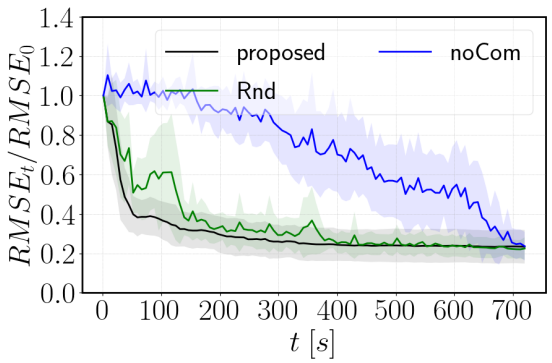

Figure 7.4. MC numerical results. Learning curve in case of time-invariant traffic: average (lines) and 68% confidence interval (shaded areas), plotted as a function of time over the whole duration of the L-long travel.

end up with the same estimation quality, due to the fact that the traffic density does not change in time and, at the end of the route, all methods have gathered the same amount of information. From one side, this means that the multi-source approaches (i.e., the proposed algorithm and *Rnd*) store redundant data, as clearly shown in Fig. 7.4: the two learning curves are characterized by a very slow decay after an initial dramatic drop. Thus, most of the PAM learning process is completed after 200 s (approximately); then, some minor corrections to the model parameters are executed. Even though multi-source approached induce a certain level of redundancy, they also attain a faster and smoother convergence. In particular, during the entire MC experiment, the RMSE ratio between the proposed approach and the *noCom* baseline is smaller than 1 for the 96% of the times (the remaining 4% is caused by small oscillations at steady-state), while it is smaller than 1 for the 66 of the times, if compared with *Rnd*. Therefore, the main benefit of an uncertainty-aware acquisition function is that it drives the data collection process towards the highest information gain, ensuring fast and high quality mapping performance. On the other hand, the absence of any communication with the external sources (i.e., *noCom*) induces the slowest convergence. In particular, the platform does not have spatial predictive capabilities; hence, it needs to reach the end of the path to have enough information to accurately reconstruct the latent function.

Impact of uncertainty-aware sampling The main benefits of the proposed strategy stand out when the traffic conditions are time-varying. The local information exploited by *NoCom* and its poor adaptivity capabilities (see Sec. 4.4.2), make this algorithm not suitable for dynamic traffic conditions, as shown in Fig. 7.5. Therefore, in real-life scenarios it is fundamental to exploit information coming from external sources. In particular, the fast adaptation capabilities of multi-source data collection, coupled with an active sampling strategy, guarantees the highest estimation performance. As described in Sec. 7.3.3 and in Alg. 1, when traffic density variations are detected, the corresponding obsolete datapoints are removed from the dataset. This increases the uncertainty of the GP model increases around the locations of the removed samples; consequently, the policy in (7.27) prioritises new query points coming from those regions. This leads to a faster recovery of the estimation process w.r.t. to a random sampling process. Indeed, Fig. 7.5 shows a remarkable performance difference between the proposed algorithm and the baseline *Rnd*; more specifically, the ratio between the RMSE of the proposed

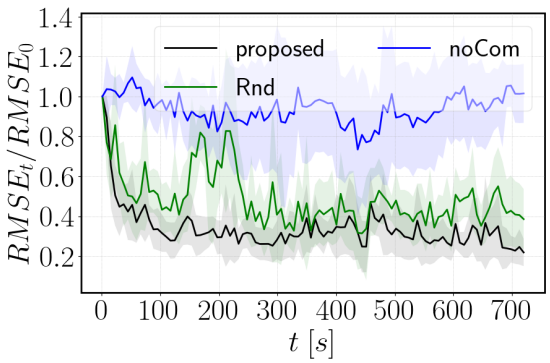

Figure 7.5. MC numerical results. Learning curve in case of time-varying traffic: average (lines) and 68% confidence interval (shaded area), plotted as a function of time over the whole duration of the L-long travel.

algorithm and that of *Rnd* is smaller than 1 for the 80% of the times (14% more than the time-invariant case previously described).

Impact of data selection Fig. 7.6 shows the processing time ratio; it highlights the benefit of performing data selection when GPR meets big data. Indeed, the time required to fit the GP model, and to estimate the underlying function, is 83% of the times smaller when using the proposed approach rather than a non selective strategy (i.e., *noSel*). Notably, the computational gap between the two methods increases (on average) as the dataset increases, as expected from (7.17).

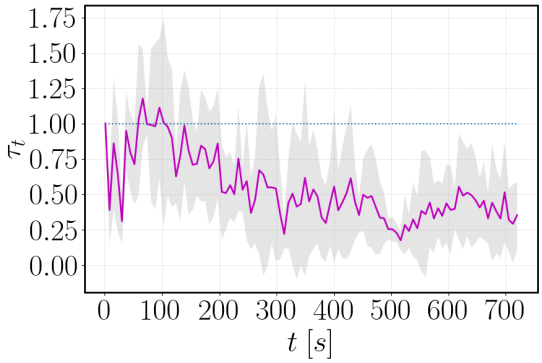

Figure 7.6. MC numerical results. Processing time ratio in case of time-invariant traffic: average (lines) and 68% confidence interval (shaded area), plotted as a function of time over the whole duration of the L-long travel.

7.5 Conclusion

This Chapter proposes an online and adaptive learning technique for estimating the on-street parking probability in the framework of connected vehicles. The parking availability map is incrementally learnt via GPR in a multi-source data collection scenario. To prevent computational and storage intractability issues, an information-driven active sensing module is applied to select incoming data. Numerical results validate the proposed approach in terms of adaptation capabilities, learning speed, and execution time.

The scope of this Thesis is to provide a prototypical solution to the PAME problem on a synthetic environment. To exhaustively validate the proposed algorithm, it should be tested in a real scenario which, however, might be challenging: at first, the setup requires a complex infrastructure, with reliable V2V and V2I communication; moreover, in real-life scenarios the perfect knowledge on the traffic conditions is usually not guaranteed; finally, dynamic and cluttered environments might affect the platform sensing performance, making the observation model (7.11) not applicable.

8

Conclusions and future works

Contents

This Thesis introduces a novel active perception paradigm, where sensor fusion and intelligent control play a fundamental role. Specifically, the aim of the dissertation is to develop a *unified*, *modular*, and *general* methodological framework that can be easily deployed on a wide variety of active sensing scenarios. To validate the MS-PAS framework proposed in Chapt. 3, several applications have been presented, including position estimation, video surveillance, and parking estimation in connected vehicles.

The following paragraphs summarize the main results discussed throughout this dissertation, with some conclusive considerations on the future trends in the field of smart and active sensing research.

8.1 Concluding Remarks and Summary of Contributions

In this Thesis, the main research interests concern the development of a multi-sensor probabilistic active sensing (MS-PAS) paradigm, which can serve as methodological reference for future research in active sensing applications. Specifically, the MS-APS framework described in Chapt. 3 is designed to satisfy the following properties.

- **Modularity**: the MS-PAS scheme satisfies two levels of modularity, namely *internal modularity* and *external modularity*. The former defines the possibility of aggregating a different number of internal platform components without changing the overall structure of the scheme; for instance, it is possible to change the number of sensing units, leaving all the other modules unchanged. On the other hand, external modularity is related to the number of platforms employed and to the level of coordination among them.

- **Generality**: the MS-PAS scheme has been designed with a high level of abstraction; hence, it can work as flexible reference basis from which specific implementations can take place, both in the formulation and in the task definition. In fact, each module of the MS-PAS scheme can be customized according to specific task requirements and constraints.

- **Integration**: MS-PAS is a unified scheme including different operating modules into a single system: from control, actuation and sensing, to data communication, processing and fusion.

The second objective of this work is to validate the MS-PAS scheme through different application and case studies, which allow to stress and test the properties enlisted above. From Chapt. 4 to 7, four applications have been proposed; each of them modifies the general MS-PAS scheme according to the application-specific setup; nonetheless, the overall operational behavior is left unchanged and represents the common thread among the four applications, despite their intrinsic differences.

In this regard, the main contributions of this manuscript are the following.

1. The design of the MS-PAS scheme, a unified methodological framework that allows to incorporate multi-sensory data into a probabilistic active sensing scheme. The proposed infrastructure is formulated in a general manner, making it easily adapted for any active sensing problem.

2. Implementation of the MS-PAS scheme in four different case studies: localization, tracking, video surveillance, and parking estimation for connected vehicles. However, the applicability of MS-PAS extends to other areas of robotics and artificial intelligence, such as environmental mapping, source term estimation and visual object recognition.

3. Throughout this dissertation, a particular attention is devoted the management of limited resource capabilities, and to the presence of measurement uncertainties and unstructured environments Calli et al. (2018); all these situations are handled through Bayesian reasoning and learning techniques.

4. Unlike conventional active sensing schemes, this work proposes a comparative analysis of different solutions; the theoretical and numerical investigations underline the ability of multi-modal and multi-sensor PAS schemes to refine incoming measurements, learn the sensors observation models in self-supervised manners, and keep under control the platform resource usages.

5. Various experiments were presented to validate the overall framework in both simulated and real-world environments. Results from comprehensive simulation studies reveal the effectiveness of multi-sensor solutions in resource-constrained active sensing missions, in terms of both the estimation quality and the level of adaptation to environmental changes.

Together with the contributions, it is also important to remark the main lessons that can be learnt from this manuscript.

1. The self-reconfiguration capabilities of most **Active Sensing** strategies enable *performance optimization* and *adaptation to dynamic and unstructured environments*. Moreover, the possibility to control the platform behavior allows the design of multi-tasking algorithms based on the exploration-exploitation trade-off: exploration is often used for mapping and knowledge discovery, while exploitation is employed in tracking and decision making processes. Finally, Active Sensing allows

to exploit the full potentialities of platforms degrees of freedom, enabling *resource optimization*. An example is provided in Chapt. 7, where processing and storage resources are preserved through data selection.

2. The **probabilistic** approaches are useful to account for realistic perception *uncertainties*, sensor nuisance and unmodeled dynamics. For instance, in Chapt. 5 a probabilistic model has been used to account for the resolution effects into the target visual detection process; in the same Chapter, a probabilistic model has been used to include interference and noise effects into the RSSI data collection in radio receivers. A further advantage of probabilistic approaches is the possibility to include some *a-priori knowledge* in the system, through Bayesian inference; in Chapt. 7 the a-priori knowledge on the parking availability map has been encrypted into the mean and covariance functions of a Gaussian Process. Finally, probabilistic active sensing strategies deal with (and learn from) *unexpected and novel situations*.

3. **Multi-Sensor** active sensing has been shown to attain faster, more accurate, and more robust estimation processes, as discussed in all applications results, from Chapt. 4 to Chapt. 7. Moreover, Chapt. 4 and 6 have proven that, within the robotics community, multi-modal sensing channels can be used to design *self-supervised learning techniques*. Finally, *resource management* can be achieved through task handover; for instance, in audio-visual tracking, the sensing channels can be alternated to preserve battery lifetime.

8.2 Considerations on the future of Smart and Active Sensing

The MS-PAS schemes developed in this Thesis represent a pointer for the advance in multi-sensor autonomous perception research, where there are still several issues that must be addressed. This conclusive Section provides some methodological and conceptual considerations regarding the field of smart and active sensing, highlighting the future technological directions and the main research challenges.

What is intelligence in autonomous perception? The recent advances in neural networks and deep learning have enabled AI scientists to tackle many complex tasks, such as natural language processing and computer vision. Nonetheless, over the past years, there has been much debate over the definition of intelligence and how to measure it in autonomous systems. As a matter of fact, solely measuring skill at any given task falls short of measuring intelligence, because unlimited priors or unlimited training data mask the system's generalization power. In fact, optimizing AI algorithms for specific tasks has entrenched the community in *narrow AI*. As a result, work in AI has drifted away from the original vision of developing "thinking machines": we are able to engineer systems that perform extremely well on specific tasks, but they still have limitations, being data-hungry and unable to make sense of situations that deviate slightly from their training data.

According to Chollet (2019), truly intelligent systems should be able to develop higher-level skills that can span across many tasks; they should also be endowed with a generalization power that allows to handle previously unseen situations. Unfortunately

though, the best that current AI systems achieve is *local generalization*, which is the ability of an agent to achieve goals within its own narrow domain. In fact, AI models often need thousands of examples to perform specific tasks and then fail as soon as they face corner cases.

Some AI researchers believe that the true direction for artificial intelligence research are methods that can scale with the availability of data, prior knowledge, and compute resources. In this regard, active sensing will play a fundamental role in machine learning and robotics, paving the way towards autonomous decision making over data collection campaigns and resource-aware autonomous perception systems. In particular, active sensing modules can be used to teach robots how to perceive the surrounding environment; once, an agent has learnt this preliminary and fundamental skill, it can gather environmental data that allow to learn and to accomplish a variety of secondary tasks. In this respect, explorative active sensing techniques will make it possible to incrementally learn new skills, giving robots the freedom to choose what to learn. To this aim, intelligent active sensing system should embed reasoning and abstraction capabilities; moreover, robots should re-purpose themselves to deal with novel tasks and to satisfy high-level specifications, with minimal human intervention in the loop.

Robustness with respect to adversarial attacks Most artificial intelligence researchers agree that one of the key concerns of machine learning is the lack of causality in machine learning systems. As long as neural networks focus on learning superficial statistical patterns in data, they will remain vulnerable to different forms of adversarial attacks (i.e., data manipulation techniques that cause trained models to behave erratically and in undesired ways). Adversarial attacks originate from either direct or indirect data poisoning effects: the former involve an attacker inserting corrupt data in the training dataset to compromise the performance of a machine learning model; the latter instead is typical of dynamic and unstructured environments, where the deployment conditions may differ from the training ones. As an example, consider a machine vision system employed in industrial settings where the temperature range is subject to sensible oscillations. In this scenario, the camera thermal noise may introduce patterns that are not present in training images, with consequent performance degradation (in terms of generalization capability). There is concern that data poisoning and adversarial attacks can become a serious security issue as deep learning becomes more prominent in physical tasks where autonomous agents physically interact with humans, such as robotics and self-driving cars. As a matter of fact, adversarial attacks reduce safety of neural networks in robotic vision, where the visual recognition systems rely on deep learning technologies to interact with the surrounding environment Melis, Demontis, Biggio, Brown, Fumera, and Roli (2017). Data poisoning is a thread also for self-supervised learning procedures, which have been shown to be a promising research trend in multi-sensor robotics. In fact, these systems exploit the correlation among different sensing cues to create the supervisory signals that guide the learning process. However, the inevitable amount of noise in input and output data may generate numerical artifacts during the training procedure, introducing novel error profiles in robot learning. For the aforementioned reasons, adversarial machine learning is today an active area of research Mehra, Kailkhura, Chen, and Hamm (2021). For instance, the authors of Lechner, Hasani, Grosu, Rus, and Henzinger (2021) show that this field needs new ways to improve adversarial robustness in deep

neural networks without reducing their accuracy and safety. In this framework, active sampling techniques can be applied. For instance, the uncertainty-aware active data selection algorithm presented in Chapt. 7 can be modified to filter out poisoned data, by exploiting detectors of out-of-distribution samples Lee, Lee, Lee, and Shin (2018) and adversarial attacks Feinman, Curtin, Shintre, and Gardner (2017); Lin, Liu, Sun, and Huang (2017b).

Computational and storage loads The advantages of multi-sensor and probabilistic approaches usually come at the cost of high computational and storage demands. In particular, the major computational bottleneck in the MS-PAS approaches is the map update procedure, which hinders their scalability to large-domain and high-resolution applications. For example, standard GPR strategies suffer from cubic time complexity in the number of accumulated measurements. In this regard, data selection has been used to mitigate data accumulation problems (see Chapt. 7); as an alternative, sparsification methods could be investigated towards enabling large-scale online applications. In general, the processing and memory load of machine learning and probabilistic inference models may limit the deployment of self-reconfigurable machines in industrial tasks, where real-time specifications are mandatory. Therefore, future research directions should focus on data selection and sensor switching algorithms to minimize the overall resource consumption and utilization.

Persistent Monitoring A core challenge in active sensing is performing autonomous data collection in unknown environments, subject to a limited amount of sensing resources. Along this line, persistent monitoring scenarios pose critical research challenges in terms of both on-board resource management (to minimize the maintenance costs) and environmental mapping (due to dynamic variations of the distribution of the monitored area). For resource management, online replanning capabilities are required. As done in Chapt. 6, the lifetime of wireless sensors can be prolonged by modeling and forecasting the probability of occurrence of the events of interest. For mapping, new representations are necessary, in order to capture temporal and spatial variations. In this perspective, sensor fusion approaches can be useful, as proven by the PAME algorithm in Chapt. 7: environmental variations are detected by cloud services and a multi-source architecture is responsible for the map adaptation.

Autonomous systems and advanced computing techniques In Chapt. 7 an internet connected vehicle estimates the parking availability along the road by accessing to cloud services and smart infrastructures (e.g., smart parking stations). This is just an example of a growing research and technological trend, involving the integration between autonomous systems and advanced computing techniques. As a matter of fact, autonomous sensing platforms can achieve better performance by combining their own local data with other advanced technologies, such as IoT architectures, big data centers, and federated learning systems on mobile devices Chen et al. (2019a). Specifically, an autonomous platform can extend its working area using information from IoT devices or cloud databases; moreover, it can incrementally update its inference models through federated learning services.

Multi-sensor fusion One of the main messages conveyed by this dissertation is that sensor fusion approaches induce several advantages. These advantages mainly involve the reduction in data ambiguity, while increasing the efficiency and the reliability of the

data acquisition system. When data fusion is applied to the routing process in large scale wireless sensor networks, only the result is forwarded; hence, the number of messages is reduced, collisions are avoided, and energy is saved. Information fusion is useful also because it provides efficient methods for automatically (or semi-automatically) transform information from different sources into a representation that provides effective support for human or automated decision making. In spite of the aforementioned benefits, multi-sensor architectures pose severe practical challenges related to device synchronization and real-time data aggregation. Furthermore, there are a number of issues arising from imperfection and diversity of the sensor technologies, outliers, spurious and conflicting data, data dimensionality and compression level Khaleghi et al. (2013). Another aspect of critical importance are data communication and processing protocols, as fusion processing can be performed in a centralized or decentralized manner. Finally, data collected at different sensor devices should be post-processed to tackle alignment and association issues: the former arise when different calibration settings are involved at different sensor nodes; data association, instead, aims to identify which sensors provide measurements from the same real-world objects (e.g., targets in multi-target applications).

Multi-robot planning A natural direction for future research is the extension of the ideas developed in this Thesis into active sensing problems with multi-robot teams. This field of study has grown considerably in the past several years, and will continue to do so, as the capabilities of autonomous platforms expand. In many monitoring missions, replacing a single-robot with a multi-robot system can offer improved spatial coverage, versatility, higher redundancy levels, and more robustness to failures. A swarm of UAVs can be deployed to monitor a very large agricultural field; alternatively, mobile robots could be used to simultaneously gather information from disjoint locations. However, such scenarios introduce additional considerations, including higher maintenance costs, as well as requirements for coordination, communication, and resource sharing.

Sensor modeling Active sensing is the problem of regulating an autonomous sensing platform guided by quality of perception criteria. In this context, the sensor model is critical, as it determines how future measurements will influence the evolution of the sensing task, and subsequently the value of the information acquired. A major practical issue emerging from the experiments in this Thesis is the development of accurate sensor models as a basis for predictive decision-making. This problem is particularly relevant for devices relying on deep learning architectures; these are gaining rapid momentum in many robotics applications, but they still lack in explainability and transparency (i.e., the result is often the outcome of a black-box inference model) Došilović, Brčić, and Hlupić (2018). Moreover, the majority of AS applications deal with harsh settings, where cluttering, environmental noise and time-varying conditions complicate the modeling of the perception process.

Software and hardware infrastructures The active sensing community would profit significantly from the availability of reliable and reproducible software and hardware to create benchmark problems and algorithms for researchers to compare new methods in a standardized manner. However, the lack of available software infrastructure for active sensing research is identified as a major obstacle hindering further research in this area. Furthermore, the design and implementation of robotics systems, which embody the basic prerequisites for active perception, is still hard, especially when it comes to deploy

multiple cooperating heterogeneous robots. In the future, the research community has to show that it is possible to build robotic systems with active perception capabilities; these include adaptivity to varied environments and tasks, the ability to deal with unexpected, novel situations, and learning from experience capabilities. In addition, active sensing platforms must be robust, reliable and safe Bajcsy et al. (2018).

Privacy concerns The power of sensor-rich robotic platforms is their capability to collect large amount of data regarding the surrounding environment. Nevertheless, this may pose severe questions on possible privacy violations. At the same time, the advances in modern computer vision techniques have encouraged researchers in signal processing and robotics community to employ cameras as principal (sometimes unique) sensing unit in tasks like scene perception, activity recognition and environmental mapping. Accordingly, new research lines are born around privacy-aware robotics and privacy-by-design signal processing techniques Osia et al. (2020).

A

Appendix

Contents

A.1 Gaussian Process Regression (GPR)

This Appendix provides a brief overview on Gaussian Processes and Gaussian Process Regression. The reader is referred to Rasmussen (2003) and Williams and Rasmussen (2006) for a more comprehensive introduction on Gaussian Processes in Machine Learning applications.

A Gaussian Process (GP) is a collection of random variables, any finite number of which have a joint Gaussian distribution Williams and Rasmussen (2006). Given the input vector $\mathbf{x} \in \mathbb{R}^p$, a GP is fully specified by its mean function $m(\mathbf{x})$ and covariance function $k(\mathbf{x}, \mathbf{x}')$, namely

$$f(\mathbf{x}) \sim \mathcal{GP}\left(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')\right), \tag{A.1}$$

where $f : \mathbb{R}^p \to \mathbb{R}$ is referred to as *latent function*. Mean and covariance (or kernel) functions incorporate prior knowledge (e.g., periodicity, smoothness) about the latent function. The mean function is typically constant (either zero or the mean of the training dataset), while the most commonly-used kernel functions are constant, linear, square exponential or Matern, as well as compositions of multiple kernels Williams and Rasmussen (2006). The hyperparameters in the mean and covariance functions are computed by fitting the train dataset $\mathcal{D}$ of cardinality n_{train}

$$\begin{aligned} &\mathcal{D} = \{(\mathbf{x}_j, y_j)\}_{j=1}^{n_{train}} = (\mathbf{X}, \mathbf{y}) \\ &\mathbf{X} = \begin{bmatrix} \mathbf{x}_1 & \dots & \mathbf{x}_{n_{train}} \end{bmatrix}^\top \in \mathbb{R}^{n_{train} \times p} \\ &\mathbf{y} = \begin{bmatrix} y_1 & \dots & y_{n_{train}} \end{bmatrix}^\top \in \mathbb{R}^{n_{train}}. \end{aligned} \tag{A.2}$$

Each training label y_j is a noisy measurement of the latent function $f(\mathbf{x})$, namely

$$y_j = f(\mathbf{x}_j) + \epsilon_j, \quad \epsilon_i \sim \mathcal{N}(\epsilon|0, \sigma^2), \tag{A.3}$$

where $\{\epsilon_j\}_{j=1}^{n_{train}}$ is the i.i.d. Gaussian noise in the training labels. Given (A.1) and (A.3), it is possible to model the observations as a GP Williams and Rasmussen (2006),

$$y(\mathbf{x}) \sim \mathcal{GP}\left(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}') + \sigma^2\mathbf{I}\right). \tag{A.4}$$

This GP is used as *prior* for non-parametric Bayesian inference of the latent function.

Consider the test inputs

$$\mathbf{X}_* = \begin{bmatrix} \mathbf{x}_{1,*} & \dots & \mathbf{x}_{n_{test},*} \end{bmatrix}^\top \in \mathbb{R}^{n_{test} \times p}. \tag{A.5}$$

By definition, any finite number of samples drawn from the GP are jointly Gaussian. Thus,

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{f}_* \end{bmatrix} = \mathcal{N}\left(\begin{bmatrix} \boldsymbol{\mu} \\ \boldsymbol{\mu}_* \end{bmatrix}, \begin{bmatrix} \mathbf{K}(\mathbf{X}, \mathbf{X}) + \sigma^2\mathbf{I} & \mathbf{K}(\mathbf{X}, \mathbf{X}_*) \\ \mathbf{K}(\mathbf{X}_*, \mathbf{X}) & \mathbf{K}(\mathbf{X}_*, \mathbf{X}_*) \end{bmatrix} \right) \tag{A.6}$$

where:

$$\mathbf{f}_* = \begin{bmatrix} f(\mathbf{x}_{1,*}) & \dots & f(\mathbf{x}_{n_{test},*}) \end{bmatrix}^\top \in \mathbb{R}^{n_{test}}; \tag{A.7}$$

the j-th row of $\boldsymbol{\mu}$ and $\boldsymbol{\mu}_*$ is $m(\mathbf{x}_j)$ and $m(\mathbf{x}_{j,*})$, respectively; the (j,k)-th entry of $\mathbf{K}(\mathbf{X}, \mathbf{X})$ and $\mathbf{K}(\mathbf{X}, \mathbf{X}_*)$ is $k(\mathbf{x}_j, \mathbf{x}_k)$ and $k(\mathbf{x}_j, \mathbf{x}_{k,*})$, respectively.

Making predictions about unobserved values $\mathbf{x}_{j,*}$ consists in drawing samples from the *predictive posterior distribution* of $\mathbf{f}_*$, given $\mathcal{D}$ and $\mathbf{X}_*$, that is

$$\begin{aligned} &\mathbf{f}_*|\mathcal{D}, \mathbf{X}_* \sim \mathcal{N}(\mathbf{f}_*|\boldsymbol{\mu}_{*|\mathcal{D}}, \boldsymbol{\Sigma}_{*|\mathcal{D}}) \\ &\boldsymbol{\mu}_{*|\mathcal{D}} = \boldsymbol{\mu}_* + \mathbf{K}(\mathbf{X}_*, \mathbf{X}) \left[\mathbf{K}(\mathbf{X}, \mathbf{X}) + \sigma^2\mathbf{I}\right]^{-1} (\mathbf{y} - \boldsymbol{\mu}) \\ &\boldsymbol{\Sigma}_{*|\mathcal{D}} = \mathbf{K}(\mathbf{X}_*, \mathbf{X}_*) - \mathbf{K}(\mathbf{X}_*, \mathbf{X}) \left[\mathbf{K}(\mathbf{X}, \mathbf{X}) + \sigma^2\mathbf{I}\right]^{-1} \mathbf{K}(\mathbf{X}, \mathbf{X}_*). \end{aligned} \tag{A.8}$$

A.2 Bayesian Optimization (BO)

In this Section the fundamentals of Bayesian Optimization are provided. The reader is referred to Frazier (2018) and Shahriari et al. (2015) for detailed overviews on Bayesian Optimization techniques.

Bayesian optimization (BO) Frazier (2018); Shahriari et al. (2015) is a procedure designed for derivative-free global optimization, particularly suited for objective functions that are expensive to evaluate. Formally, BO aims to solve the following optimization problem

$$\mathbf{x}^* = \arg\max_{\mathbf{x} \in \mathcal{X}} J(\mathbf{x}) \tag{A.9}$$

where $\mathcal{X}$ is a domain space of interest; the objective function $J : \mathcal{X} \to \mathbb{R}$ is unknown (i.e., black-box optimization), but it can be evaluated at any arbitrary query point $\mathbf{x}_* \in \mathcal{X}$. This evaluation produces a noise-corrupted (stochastic) output $y \in \mathbb{R}$.

To solve (A.9), BO adopts a sequential procedure.

1. **Cost function approximation**: at first, the objective function $J(\cdot)$ is approximated by a probabilistic model (e.g., a GP), easier to optimize and referred as *surrogate function*. The surrogate model represents a prior belief over the possible objective functions and it is sequentially refined via Bayesian posterior updating (e.g., GPR), as new data (query points) are collected. The new query point $\mathbf{x}_t$ produces y_t, noisy-corrupted version of $J(\mathbf{x}_t)$. Then, model updating is performed over $\mathcal{D}_t = \{\mathcal{D}_{t-1}, (\mathbf{x}_t, y_t)\}$, where

$$\mathcal{D}_{t-1} = \{(\mathbf{x}_j, y_j)\}_{j=1}^{t-1}. \tag{A.10}$$

2. **Data collection**: Each query point in $\mathcal{D}_t$ is chosen according to a suitable selection criterion $a(\mathbf{x}|\mathcal{D}_t)$ (called *acquisition function*)

$$\mathbf{x}_{t+1} = \arg\max_{\mathbf{x}\in\mathcal{X}} a(\mathbf{x}|\mathcal{D}_t). \tag{A.11}$$

   The acquisition function is designed over the surrogate function and quantifies the utility of a query point to produce a more informative posterior distribution; hence, the design of a suitable selection criterion is critical in many applications. In literature a wide variety of acquisition functions are proposed Shahriari et al. (2015): some of them prioritise samples that minimize the uncertainty of the model (or, equivalently, maximize the information gain); some are used to foster a balanced exploration of the latent function domain, while others are designed to focus the estimation quality on specified domain regions; finally, some acquisition functions are designed to take into account the cost associated to the sampling process (useful when query points have non-equal sampling cost).

BO in Active Sensing applications The Bayesian optimization framework is particularly useful when $J(\cdot)$ is costly, nonconvex and multimodal, and it is not possible to have access to derivatives with respect to $\mathbf{x}$. Moreover, BO has the capability to balance exploration with exploitation, and enables probabilistic decision making strategies that gain from past experience. Finally, the use of a Gaussian Process as surrogate function, allows to reach good approximation results with few measurements. For these reasons, BO has been successfully applied for several active sensing tasks, including path planning for target search Meera et al. (2019), cooperative multi-robot control Ghassemi and Chowdhury (2020), environment exploration Bai et al. (2016), and source localization Carpin et al. (2015).

A.3 Recursive Bayesian Estimation (RBE)

This Section serves as brief tutorial on Recursive Bayesian Estimation and Particle Filtering in target tracking applications. The objective is to estimate the probability density function (PDF) of a given variable (in this case, the target position $\mathbf{p}_t$), recursively over time and using incoming measurements $\mathbf{z}_{1:t}$ (in this Section a generic sensor is

considered; hence, no specific assumptions are made on the observation model). The reader is referred Smith (2013) for detailed overviews on Sequential Monte Carlo methods and Bayesian statistical inference.

Given the observations $\mathbf{z}_{1:t}$, RBE provides a two-stage procedure to recursively update the target belief state, namely the posterior distribution $p(\mathbf{p}_t|\mathbf{z}_{1:t})$. The prediction stage involves using the target process model $p(\mathbf{p}_t|\mathbf{p}_{t-1})$ (see Chapt. 5) to obtain the prior of the target position via the Chapman-Kolmogorov equation Radmard and Croft (2017)

$$p(\mathbf{p}_t|\mathbf{z}_{1:t-1}) = \int p(\mathbf{p}_t|\mathbf{p}_{t-1})p(\mathbf{p}_{t-1}|\mathbf{z}_{1:t-1})d\mathbf{p}_{t-1}. \tag{A.12}$$

As a new observation $\mathbf{z}_t$ becomes available, Bayes rule Smith (2013) updates the target belief state

$$p(\mathbf{p}_t|\mathbf{z}_{1:t}) = \frac{p(\mathbf{p}_t|\mathbf{z}_{1:t-1})p(\mathbf{z}_t|\mathbf{p}_t)}{\int p(\mathbf{z}_t|\mathbf{p}_{t-1})p(\mathbf{p}_{t-1}|\mathbf{z}_{1:t-1})d\mathbf{p}_{t-1}}, \tag{A.13}$$

with $p(\mathbf{p}_0)$ as initial target state belief.

Particle filtering Smith (2013) is a Monte Carlo approximation of the density $p(\mathbf{p}_t|\mathbf{z}_{1:t})$ with a sum of N_s Dirac functions centered in the particles $\{\mathbf{p}_t^{(i)}\}_{i=1}^{N_s}$, that is

$$p(\mathbf{p}_t|\mathbf{z}_{1:t}) \approx \sum_{i=1}^{N_s} w_t^{(i)} \delta\left(\mathbf{p}_t - \mathbf{p}_t^{(i)}\right), \tag{A.14}$$

where $w_t^{(i)}$ is the weight of particle $\mathbf{p}_t^{(i)}$ and it holds

$$\mathbf{p}_t^{(i)} = f\left(\mathbf{p}_{t-1}^{(i)}, \boldsymbol{\eta}_{t-1}\right) \tag{A.15a}$$

$$w_t^{(i)} \propto w_{t-1}^{(i)} p\left(\mathbf{z}_t|\mathbf{p}_t^{(i)}\right) \tag{A.15b}$$

for $i = 1, \ldots, N_s$. The i-th weight is therefore proportional to the likelihood function evaluated on the i-th particle, namely $p(\mathbf{z}_t|\mathbf{p}_t^{(i)})$. Under this scheme, a few particles will concentrate most of the weight over time, leading to a phenomenon called degeneracy. To alleviate this problem, systematic resampling is adopted, which involves resampling from the particle set when the effective number of particles

$$N_{eff} = \frac{1}{\sum_{i=1}^{N_s} w_i^2} \tag{A.16}$$

is low (i.e., $N_{eff} < N_s/2$). Particles with higher weights are more likely to be sampled, but the procedure preserves some low weight particles as well Smith (2013).

MAP and MMSE estimators Once the discrete approximation of $p(\mathbf{p}_t|\mathbf{z}_{1:t})$ has been reconstructed, it is possible to compute a punctual estimate of the variable $\mathbf{p}_t$.

- Maximum A-Posteriori (MAP) estimate: mode of the posterior distribution, namely

$$\widehat{\mathbf{p}}_t = \mathbf{p}_t^{(i^*)};\ i^* = \arg\max_{i \in [1, N_s]} w_t^{(i)}. \tag{A.17}$$

- Minimum Mean Square Error (MMSE) estimate: Monte Carlo approximation of the expectation, namely

$$\widehat{\mathbf{p}}_t = \frac{1}{N_s} \sum_{i=1}^{N_s} \omega_t^{(i)} \mathbf{p}_t^{(i)} \approx \mathbb{E}[\mathbf{p}_t | \mathbf{z}_{1:t}]. \tag{A.18}$$

“ It’s this trust (in AI) that bothers me. Not because of how often AI gets it wrong. It’s how badly it gets it wrong when it makes a mistake that has me worried. These systems do not fail gracefully.”

Peter Haas

Bibliography

Achroufene A., Amirat Y., and Chibani A. Rss-based indoor localization using belief function theory. *IEEE Transactions on Automation Science and Engineering*, 16 (3):1163–1180, 2018.

Aghajan H. and Cavallaro A. *Multi-camera networks: principles and applications*. Academic press, 2009.

Aghdasi H. S., Nasseri S., and Abbaspour M. Energy efficient camera node activation control in multi-tier wireless visual sensor networks. *Wireless networks*, 19 (5):725–740, 2013.

Alemdar H., Durmus Y., and Ersoy C. Wireless healthcare monitoring with rfid-enhanced video sensor networks. *Int. Journal of Distributed Sensor Networks*, 6(1): 473037, 2010.

Aoki E. H., Bagchi A., Mandal P., and Boers Y. A theoretical look at information-driven sensor management criteria. In *14th International Conference on Information Fusion*, pages 1–8. IEEE, 2011a.

Aoki E., Bagchi A., Mandal P., and Boers Y. On the "near-universal proxy" argument for theoretical justification of information-driven sensor management. In *2011 IEEE Statistical Signal Processing Workshop (SSP)*, pages 245–248. IEEE, 2011b.

Aranda M., López-Nicolás G., Sagüés C., and Mezouar Y. Formation control of mobile robots using multiple aerial cameras. *IEEE Transactions on Robotics*, 31(4): 1064–1071, 2015.

Arndt K., Hazara M., Ghadirzadeh A., and Kyrki V. Meta reinforcement learning for sim-to-real domain adaptation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2725–2731. IEEE, 2020.

Arulampalam M. S., Maskell S., Gordon N., and Clapp T. A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking. *IEEE Transactions on signal processing*, 50(2):174–188, 2002.

Aryandoust A., van Vliet O., and Patt A. City-scale car traffic and parking density maps from uber movement travel time data. *Scientific data*, 6(1):1–18, 2019.

Atanasov N. A., Le Ny J., and Pappas G. J. Distributed algorithms for stochastic source seeking with mobile robot networks. *Journal of Dynamic Systems, Measurement, and Control*, 137(3), 2015.

Aughenbaugh J. M. and LaCour B. R. Sensor management for particle filter tracking. *IEEE Transactions on Aerospace and Electronic Systems*, 47(1):503–523, 2011.

Avula S., Chadwick G., Arguello J., and Capra R. Searchbots: User engagement with chatbots during collaborative search. In *Proceedings of the 2018 conference on human information interaction & retrieval*, pages 52–61, 2018.

Bai S., Wang J., Chen F., and Englot B. Information-theoretic exploration with bayesian optimization. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1816–1822. IEEE, 2016.

Bajcsy R. Active perception. *Proceedings of the IEEE*, 76(8):966–1005, 1988.

Bajcsy R., Aloimonos Y., and Tsotsos J. K. Revisiting active perception. *Autonomous Robots*, 42(2):177–196, 2018.

Bajsy R. Active perception vs. passive perception. In *Proc. IEEE Workshop on Computer Vision*, 1985.

Balsamo D., Merrett G. V., Zaghari B., Wei Y., Ramchurn S., Stein S., Weddell A. S., and Beeby S. Wearable and autonomous computing for future smart cities: Open challenges. In *2017 25th International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, pages 1–5. IEEE, 2017.

Baltrušaitis T., Ahuja C., and Morency L.-P. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41(2):423–443, 2018.

Bemporad A., Morari M., and Ricker N. L. Model predictive control. *MathWorks.–2010.–205 p*, 2005.

Bernardini S., Fox M., and Long D. Combining temporal planning with probabilistic reasoning for autonomous surveillance missions. *Autonomous Robots*, 41(1):181–203, 2017.

Bischoff J., Maciejewski M., Schlenther T., and Nagel K. Autonomous vehicles and their impact on parking search. *IEEE Intelligent Transportation Systems Magazine*, 11(4):19–27, 2018.

Blondin M. J., Sáez J. S., and Pardalos P. M. Control engineering from classical to intelligent control theory—an overview. *Computational Intelligence and Optimization Methods for Control Engineering*, pages 1–30, 2019.

Boers Y., Driessen H., Bagchi A., and Mandal P. Particle filter based entropy. In *2010 13th International Conference on Information Fusion*, pages 1–8. IEEE, 2010.

Bonatti R., Ho C., Wang W., Choudhury S., and Scherer S. Towards a robust aerial cinematography platform: Localizing and tracking moving targets in unstructured environments. *arXiv preprint arXiv:1904.02319*, 2019.

Bopardikar S. D., Bullo F., and Hespanha J. P. Cooperative pursuit with sensing limitations. In *2007 American Control Conference*, pages 5394–5399. IEEE, 2007.

Bordel B., Alcarria R., Robles T., and Martín D. Cyber–physical systems: Extending pervasive sensing from control theory to the internet of things. *Pervasive and mobile computing*, 40:156–184, 2017.

Bourgault F., Furukawa T., and Durrant-Whyte H. F. Coordinated decentralized search for a lost target in a bayesian world. In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, volume 1, pages 48–53. IEEE, 2003.

Bourne J. R., Pardyjak E. R., and Leang K. K. Coordinated bayesian-based bioinspired plume source term estimation and source seeking for mobile robots. *IEEE Transactions on Robotics*, 35(4):967–986, 2019.

Bregonzio M., Taj M., and Cavallaro A. Multi-modal particle filtering tracking using appearance, motion and audio likelihoods. In *2007 IEEE International Conference on Image Processing*, volume 5, pages V–33. IEEE, 2007.

Brendel A. and Kellermann W. Learning-based acoustic source-microphone distance estimation using the coherent-to-diffuse power ratio. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 61–65. IEEE, 2018.

Brochu E., Cora V. M., and De Freitas N. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.

Brutti A. and Cavallaro A. Unsupervised cross-modal deep-model adaptation for audio-visual re-identification with wearable cameras. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 438–445, 2017.

Buşoniu L., Babuška R., and De Schutter B. Multi-agent reinforcement learning: An overview. In *Innovations in multi-agent systems and applications-1*, pages 183–221. Springer, 2010.

Calli B., Caarls W., Wisse M., and Jonker P. P. Active vision via extremum seeking for robots in unstructured environments: Applications in object recognition and manipulation. *IEEE Transactions on Automation Science and Engineering*, 15(4): 1810–1822, 2018.

Cantón Paterna V., Calveras Augé A., Paradells Aspas J., and Pérez Bullones M. A bluetooth low energy indoor positioning system with channel diversity, weighted trilateration and Kalman filtering. *Sensors*, 17(12):2927, 2017a.

Cantón Paterna V., Calveras Auge A., Paradells Aspas J., and Perez Bullones M. A. A bluetooth low energy indoor positioning system with channel diversity, weighted trilateration and kalman filtering. *Sensors*, 17(12):2927, 2017b.

Cao J. and Menendez M. System dynamics of urban traffic based on its parking-related-states. *Transportation Research Part B: Methodological*, 81:718–736, 2015.

Cao Z., Tan M., Li L., Gu N., and Wang S. Cooperative hunting by distributed mobile robots based on local interaction. *IEEE Transactions on Robotics*, 22(2): 402–406, 2006.

Carpin M., Rosati S., Khan M. E., and Rimoldi B. Uavs using bayesian optimization to locate wifi devices. *arXiv preprint:1510.03592*, 2015.

Carron A., Todescato M., Carli R., Schenato L., and Pillonetto G. Multi-agents adaptive estimation and coverage control using gaussian regression. In *2015 European Control Conference (ECC)*, pages 2490–2495. IEEE, 2015.

Cenedese A., Ortolan G., and Bertinato M. Low-density wireless sensor networks for localization and tracking in critical environments. *IEEE Transactions on Vehicular Technology*, 59(6):2951–2962, 2010.

Cenedese A. and Varotto L. A distributed approach to 3d reconstruction in marker motion capture systems. In *Proceedings of the 13th International Conference on Distributed Smart Cameras*, ICDSC 2019, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450371896. URL https://doi.org/10.1145/3349801.3349818.

Chang Y.-Y., Mazzon R., and Cavallaro A. Real-time quality assessment of videos from body-worn cameras. In *2018 26th European Signal Processing Conference (EUSIPCO)*, pages 2160–2164, 2018.

Chen H., Zhao X., and Tan M. A novel pan-tilt camera control approach for visual tracking. In *Proceeding of the 11th World Congress on Intelligent Control and Automation*, pages 2860–2865. IEEE, 2014.

Chen J., Abbod M., and Shieh J.-S. Integrations between autonomous systems and modern computing techniques: a mini review. *Sensors*, 19(18):3897, 2019a.

Chen J., Zhang X., Xin B., and Fang H. Coordination between unmanned aerial and ground vehicles: A taxonomy and optimization perspective. *IEEE transactions on cybernetics*, 46(4):959–972, 2015.

Chen S., Li Y., and Kwok N. M. Active vision in robotic systems: A survey of recent developments. *The International Journal of Robotics Research*, 30(11):1343–1377, 2011.

Chen Y., Xu S., Ren Z., and Chirarattananon P. Collision resilient insect-scale soft-actuated aerial robots with high agility. *IEEE Transactions on Robotics*, 2021.

Chen Z. and others . Bayesian filtering: From kalman filters to particle filters, and beyond. *Statistics*, 182(1):1–69, 2003.

Chen Z., Zou H., Yang J., Jiang H., and Xie L. Wifi fingerprinting indoor localization using local feature-based deep lstm. *IEEE Systems Journal*, 2019b.

Chibani A., Amirat Y., Mohammed S., Matson E., Hagita N., and Barreto M. Ubiquitous robotics: Recent challenges and future trends. *Robotics and Autonomous Systems*, 61(11):1162–1172, 2013.

Choi J., Choi Y.-S., and Talwar S. Unsupervised learning techniques for trilateration: From theory to android app implementation. *IEEE Access*, 7:134525–134538, 2019.

Chollet F. On the measure of intelligence. *arXiv preprint arXiv:1911.01547*, 2019.

Chung C. F. and Furukawa T. Coordinated pursuer control using particle filters for autonomous search-and-capture. *Robotics and Autonomous Systems*, 57(6-7):700–711, 2009.

Chung S.-J., Paranjape A. A., Dames P., Shen S., and Kumar V. A survey on aerial swarm robotics. *IEEE Transactions on Robotics*, 34(4):837–855, 2018.

Chung T. H. and Burdick J. W. Analysis of search decision making using probabilistic search strategies. *IEEE Transactions on Robotics*, 28(1):132–144, 2011.

Chung T. H., Hollinger G. A., and Isler V. Search and pursuit-evasion in mobile robotics. *Autonomous robots*, 31(4):299, 2011.

Chung T. H., Kress M., and Royset J. O. Probabilistic search optimization and mission assignment for heterogeneous autonomous agents. In *2009 IEEE International Conference on Robotics and Automation*, pages 939–945. IEEE, 2009.

Čičić M., Barreau M., and Johansson K. H. Numerical investigation of traffic state reconstruction and control using connected automated vehicles. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–6. IEEE, 2020.

Cliff O. M., Saunders D. L., and Fitch R. Robotic ecology: Tracking small dynamic animals with an autonomous aerial vehicle. *Science Robotics*, 3(23):eaat8409, 2018.

Cobos M., Antonacci F., Alexandridis A., Mouchtaris A., and Lee B. A survey of sound source localization methods in wireless acoustic sensor networks. *Wireless Communications and Mobile Computing*, 2017, 2017.

Coluccia A. and Ricciato F. Rss-based localization via bayesian ranging and iterative least squares positioning. *IEEE Communications letters*, 18(5):873–876, 2014.

Cover T. M. *Elements of information theory*. John Wiley & Sons, 1999.

Cuevas-Velasquez H., Li N., Tylecek R., Saval-Calvo M., and Fisher R. B. Hybrid multi-camera visual servoing to moving target. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1132–1137. IEEE, 2018.

Czarnetzki S., Kerner S., and Kruse M. Real-time active vision by entropy minimization applied to localization. In *Robot Soccer World Cup*, pages 266–277. Springer, 2010.

Das S., Roy S., and Sambasivan R. Fast gaussian process regression for big data. *Big data research*, 14:12–26, 2018.

Davison A. J., Mayol W. W., and Murray D. W. Real-time localization and mapping with wearable active vision. In *The Second IEEE and ACM International Symposium on Mixed and Augmented Reality, 2003. Proceedings.*, pages 18–27. IEEE, 2003.

Davoodi M. and Velni J. M. Heterogeneity-aware graph partitioning for distributed deployment of multiagent systems. *IEEE Transactions on Cybernetics*, 2020.

de San Bernabé A., Martinez-de Dios J., and Ollero A. Efficient integration of rssi for tracking using wireless camera networks. *Information Fusion*, 36:296–312, 2017.

Deak G., Curran K., and Condell J. A survey of active and passive indoor localisation systems. *Computer Communications*, 35(16):1939–1954, 2012.

Dehghan S. M. M., Farmani M., and Moradi H. Aerial localization of an rf source in nlos condition. In *2012 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pages 1146–1151. IEEE, 2012.

Denzler J., Zobel M., and Niemann H. Information theoretic focal length selection for real-time active 3-d object tracking. In *null*, page 400. IEEE, 2003.

Deyle T., Nguyen H., Reynolds M., and Kemp C. Rfid-guided robots for pervasive automation. *IEEE Pervasive Computing*, 9(2):37–45, 2010.

DiBiase J. H. *A high-accuracy, low-latency technique for talker localization in reverberant environments using microphone arrays*. Brown University Providence, RI, 2000.

Ding C., Song B., Morye A., Farrell J. A., and Roy-Chowdhury A. K. Collaborative sensing in a distributed ptz camera network. *IEEE Transactions on Image Processing*, 21(7):3282–3295, 2012.

Dogancay K. Uav path planning for passive emitter localization. *IEEE Transactions on Aerospace and Electronic systems*, 48(2):1150–1166, 2012.

Doğançay K., Hmam H., Drake S. P., and Finn A. Centralized path planning for unmanned aerial vehicles with a heterogeneous mix of sensors. In *2009 International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP)*, pages 91–96. IEEE, 2009.

Dore A., Cattoni A. F., and Regazzoni C. S. A particle filter based fusion framework for video-radio tracking in smart spaces. In *2007 IEEE Conference on Advanced Video and Signal Based Surveillance*, pages 99–104. IEEE, 2007a.

Dore A., Cattoni A. F., and Regazzoni C. S. A particle filter based fusion framework for video-radio tracking in smart spaces. In *2007 IEEE Conference on Advanced Video and Signal Based Surveillance*, pages 99–104. IEEE, 2007b.

Došilović F. K., Brčić M., and Hlupić N. Explainable artificial intelligence: A survey. In *2018 41st International convention on information and communication technology, electronics and microelectronics (MIPRO)*, pages 0210–0215. IEEE, 2018.

Du Y.-C., Zhang M.-X., Ling H.-F., and Zheng Y.-J. Evolutionary planning of multi-uav search for missing tourists. *IEEE Access*, 7:73480–73492, 2019.

Engmann F., Katsriku F. A., Abdulai J.-D., Adu-Manu K. S., and Banaseka F. K. Prolonging the lifetime of wireless sensor networks: A review of current techniques. *Wireless Communications and Mobile Computing*, 2018, 2018.

Esterle L., Lewis P. R., McBride R., and Yao X. The future of camera networks: Staying smart in a chaotic world. In *Proc. of the 11th International Conf. on Distributed Smart Cameras*, pages 163–168, 2017.

Fabris M. *Distributed Optimization Strategies for Mobile Multi-Agent Systems*. PhD thesis, University of Padova, 2021.

Fabris M. and Cenedese A. Distributed strategies for dynamic coverage with limited sensing capabilities. In *2019 27th Mediterranean Conference on Control and Automation (MED)*, pages 209–214. IEEE, 2019.

Farmani M., Pedersen M. S., Tan Z.-H., and Jensen J. On the influence of microphone array geometry on hrtf-based sound source localization. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 439–443. IEEE, 2015.

Feinman R., Curtin R. R., Shintre S., and Gardner A. B. Detecting adversarial samples from artifacts. *arXiv preprint arXiv:1703.00410*, 2017.

Ferrari S., Fierro R., and Tolic D. A geometric optimization approach to tracking maneuvering targets using a heterogeneous mobile sensor network. In *Proceedings of the 48h IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*, pages 1080–1087. IEEE, 2009.

Ferrarotti L., Breschi V., and Bemporad A. The benefits of sharing: a cloud-aided performance-driven framework to learn optimal feedback policies. In *Learning for Dynamics and Control*, pages 87–98. PMLR, 2021.

Fink J. and Kumar V. Online methods for radio signal mapping with mobile robots. In *2010 IEEE International Conference on Robotics and Automation*, pages 1940–1945. IEEE, 2010.

Frazier P. I. A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.

Frew E. W. and Elston J. Target assignment for integrated search and tracking by active robot networks. In *2008 IEEE International Conference on Robotics and Automation*, pages 2354–2359. IEEE, 2008.

Fricke G. M., Hecker J. P., Griego A. D., Tran L. T., and Moses M. E. A distributed deterministic spiral search algorithm for swarms. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4430–4436. IEEE, 2016.

Friedewald M. and Raabe O. Ubiquitous computing: An overview of technology impacts. *Telematics and Informatics*, 28(2):55–65, 2011.

Furukawa T., Bourgault F., Lavis B., and Durrant-Whyte H. F. Recursive bayesian search-and-tracking using coordinated uavs for lost targets. In *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.*, pages 2521–2526. IEEE, 2006.

Furukawa T., Mak L. C., Durrant-Whyte H., and Madhavan R. Autonomous bayesian search and tracking, and its experimental validation. *Advanced Robotics*, 26 (5-6):461–485, 2012.

Gage D. W. Randomized search strategies with imperfect sensors. In *Mobile Robots VIII*, volume 2058, pages 270–279. International Society for Optics and Photonics, 1994.

Gao W., Wang W., Zhu H., Huang G., Wu D., and Du Z. Robust radiation sources localization based on the peak suppressed particle filter for mixed multi-modal environments. *Sensors*, 18(11):3784, 2018.

Ge X., Yang F., and Han Q.-L. Distributed networked control systems: A brief overview. *Information Sciences*, 380:117–131, 2017.

Ghassemi P. and Chowdhury S. An extended bayesian optimization approach to decentralized swarm robotic search. *Journal of Computing and Information Science in Engineering*, 20(5), 2020.

Ghods R., Durkin W. J., and Schneider J. Multi-agent active search using realistic depth-aware noise model. *arXiv preprint arXiv:2011.04825*, 2020.

Goldsmith A. *Wireless Communications*. Cambridge University Press, 2005.

Gómez-Tornero J. L., Cañete-Rebenaque D., López-Pastor J. A., and Martínez-Sala A. S. Hybrid analog-digital processing system for amplitude-monopulse rssi-based mimo wifi direction-of-arrival estimation. *IEEE Journal of Selected Topics in Signal Processing*, 12(3):529–540, 2018.

Gopalswamy S. and Rathinam S. Infrastructure enabled autonomy: A distributed intelligence architecture for autonomous vehicles. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 986–992. IEEE, 2018.

Gordon N., Ristic B., and Arulampalam S. Beyond the kalman filter: Particle filters for tracking applications. *Artech House, London*, 830(5):1–4, 2004.

Grünwald P. D., Dawid A. P., and others . Game theory, maximum entropy, minimum discrepancy and robust bayesian decision theory. *the Annals of Statistics*, 32 (4):1367–1433, 2004.

Gu Y. and Ren F. Energy-efficient indoor localization of smart hand-held devices using bluetooth. *IEEE Access*, 3:1450–1461, 2015.

Guérin F., Guinand F., Brethé J.-F., Pelvillain H., and others . Towards an autonomous warehouse inventory scheme. In *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1–8. IEEE, 2016.

Gupta A. K., Smith K. G., and Shalley C. E. The interplay between exploration and exploitation. *Academy of management journal*, 49(4):693–706, 2006.

Häne C., Heng L., Lee G. H., Fraundorfer F., Furgale P., Sattler T., and Pollefeys M. 3d visual perception for self-driving cars using a multi-camera system: Calibration, mapping, localization, and obstacle detection. *Image and Vision Computing*, 68:14–27, 2017.

Hasanzade M., Herekoğlu Ö., Yeniçeri R., Koyuncu E., and İnalhan G. RF source localization using unmanned aerial vehicle with particle filter. In *2018 9th International Conference on Mechanical and Aerospace Engineering (ICMAE)*, pages 284–289. IEEE, 2018.

Haubner T., Schmidt A., and Kellermann W. Active acoustic source tracking exploiting particle filtering and monte carlo tree search. In *27th European Signal Processing Conference*, pages 1–5. IEEE, 2019a.

Haubner T., Schmidt A., and Kellermann W. Active acoustic source tracking exploiting particle filtering and Monte Carlo tree search. In *2019 27th European Signal Processing Conference (EUSIPCO)*, pages 1–5. IEEE, 2019b.

Haugen J. and Imsland L. Monitoring moving objects using aerial mobile sensors. *IEEE Transactions on Control Systems Technology*, 24(2):475–486, 2015.

Heng L., Choi B., Cui Z., Geppert M., Hu S., Kuan B., Liu P., Nguyen R., Yeo Y. C., Geiger A., and others . Project autovision: Localization and 3d scene perception for an autonomous vehicle with a multi-camera system. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 4695–4702. IEEE, 2019.

Hernandez-Leal P., Kartal B., and Taylor M. E. A survey and critique of multiagent deep reinforcement learning. *Autonomous Agents and Multi-Agent Systems*, 33 (6):750–797, 2019.

Hero A. O., Kreucher C. M., and Blatt D. Information theoretic approaches to sensor management. In *Foundations and applications of sensor management*, pages 33–57. Springer, 2008.

Herrera M., Pérez-Hernández M., Kumar Parlikad A., and Izquierdo J. Multiagent systems and complex networks: Review and applications in systems engineering. *Processes*, 8(3):312, 2020.

Hoang M. T., Yuen B., Dong X., Lu T., Westendorp R., and Reddy K. Recurrent neural networks for accurate rssi indoor localization. *IEEE Internet of Things Journal*, 6(6):10639–10651, 2019.

Hoffmann G. M. and Tomlin C. J. Mobile sensor network control using mutual information methods and particle filters. *IEEE Transactions on Automatic Control*, 55 (1):32–47, 2009.

Hofmann P., Samp C., and Urbach N. Robotic process automation. *Electronic Markets*, 30(1):99–106, 2020.

Hollinger G. and Singh S. Proofs and experiments in scalable, near-optimal search by multiple robots. *Proceedings of Robotics: Science and Systems IV, Zurich, Switzerland*, 1, 2008.

Hollinger G., Singh S., Djugash J., and Kehagias A. Efficient multi-robot search for a moving target. *The International Journal of Robotics Research*, 28(2):201–219, 2009.

Hollinger G. A., Yerramalli S., Singh S., Mitra U., and Sukhatme G. S. Distributed data fusion for multirobot search. *IEEE Transactions on Robotics*, 31(1): 55–66, 2014.

Hornik K. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):251–257, 1991.

Hornung A., Zeng B., and Kobbelt L. Image selection for improved multi-view stereo. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008.

How J. P., Fraser C., Kulling K. C., Bertuccelli L. F., Toupet O., Brunet L., Bachrach A., and Roy N. Increasing autonomy of uavs. *IEEE Robotics & Automation Magazine*, 16(2):43–51, 2009.

Hu J., Xie L., Xu J., and Xu Z. Multi-agent cooperative target search. *Sensors*, 14 (6):9408–9428, 2014.

Hu X., Hu H., Verma S., and Zhang Z.-L. Physics-guided deep neural networks for powerflow analysis. *arXiv preprint arXiv:2002.00097*, 2020.

Hue C., Le Cadre J.-P., and Pérez P. Tracking multiple objects with particle filtering. *IEEE transactions on aerospace and electronic systems*, 38(3):791–812, 2002.

Hutchinson M., Liu C., and Chen W.-H. Source term estimation of a hazardous airborne release using an unmanned aerial vehicle. *Journal of Field Robotics*, 36(4): 797–817, 2019a.

Hutchinson M., Liu C., Thomas P., and Chen W.-H. Unmanned aerial vehicle-based hazardous materials response: Information-theoretic hazardous source search and reconstruction. *IEEE Robotics & Automation Magazine*, 2019b.

Hutchinson M., Oh H., and Chen W.-H. A review of source term estimation methods for atmospheric dispersion events using static or mobile sensors. *Information Fusion*, 36:130–148, 2017.

Iashin V., Palermo F., Solak G., and Coppola C. Top-1 corsmal challenge 2020 submission: Filling mass estimation using multi-modal observations of human-robot handovers. *arXiv preprint arXiv:2012.01311*, 2020.

Ishikawa R., Nagao Y., Hachiuma R., and Saito H. Audio-visual hybrid approach for filling mass estimation. In *25th International Conference on Pattern Recognition Workshops, ICPR 2020*, pages 437–450. Springer Science and Business Media Deutschland GmbH, 2021.

Isler V., Kannan S., and Khanna S. Randomized pursuit-evasion in a polygonal environment. *IEEE Transactions on Robotics*, 21(5):875–884, 2005.

Jafari M. and Xu H. Intelligent control for unmanned aerial systems with system uncertainties and disturbances using artificial neural network. *Drones*, 2(3):30, 2018.

Jaimes A. and Sebe N. Multimodal human–computer interaction: A survey. *Computer vision and image understanding*, 108(1-2):116–134, 2007.

Jais M., Ehkan P., Ahmad R., Ismail I., Sabapathy T., and Jusoh M. Review of angle of arrival (aoa) estimations through received signal strength indication (rssi) for wireless sensors network (wsn). In *2015 International Conference on Computer, Communications, and Control Technology (I4CT)*, pages 354–359. IEEE, 2015.

Jaradat M., Jarrah M., Bousselham A., Jararweh Y., and Al-Ayyoub M. The internet of energy: smart sensor networks and big data management for smart grid. *Procedia Computer Science*, 56:592–597, 2015.

Jelicic V., Magno M., Brunelli D., Bilas V., and Benini L. Benefits of wake-up radio in energy-efficient multimodal surveillance wireless sensor network. *IEEE Sensors Journal*, 14(9):3210–3220, 2014.

Jeon B. F., Shim D., and Kim H. J. Detection-aware trajectory generation for a drone cinematographer. *arXiv preprint arXiv:2009.01565*, 2020.

Ji S.-Q., Huang M.-B., and Huang H.-P. Robot intelligent grasp of unknown objects based on multi-sensor information. *Sensors*, 19(7):1595, 2019.

Jiang H., Cai C., Ma X., Yang Y., and Liu J. Smart home based on wifi sensing: A survey. *IEEE Access*, 6:13317–13325, 2018.

Joy J. and Gerla M. Internet of vehicles and autonomous connected car-privacy and security issues. In *2017 26th International Conference on Computer Communication and Networks (ICCCN)*, pages 1–9. IEEE, 2017.

Julian B. J., Angermann M., Schwager M., and Rus D. Distributed robotic sensor networks: An information-theoretic approach. *The International Journal of Robotics Research*, 31(10):1134–1154, 2012.

Kandris D., Nakas C., Vomvas D., and Koulouras G. Applications of wireless sensor networks: an up-to-date survey. *Applied System Innovation*, 3(1):14, 2020.

Katenka N., Levina E., and Michailidis G. Tracking multiple targets using binary decisions from wireless sensor networks. *Journal of the American Statistical Association*, 108(502):398–410, 2013.

Katsilieris F., Boers Y., and Driessen H. Optimal search: a practical interpretation of information-driven sensor management. In *2012 15th International Conference on Information Fusion*, pages 439–446. IEEE, 2012.

Kemna S., Rogers J. G., Nieto-Granda C., Young S., and Sukhatme G. S. Multi-robot coordination through dynamic voronoi partitioning for informative adaptive sampling in communication-constrained environments. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2124–2130. IEEE, 2017.

Kermorgant O. and Chaumette F. Multi-sensor data fusion in sensor-based control: application to multi-camera visual servoing. In *2011 IEEE International Conference on Robotics and Automation*, pages 4518–4523. IEEE, 2011.

Kersting K., Plagemann C., Pfaff P., and Burgard W. Most likely heteroscedastic gaussian process regression. In *Proc. of the 24th Int. Conf. on Machine Learning*, pages 393–400, 2007.

Khalastchi E. and Kalech M. Fault detection and diagnosis in multi-robot systems: a survey. *Sensors*, 19(18):4019, 2019.

Khaleghi B., Khamis A., Karray F. O., and Razavi S. N. Multisensor data fusion: A review of the state-of-the-art. *Information fusion*, 14(1):28–44, 2013.

Khan A., Rinner B., and Cavallaro A. Cooperative robots to observe moving targets. *IEEE transactions on cybernetics*, 48(1):187–198, 2016.

Khan A., Yanmaz E., and Rinner B. Information merging in multi-uav cooperative search. In *2014 ieee international conference on robotics and automation (icra)*, pages 3122–3129. IEEE, 2014.

Kim J.-H., Kim Y.-D., and Lee K.-H. The third generation of robotics: Ubiquitous robot. In *Proc of the 2nd Int Conf on Autonomous Robots and Agents*. Citeseer, 2004.

Kim J., Cho Y., and Kim A. Exposure control using bayesian optimization based on entropy weighted image gradient. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 857–864. IEEE, 2018a.

Kim S. J., Jeong Y., Park S., Ryu K., and Oh G. A survey of drone use for entertainment and avr (augmented and virtual reality). In *Augmented Reality and Virtual Reality*, pages 339–352. Springer, 2018b.

Kindt P. H., Saur M., Balszun M., and Chakraborty S. Neighbor discovery latency in ble-like protocols. *IEEE Transactions on Mobile Computing*, 17(3):617–631, 2017.

Kleisouris K., Chen Y., Yang J., and Martin R. P. The impact of using multiple antennas on wireless localization. In *2008 5th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks*, pages 55–63. IEEE, 2008.

Koch W. On exploiting 'negative'sensor evidence for target tracking and sensor data fusion. *Information Fusion*, 8(1):28–39, 2007.

Konings D., Faulkner N., Alam F., Noble F., and Lai E. Do rssi values reliably map to rss in a localization system? In *2017 2nd Workshop on Recent Trends in Telecommunications Research (RTTR)*, pages 1–5. IEEE, 2017.

Koohifar F., Guvenc I., and Sichitiu M. L. Autonomous tracking of intermittent RF source using a uav swarm. *IEEE Access*, 6:15884–15897, 2018.

Koohifar F., Kumbhar A., and Guvenc I. Receding horizon multi-uav cooperative tracking of moving rf source. *IEEE Communications Letters*, 21(6):1433–1436, 2016.

Koopman B. O. Search and screening, operations evaluation group report 56. *Center for Naval Analysis, Alexandria, Virginia*, 1946.

Krasniqi X. and Hajrizi E. Use of iot technology to drive the automotive industry from connected to full autonomous vehicles. *IFAC-PapersOnLine*, 49(29):269–274, 2016.

Krause A. and Guestrin C. Submodularity and its applications in optimized information gathering. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(4):1–20, 2011.

Kreucher C., Hero A. O., and Kastella K. A comparison of task driven and information driven sensor management for target tracking. In *Proceedings of the 44th IEEE Conference on Decision and Control*, pages 4004–4009. IEEE, 2005a.

Kreucher C., Kastella K., and Hero Iii A. O. Sensor management using an active sensing approach. *Signal Processing*, 85(3):607–624, 2005b.

Kreucher C., Kastella K., and Hero Iii A. O. Sensor management using an active sensing approach. *Signal Processing*, 85(3):607–624, 2005c.

Kumar G. A., Lee J. H., Hwang J., Park J., Youn S. H., and Kwon S. Lidar and camera fusion approach for object distance estimation in self-driving vehicles. *Symmetry*, 12(2):324, 2020.

Kyung C.-M. and others . *Theory and applications of smart cameras*. Springer, 2016.

Lanillos P., Besada-Portas E., Pajares G., and Ruz J. J. Minimum time search for lost targets using cross entropy optimization. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 602–609. IEEE, 2012.

Lanillos P., Gan S. K., Besada-Portas E., Pajares G., and Sukkarieh S. Multi-uav target search using decentralized gradient-based negotiation with expected observation. *Information Sciences*, 282:92–110, 2014.

Lanillos P., Yañez-Zuluaga J., Ruz J. J., and Besada-Portas E. A bayesian approach for constrained multi-agent minimum time search in uncertain dynamic domains. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, pages 391–398, 2013.

Lathuilière S., Massé B., Mesejo P., and Horaud R. Neural network based reinforcement learning for audio–visual gaze control in human–robot interaction. *Pattern Recognition Letters*, 118:61–71, 2019.

Lechner M., Hasani R., Grosu R., Rus D., and Henzinger T. A. Adversarial training is not ready for robot learning. *arXiv preprint arXiv:2103.08187*, 2021.

Lee K., Lee K., Lee H., and Shin J. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in neural information processing systems*, 31, 2018.

Lewis B. P. and Beard R. W. A framework for visual return-to-home capability in gps-denied environments. In *2016 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 633–642. IEEE, 2016.

Lewis P. R., Goldingay H., and Nallur V. It's good to be different: Diversity, heterogeneity, and dynamics in collective systems. In *2014 IEEE Eighth International Conference on Self-Adaptive and Self-Organizing Systems Workshops*, pages 84–89. IEEE, 2014.

Li B., Yang L., Xiao J., Valde R., Wrenn M., and Leflar J. Collaborative mapping and autonomous parking for multi-story parking garage. *IEEE Transactions on Intelligent Transportation Systems*, 19(5):1629–1639, 2018a.

Li G., Geng E., Ye Z., Xu Y., Lin J., and Pang Y. Indoor positioning algorithm based on the improved RSSI distance model. *Sensors*, 18(9):2820, 2018b.

Li S., Da Xu L., and Zhao S. The internet of things: a survey. *Information Systems Frontiers*, 17(2):243–259, 2015.

Li S., He C., Liu M., Wan Y., Gu Y., Xie J., Fu S., and Lu K. Design and implementation of aerial communication using directional antennas: learning control in unknown communication environments. *IET Control Theory & Applications*, 13(17): 2906–2916, 2019.

Li T., Sahu A. K., Talwalkar A., and Smith V. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020.

Li X. R. and Jilkov V. P. Survey of maneuvering target tracking. part i. dynamic models. *IEEE Transactions on aerospace and electronic systems*, 39(4):1333–1364, 2003.

Li X., Girin L., Horaud R., and Gannot S. Estimation of the direct-path relative transfer function for supervised sound-source localization. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 24(11):2171–2186, 2016.

LiKamWa R., Priyantha B., Philipose M., Zhong L., and Bahl P. Energy characterization and optimization of image sensing toward continuous mobile vision. In *Proc. of the 11th Annual Int. Conf. on Mobile Systems, Applications, and Services*, pages 69–82, 2013.

Lin K., Chen M., Deng J., Hassan M. M., and Fortino G. Enhanced fingerprinting and trajectory prediction for iot localization in smart buildings. *IEEE Transactions on Automation Science and Engineering*, 13(3):1294–1307, 2016.

Lin L. and Goodrich M. A. Uav intelligent path planning for wilderness search and rescue. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 709–714. IEEE, 2009.

Lin T., Rivano H., and Le Mouël F. A survey of smart parking solutions. *IEEE Transactions on Intelligent Transportation Systems*, 18(12):3229–3253, 2017a.

Lin Y.-C., Liu M.-Y., Sun M., and Huang J.-B. Detecting adversarial attacks on neural network policies with visual foresight. *arXiv preprint arXiv:1710.00814*, 2017b.

Lissandrini N., Michieletto G., Antonello R., Galvan M., Franco A., and Cenedese A. Cooperative optimization of uavs formation visual tracking. *Robotics*, 8(3):52, 2019.

Liu C. and Hedrick J. K. Model predictive control-based target search and tracking using autonomous mobile robot with limited sensing domain. In *2017 American Control Conference (ACC)*, pages 2937–2942. IEEE, 2017.

Liu C., Li S. E., and Hedrick J. K. Model predictive control-based probabilistic search method for autonomous ground robot in a dynamic environment. In *ASME 2015 Dynamic Systems and Control Conference*. American Society of Mechanical Engineers Digital Collection, 2015.

Liu C. H., Chen Z., Tang J., Xu J., and Piao C. Energy-efficient UAV control for effective and fair communication coverage: A deep reinforcement learning approach. *IEEE Journal on Selected Areas in Communications*, 36(9):2059–2070, 2018.

Liu H., Ong Y.-S., Shen X., and Cai J. When gaussian process meets big data: A review of scalable gps. *IEEE transactions on neural networks and learning systems*, 31 (11):4405–4423, 2020.

Liu H., Darabi H., Banerjee P., and Liu J. Survey of wireless indoor positioning techniques and systems. *IEEE Trans. on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 37(6):1067–1080, 2007.

Liu J., Barreau M., Čičić M., and Johansson K. H. Learning-based traffic state reconstruction using probe vehicles. *IFAC-PapersOnLine*, 54(2):87–92, 2021a.

Liu Q., Feng F., Lan C., and Chan R. H. Va2mass: Towards the fluid filling mass estimation via integration of vision and audio learning. In *25th International Conference on Pattern Recognition Workshops (ICPRW 2020)*, pages 451–463. Springer Nature Switzerland AG, 2021b.

Lopes C. E. R. and Ruiz L. B. On the development of a multi-tier, multimodal wireless sensor network for wild life monitoring. In *IFIP Wireless Days*, pages 1–5, 2008.

Luo S., Bimbo J., Dahiya R., and Liu H. Robotic tactile perception of object properties: A review. *Mechatronics*, 48:54–67, 2017.

Luo W., Sun P., Zhong F., Liu W., Zhang T., and Wang Y. End-to-end active object tracking via reinforcement learning. In *International Conference on Machine Learning*, pages 3286–3295. PMLR, 2018.

Ma K.-C., Liu L., and Sukhatme G. S. Informative planning and online learning with sparse gaussian processes. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4292–4298. IEEE, 2017.

Magno M., Tombari F., Brunelli D., Di Stefano L., and Benini L. Multimodal video analysis on self-powered resource-limited wireless smart camera. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 3(2):223–235, 2013.

Mamduhi M. H., Hashemi E., Baras J. S., and Johansson K. H. Event-triggered add-on safety for connected and automated vehicles using road-side network infrastructure. *IFAC-PapersOnLine*, 53(2):15154–15160, 2020.

Marler R. T. and Arora J. S. Survey of multi-objective optimization methods for engineering. *Structural and multidisciplinary optimization*, 26(6):369–395, 2004.

Martinson E., Apker T., and Bugajska M. Optimizing a reconfigurable robotic microphone array. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 125–130. IEEE, 2011.

Masehian E., Jannati M., and Hekmatfar T. Cooperative mapping of unknown environments by multiple heterogeneous mobile robots with limited sensing. *Robotics and Autonomous Systems*, 87:188–218, 2017.

Masson J., Bechet M. B., and Vergassola M. Chasing information to search in random environments. *Journal of Physics A: Mathematical and Theoretical*, 42(43): 434009, 2009.

Mathew N., Smith S. L., and Waslander S. L. Multirobot rendezvous planning for recharging in persistent tasks. *IEEE Transactions on Robotics*, 31(1):128–142, 2015.

Mathur S., Jin T., Kasturirangan N., Chandrasekaran J., Xue W., Gruteser M., and Trappe W. Parknet: drive-by sensing of road-side parking statistics. In *Proceedings of the 8th international conference on Mobile systems, applications, and services*, pages 123–136, 2010.

Matthiesen D. J. Efficient beam scanning, energy allocation, and time allocation for search and detection. In *2010 IEEE International Symposium on Phased Array Systems and Technology*, pages 361–368. IEEE, 2010.

Mavrinac A. and Chen X. Modeling coverage in camera networks: A survey. *International journal of computer vision*, 101(1):205–226, 2013.

Mavrommati A., Tzorakoleftherakis E., Abraham I., and Murphey T. D. Real-time area coverage and target localization using receding-horizon ergodic exploration. *IEEE Transactions on Robotics*, 34(1):62–80, 2017.

Maza I. and Ollero A. Multiple uav cooperative searching operation using polygon area decomposition and efficient coverage algorithms. In *Distributed Autonomous Robotic Systems 6*, pages 221–230. Springer, 2007.

McHutchon A. and Rasmussen C. Gaussian process training with input noise. *Advances in Neural Information Processing Systems*, 24:1341–1349, 2011.

Meera A. A., Popović M., Millane A., and Siegwart R. Obstacle-aware adaptive informative path planning for uav-based target search. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 718–724. IEEE, 2019.

Mehra A., Kailkhura B., Chen P.-Y., and Hamm J. How robust are randomized smoothing based defenses to data poisoning? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13244–13253, 2021.

Melis M., Demontis A., Biggio B., Brown G., Fumera G., and Roli F. Is deep learning safe for robot vision? adversarial examples against the icub humanoid. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 751–759, 2017.

Micheloni C., Rinner B., and Foresti G. L. Video analysis in pan-tilt-zoom camera networks. *IEEE Signal Processing Magazine*, 27(5):78–90, 2010.

Michieletto G. *Multi-Agent Systems in Smart Environments-from sensor networks to aerial platform formations.* PhD thesis, University of Padova, 2018.

Miki T., Popović M., Gawel A., Hitz G., and Siegwart R. Multi-agent time-based decision-making for the search and action problem. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2365–2372. IEEE, 2018.

Milighetti G., Emter T., Kuntze H.-B., Bechler D., and Kroschel K. Combined visual-acoustic grasping for humanoid robots. In *2006 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, pages 1–6. IEEE, 2006.

Mohamed A. E.-M. A. and El-Hadidy M. A. A. Coordinated search for a conditionally deterministic target motion in the plane. *European Journal of Mathematical Sciences*, 2(3):272–295, 2013.

Monroy J., Ruiz-Sarmiento J.-R., Moreno F.-A., Melendez-Fernandez F., Galindo C., and Gonzalez-Jimenez J. A semantic-based gas source localization with a mobile robot combining vision and chemical sensing. *Sensors*, 18(12):4174, 2018.

Moons B., Bankman D., and Verhelst M. Embedded deep learning. *Embedded Deep Learning*, 2019.

Moore C. J., Chua A. J., Berry C. P., and Gair J. R. Fast methods for training gaussian processes on large datasets. *Royal Society open science*, 3(5):160125, 2016.

Mousavi H. K., Liu G., Yuan W., Takáč M., Muñoz-Avila H., and Motee N. A layered architecture for active perception: Image classification using deep reinforcement learning. *arXiv preprint arXiv:1909.09705*, 2019.

Mukerjee M. K., Naylor D., Jiang J., Han D., Seshan S., and Zhang H. Practical, real-time centralized control for cdn-based live video delivery. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, pages 311–324, 2015.

Murphy S. Ó., Sreenan C., and Brown K. N. Autonomous unmanned aerial vehicle for search and rescue using software defined radio. In *2019 IEEE 89th Vehicular Technology Conference (VTC2019-Spring)*, pages 1–6. IEEE, 2019.

Nagaty A., Thibault C., Trentini M., and Li H. Probabilistic cooperative target localization. *IEEE Transactions on Automation Science and Engineering*, 12(3):786–794, 2015.

Nakamura E. F., Loureiro A. A., and Frery A. C. Information fusion for wireless sensor networks: Methods, models, and classifications. *ACM Computing Surveys (CSUR)*, 39(3):9–es, 2007.

Nava M., Guzzi J., Chavez-Garcia R. O., Gambardella L. M., and Giusti A. Learning long-range perception using self-supervision from short-range sensors and odometry. *IEEE Robotics and Automation Letters*, 4(2):1279–1286, 2019.

Nayak J., Gonzalez-Argueta L., Song B., Roy-Chowdhury A., and Tuncel E. Multi-target tracking through opportunistic camera control in a resource constrained multimodal sensor network. In *2008 Second ACM/IEEE International Conference on Distributed Smart Cameras*, pages 1–10. IEEE, 2008.

Nguyen J., Lawrance N., Fitch R., and Sukkarieh S. Energy-constrained motion planning for information gathering with autonomous aerial soaring. In *2013 IEEE International Conference on Robotics and Automation*, pages 3825–3831. IEEE, 2013.

Noori N., Renzaglia A., Vander Hook J., and Isler V. Constrained probabilistic search for a one-dimensional random walker. *IEEE Transactions on Robotics*, 32(2): 261–274, 2016.

O'Callaghan S., Ramos F. T., and Durrant-Whyte H. Contextual occupancy maps using gaussian processes. In *2009 IEEE International Conference on Robotics and Automation*, pages 1054–1060. IEEE, 2009.

Osia S. A., Shahin Shamsabadi A., Sajadmanesh S., Taheri A., Katevas K., Rabiee H. R., Lane N. D., and Haddadi H. A hybrid deep learning architecture for privacy-preserving mobile analytics. *IEEE Internet of Things Journal*, 7(5):4505–4518, 2020.

Ousingsawat J. and Earl M. G. Modified lawn-mower search pattern for areas comprised of weighted regions. In *2007 American Control Conference*, pages 918–923. IEEE, 2007.

Pack D. J., DeLima P., Toussaint G. J., and York G. Cooperative control of uavs for localization of intermittently emitting mobile targets. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(4):959–970, 2009.

Pang Y. L., Xompero A., Oh C., and Cavallaro A. Towards safe human-to-robot handovers of unknown containers, 2021.

Papanikolopoulos N. The framework of controlled active vision. *Mathematical and computer modelling*, 24(5-6):145–163, 1996.

Park J. H., Chao H.-C., Hussain S., and Yen N. Y. Pervasive sensing technologies and emerging trends, 2014.

Park M. and Oh H. Cooperative information-driven source search and estimation for multiple agents. *Information Fusion*, 54:72–84, 2020.

Park S.-M., Lee J.-H., and Chwa K.-Y. Visibility-based pursuit-evasion in a polygonal region by a searcher. In *International Colloquium on Automata, Languages, and Programming*, pages 456–468. Springer, 2001.

Parker L. E., Rus D., and Sukhatme G. S. Multiple mobile robot systems. In *Springer Handbook of Robotics*, pages 1335–1384. Springer, 2016.

Patel M., Girgensohn A., and Biehl J. Fusing map information with a probabilistic sensor model for indoor localization using RF beacons. In *Int. Conf. on Indoor Positioning and Indoor Navigation (IPIN)*, pages 1–8. IEEE, 2018.

Patten T., Martens W., and Fitch R. Monte carlo planning for active object classification. *Autonomous Robots*, 42(2):391–421, 2018.

Perez-Carabaza S., Besada-Portas E., Lopez-Orozco J. A., and Jesus M. Ant colony optimization for multi-uav minimum time search in uncertain domains. *Applied Soft Computing*, 62:789–806, 2018.

Pérez-Carabaza S., Besada-Portas E., Lopez-Orozco J. A., and Pajares G. Minimum time search in real-world scenarios using multiple uavs with onboard orientable cameras. *Journal of Sensors*, 2019, 2019a.

Pérez-Carabaza S., Scherer J., Rinner B., López-Orozco J. A., and Besada-Portas E. UAV trajectory optimization for minimum time search with communication constraints and collision avoidance. *Engineering Applications of Artificial Intelligence*, 85:357–371, 2019b.

Peri N., Khorramshahi P., Saketh Rambhatla S., Shenoy V., Rawat S., Chen J.-C., and Chellappa R. Towards real-time systems for vehicle re-identification, multi-camera tracking, and anomaly detection. In *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition Workshops*, pages 622–623, 2020.

Popovic M. *Enviromental Mapping and Informative Path Planning for UAV-based Active Sensing.* PhD thesis, ETH Zurich, 2019.

Qian X., Brutti A., Lanz O., Omologo M., and Cavallaro A. Audio-visual tracking of concurrent speakers. *IEEE Transactions on Multimedia*, pages 1–1, 2021a.

Qian X., Brutti A., Lanz O., Omologo M., and Cavallaro A. Audio-visual tracking of concurrent speakers. *IEEE Transactions on Multimedia*, 2021b.

Qin F., Dai X., and Mitchell J. E. Effective-snr estimation for wireless sensor network using kalman filter. *Ad Hoc Networks*, 11(3):944–958, 2013.

Queralta J. P., Taipalmaa J., Pullinen B. C., Sarker V. K., Gia T. N., Tenhunen H., Gabbouj M., Raitoharju J., and Westerlund T. Collaborative multi-robot systems for search and rescue: Coordination and perception. *arXiv preprint arXiv:2008.12610*, 2020.

Radak J., Baulig L., Bijak D., Schowalter C., and Frey H. Moving towards wireless sensors using rssi measurements and particle filtering. In *Proceedings of the 14th ACM Symposium on Performance Evaluation of Wireless Ad Hoc, Sensor, & Ubiquitous Networks*, pages 33–40, 2017.

Radmard S. and Croft E. A. Active target search for high dimensional robotic systems. *Autonomous Robots*, 41(1):163–180, 2017.

Raj T., Hashim F. H., Huddin A. B., Ibrahim M. F., and Hussain A. A survey on lidar scanning mechanisms. *Electronics*, 9(5):741, 2020.

Rajabioun T. and Ioannou P. A. On-street and off-street parking availability prediction using multivariate spatiotemporal models. *IEEE Transactions on Intelligent Transportation Systems*, 16(5):2913–2924, 2015.

Rajasekaran R. K., Ahmed N., and Frew E. Bayesian fusion of unlabeled vision and rf data for aerial tracking of ground targets. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1629–1636. IEEE, 2020.

Ramirez J.-P., Doucette E. A., Curtis J. W., and Gans N. Moving target acquisition through state uncertainty minimization. In *2014 American Control Conference*, pages 3425–3430. IEEE, 2014.

Rao N. S. V. On fusers that perform better than best sensor. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(8):904–909, 2001.

Rasmussen C. E. Gaussian processes in machine learning. In *Summer School on Machine Learning*, pages 63–71. Springer, 2003.

Redmon J., Divvala S., Girshick R., and Farhadi A. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 779–788, 2016.

Rinner B., Esterle L., Simonjan J., Nebehay G., Pflugfelder R., Domínguez G. F., and Lewis P. R. Self-aware and self-expressive camera networks. *Computer*, 48(7):21–28, 2015.

Rinner B., Winkler T., Schriebl W., Quaritsch M., and Wolf W. The evolution from single to pervasive smart cameras. In *2008 Second ACM/IEEE International Conference on Distributed Smart Cameras*, pages 1–10. IEEE, 2008.

Rinner B. and Wolf W. An introduction to distributed smart cameras. *Proc. of the IEEE*, 96(10):1565–1575, 2008.

Ristic B., Angley D., Moran B., and Palmer J. L. Autonomous multi-robot search for a hazardous source in a turbulent environment. *Sensors*, 17(4):918, 2017.

Ristic B., Skvortsov A., and Gunatilaka A. A study of cognitive strategies for an autonomous search. *Information Fusion*, 28:1–9, 2016.

Rizk Y., Awad M., and Tunstel E. W. Decision making in multiagent systems: A survey. *IEEE Transactions on Cognitive and Developmental Systems*, 10(3):514–529, 2018.

Rizk Y., Awad M., and Tunstel E. W. Cooperative heterogeneous multi-robot systems: a survey. *ACM Computing Surveys (CSUR)*, 52(2):1–31, 2019.

Röbesaat J., Zhang P., Abdelaal M., and Theel O. An improved BLE indoor localization with Kalman-based fusion: An experimental study. *Sensors*, 17(5):951, 2017.

Robin C. and Lacroix S. Multi-robot target detection and tracking: taxonomy and survey. *Autonomous Robots*, 40(4):729–760, 2016.

Rodríguez A., Negro J. J., Mulero M., Rodríguez C., Hernández-Pliego J., and Bustamante J. The eye in the sky: combined use of unmanned aerial systems and gps data loggers for ecological research and conservation of small birds. *PLoS One*, 7(12):e50336, 2012.

Rodríguez M., Al-Kaff A., Madridano Á., Martín D., and de la Escalera A. Wilderness search and rescue with heterogeneous multi-robot systems. In *2020 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 110–116. IEEE, 2020.

Rodríguez-Molina A., Mezura-Montes E., Villarreal-Cervantes M. G., and Aldape-Pérez M. Multi-objective meta-heuristic optimization in intelligent control: A survey on the controller tuning problem. *Applied Soft Computing*, 93:106342, 2020.

Rondón R., Mahmood A., Grimaldi S., and Gidlund M. Understanding the performance of bluetooth mesh: Reliability, delay, and scalability analysis. *IEEE Internet of Things Journal*, 7(3):2089–2101, 2019.

Rosić M. B., Simić M. I., and Pejović P. V. Passive target localization problem based on improved hybrid adaptive differential evolution and nelder-mead algorithm. *Journal of Sensors*, 2020, 2020.

Ryan A. and Hedrick J. K. Particle filter based information-theoretic active sensing. *Robotics and Autonomous Systems*, 58(5):574–584, 2010.

Ryan A. D. *Information-theoretic control for mobile sensor teams.* University of California, Berkeley, 2008.

Sa I., Popović M., Khanna R., Chen Z., Lottes P., Liebisch F., Nieto J., Stachniss C., Walter A., and Siegwart R. Weedmap: A large-scale semantic weed mapping framework using aerial multispectral imaging and deep neural network for precision farming. *Remote Sensing*, 10(9):1423, 2018.

SAE On-Road Automated Vehicle Standards Committee and others . Taxonomy and definitions for terms related to on-road motor vehicle automated driving systems. *SAE Standard J*, 3016:1–16, 2014.

Sahar A. and Han D. An lstm-based indoor positioning method using wi-fi signals. In *Proceedings of the 2nd International Conference on Vision, Image and Signal Processing*, pages 1–5, 2018.

Sandino J., Vanegas F., Gonzalez F., and Maire F. Autonomous uav navigation for active perception of targets in uncertain and cluttered environments. In *2020 IEEE Aerospace Conference*, pages 1–12. IEEE, 2020.

SanMiguel J. C. and Cavallaro A. Cost-aware coalitions for collaborative tracking in resource-constrained camera networks. *IEEE Sensors Journal*, 15(5):2657–2668, 2014.

SanMiguel J. C. and Cavallaro A. Efficient estimation of target detection quality. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 915–919. IEEE, 2017.

SanMiguel J. C., Micheloni C., Shoop K., Foresti G. L., and Cavallaro A. Self-reconfigurable smart camera networks. *Computer*, 47(5):67–73, 2014.

SanMiguel J. C. and Cavallaro A. Energy consumption models for smart camera networks. *IEEE Transactions on Circuits and Systems for Video Technology*, 27(12): 2661–2674, 2016.

Saragih C. F. D., Kinasih F. M. T. R., Machbub C., Rusmin P. H., and Rohman A. S. Visual servo application using model predictive control (mpc) method on pan-tilt camera platform. In *2019 6th International Conference on Instrumentation, Control, and Automation (ICA)*, pages 1–7. IEEE, 2019a.

Saragih C. F. D., Kinasih F. M. T. R., Machbub C., Rusmin P. H., and Rohman A. S. Visual servo application using model predictive control (mpc) method on pan-tilt camera platform. In *2019 6th International Conference on Instrumentation, Control, and Automation (ICA)*, pages 1–7. IEEE, 2019b.

Schlotfeldt B., Atanasov N., and Pappas G. J. Maximum information bounds for planning active sensing trajectories. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4913–4920. IEEE, 2019.

Schmidt R. Multiple emitter location and signal parameter estimation. *IEEE transactions on antennas and propagation*, 34(3):276–280, 1986.

Schwager M., Julian B. J., Angermann M., and Rus D. Eyes in the sky: Decentralized control for the deployment of robotic camera networks. *Proceedings of the IEEE*, 99(9):1541–1561, 2011.

Senanayake M., Senthooran I., Barca J. C., Chung H., Kamruzzaman J., and Murshed M. Search and tracking algorithms for swarms of robots: A survey. *Robotics and Autonomous Systems*, 75:422–434, 2016.

Senanayake R., O'Callaghan S., and Ramos F. Learning highly dynamic environments with stochastic variational inference. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2532–2539. IEEE, 2017.

Shahidian S. A. A. and Soltanizadeh H. Single-and multi-UAV trajectory control in RF source localization. *Arabian Journal for Science and Engineering*, 42(2):459–466, 2017.

Shahriari B., Swersky K., Wang Z., Adams R. P., and De Freitas N. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015.

Shubina K. and Tsotsos J. K. Visual search for an object in a 3d environment using a mobile robot. *Computer Vision and Image Understanding*, 114(5):535–547, 2010.

Sibanyoni S. V., Ramotsoela D. T., Silva B. J., and Hancke G. P. A 2-d acoustic source localization system for drones in search and rescue missions. *IEEE Sensors Journal*, 19(1):332–341, 2018.

Siegel J. E., Erb D. C., and Sarma S. E. A survey of the connected vehicle landscape—architectures, enabling technologies, applications, and development areas. *IEEE Transactions on Intelligent Transportation Systems*, 19(8):2391–2406, 2017.

Siekkinen M., Hiienkari M., Nurminen J. K., and Nieminen J. How low energy is bluetooth low energy? comparative measurements with zigbee/802.15. 4. In *2012 IEEE wireless communications and networking conference workshops (WCNCW)*, pages 232–237. IEEE, 2012.

Skolnik M. I. and others . *Introduction to radar systems*, volume 3. McGraw-hill New York, 1980.

Smith A. *Sequential Monte Carlo methods in practice*. Springer Science & Business Media, 2013.

Smith D. and Singh S. Approaches to multisensor data fusion in target tracking: A survey. *IEEE transactions on knowledge and data engineering*, 18(12):1696–1710, 2006.

Snoek J., Larochelle H., and Adams R. P. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pages 2951–2959, 2012.

Sokullu R., Balcı A., and Demir E. The role of drones in ambient assisted living systems for the elderly. In *Enhanced Living Environments*, pages 295–321. Springer, 2019.

Stone L. D., Keller C. M., Kratzke T. M., and Strumpfer J. P. Search analysis for the underwater wreckage of air france flight 447. In *14th International Conference on Information Fusion*, pages 1–8. IEEE, 2011.

Stoyanova T., Kerasiotis F., Prayati A., and Papadopoulos G. Evaluation of impact factors on rss accuracy for localization and tracking applications. In *Proceedings*

of the 5th ACM international workshop on Mobility management and wireless access, pages 9–16, 2007.

Sui K., Zhou M., Liu D., Ma M., Pei D., Zhao Y., Li Z., and Moscibroda T. Characterizing and improving wifi latency in large-scale operational networks. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*, pages 347–360, 2016.

Sun Y., Xiao J., Li X., and Cabrera-Mora F. Adaptive source localization by a mobile robot using signal power gradient in sensor networks. In *IEEE GLOBECOM 2008-2008 IEEE Global Telecommunications Conference*, pages 1–5. IEEE, 2008.

Svečko J., Malajner M., and Gleich D. Distance estimation using rssi and particle filter. *ISA transactions*, 55:275–285, 2015.

Tadesse G. A., Bent O., Marcenaro L., Weldemariam K., and Cavallaro A. Privacy-aware human activity recognition from a wearable camera: Highlights from the ieee video and image processing cup 2019 student competition [sp competitions]. *IEEE Signal Processing Magazine*, 37(3):168–172, 2020.

Talebpour A. and Mahmassani H. S. Influence of connected and autonomous vehicles on traffic flow stability and throughput. *Transportation Research Part C: Emerging Technologies*, 71:143–163, 2016.

Tavafoghi H., Poolla K., and Varaiya P. A queuing approach to parking: Modeling, verification, and prediction. *arXiv preprint arXiv:1908.11479*, 2019.

Tenney R. R. and Sandell N. R. Detection with distributed sensors. *IEEE Transactions on Aerospace and Electronic Systems*, AES-17(4):501–510, 1981.

Tosi J., Taffoni F., Santacatterina M., Sannino R., and Formica D. Performance evaluation of bluetooth low energy: A systematic review. *Sensors*, 17(12):2898, 2017.

Treurniet J. J., Sarkar C., Prasad R. V., and De Boer W. Energy consumption and latency in ble devices under mutual interference: An experimental study. In *2015 3rd International Conference on Future Internet of Things and Cloud*, pages 333–340. IEEE, 2015.

Tron R., Vidal R., and Terzis A. Distributed pose averaging in camera networks via consensus on se (3). In *2008 Second ACM/IEEE International Conference on Distributed Smart Cameras*, pages 1–10. IEEE, 2008.

Trummel K. and Weisinger J. The complexity of the optimal searcher path problem. *Operations Research*, 34(2):324–327, 1986.

Twigg J. N., Fink J. R., Paul L. Y., and Sadler B. M. Rss gradient-assisted frontier exploration and radio source localization. In *2012 IEEE International Conference on Robotics and Automation*, pages 889–895. IEEE, 2012.

Ucinski D. *Optimal measurement methods for distributed parameter system identification.* CRC press, 2004.

Uhlemann E. Introducing connected vehicles [connected vehicles]. *IEEE Vehicular Technology Magazine*, 10(1):23–31, 2015.

Uhlemann E. Time for autonomous vehicles to connect [connected vehicles]. *IEEE vehicular technology magazine*, 13(3):10–13, 2018.

Van Nguyen H., Chen F., Chesser J., Rezatofighi H., and Ranasinghe D. Lavapilot: Lightweight uav trajectory planner with situational awareness for embedded autonomy to track and locate radio-tags. *arXiv preprint arXiv:2007.15860*, 2020.

Vander Hook J., Tokekar P., and Isler V. Algorithms for cooperative active localization of static targets with mobile bearing sensors under communication constraints. *IEEE Transactions on Robotics*, 31(4):864–876, 2015.

Vargas E., Brown K., and Subr K. Impact of microphone array configurations on robust indirect 3d acoustic source localization. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3221–3225. IEEE, 2018.

Varotto L., Cenedese A., and Cavallaro A. Probabilistic radio-visual active sensing for search and tracking, 2021.

Varotto L. and Cenedese A. Online and adaptive parking availability mapping: An uncertainty-aware active sensing approach for connected vehicles, 2021a.

Varotto L. and Cenedese A. Probabilistic rf-assisted camera wake-up through self-supervised gaussian process regression, 2021b.

Varotto L. and Cenedese A. Transmitter discovery through radio-visual probabilistic active sensing, 2021c.

Varotto L., Fabris M., Michieletto G., and Cenedese A. Distributed dual quaternion based localization of visual sensor networks. In *2019 18th European Control Conference (ECC)*, pages 1836–1841. IEEE, 2019a.

Varotto L., Zampieri A., and Cenedese A. Street sensors set selection through road network modeling and observability measures. In *2019 27th Mediterranean Conference on Control and Automation (MED)*, pages 392–397, 2019b.

Vergassola M., Villermaux E., and Shraiman B. I. 'infotaxis' as a strategy for searching without gradients. *Nature*, 445(7126):406–409, 2007.

Vicentini F. Collaborative robotics: a survey. *Journal of Mechanical Design*, 143(4), 2021.

Viola P. and Jones M. Rapid object detection using a boosted cascade of simple features. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages I–I, 2001.

Viseras A., Wiedemann T., Manss C., Magel L., Mueller J., Shutin D., and Merino L. Decentralized multi-agent exploration with online-learning of gaussian processes. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4222–4229. IEEE, 2016.

Vollmer M. and Möllmann K.-P. High speed and slow motion: the technology of modern high speed cameras. *Physics Education*, 46(2):191, 2011.

Voulodimos A., Doulamis N., Doulamis A., and Protopapadakis E. Deep learning for computer vision: A brief review. *Computational intelligence and neuroscience*, 2018, 2018.

Wang L., Sanchez-Matilla R., and Cavallaro A. Tracking a moving sound source from a multi-rotor drone. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2511–2516. IEEE, 2018.

Wang L., Sanchez-Matilla R., and Cavallaro A. Audio-visual sensing from a quadcopter: dataset and baselines for source localization and sound enhancement. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5320–5325. IEEE, 2019.

Wang L. *Model predictive control system design and implementation using MATLAB®*. Springer Science & Business Media, 2009.

Weiser M., Gold R., and Brown J. S. The origins of ubiquitous computing research at parc in the late 1980s. *IBM systems journal*, 38(4):693–696, 1999.

Westfechtel T., Ohno K., Mizuno N., Hamada R., Kojima S., and Tadokoro S. Parking spot estimation and mapping method for mobile robots. *IEEE Robotics and Automation Letters*, 3(4):3371–3378, 2018.

White K., Williams J., and Hoffensetz P. Radar sensor management for detection and tracking. In *2008 11th International Conference on Information Fusion*, pages 1–8. IEEE, 2008.

Williams C. K. and Rasmussen C. E. *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006.

Wong E.-M., Bourgault F., and Furukawa T. Multi-vehicle bayesian search for multiple lost targets. In *Proceedings of the 2005 ieee international conference on robotics and automation*, pages 3169–3174. IEEE, 2005.

Wong K., Gu Y., and Kamijo S. Mapping for autonomous driving: Opportunities and challenges. *IEEE Intelligent Transportation Systems Magazine*, 13(1), 2021.

Wu J., Chen X.-Y., Zhang H., Xiong L.-D., Lei H., and Deng S.-H. Hyperparameter optimization for machine learning models based on bayesian optimization. *Journal of Electronic Science and Technology*, 17(1):26–40, 2019.

Xiong J. and Jamieson K. Arraytrack: A fine-grained indoor location system. In *Presented as part of the 10th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 13)*, pages 71–84, 2013.

Xu B., Zhu X., and Zhu H. An efficient indoor localization method based on the long short-term memory recurrent neuron network. *IEEE Access*, 7:123912–123921, 2019.

Xu J. and Zhang J. Exploration-exploitation tradeoffs in metaheuristics: Survey and analysis. In *Proceedings of the 33rd Chinese Control Conference*, pages 8633–8638. IEEE, 2014.

Xue T., Wang W., Ma J., Liu W., Pan Z., and Han M. Progress and prospects of multimodal fusion methods in physical human–robot interaction: A review. *IEEE Sensors Journal*, 20(18):10355–10370, 2020.

Yang H., Han Q.-L., Ge X., Ding L., Xu Y., Jiang B., and Zhou D. Fault-tolerant cooperative control of multiagent systems: A survey of trends and methodologies. *IEEE Transactions on Industrial Informatics*, 16(1):4–17, 2019a.

Yang T., Yi X., Wu J., Yuan Y., Wu D., Meng Z., Hong Y., Wang H., Lin Z., and Johansson K. H. A survey of distributed optimization. *Annual Reviews in Control*, 47:278–305, 2019b.

Ye Y. and Tsotsos J. K. Sensor planning for 3d object search. *Computer Vision and Image Understanding*, 73(2):145–168, 1999.

Yi W., Morelande M. R., Kong L.-J., and Yang J.-Y. Target tracking for an unknown and time-varying number of targets via particle filtering. In *2012 15th International Conference on Information Fusion*, pages 309–316. IEEE, 2012.

Yu K., Budhiraja A. K., and Tokekar P. Algorithms for routing of unmanned aerial vehicles with mobile recharging stations. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–5. IEEE, 2018.

Yu W., Liang F., He X., Hatcher W. G., Lu C., Lin J., and Yang X. A survey on the edge computing for the internet of things. *IEEE access*, 6:6900–6919, 2017.

Zafari F., Gkelias A., and Leung K. K. A survey of indoor localization systems and technologies. *IEEE Communications Surveys & Tutorials*, 21(3):2568–2599, 2019.

Zafari F., Papapanagiotou I., and Hacker T. J. A novel bayesian filtering based algorithm for rssi-based indoor localization. In *2018 IEEE International Conference on Communications (ICC)*, pages 1–7. IEEE, 2018.

Zanella A. Best practice in RSS measurements and ranging. *IEEE Communications Surveys & Tutorials*, 18(4):2662–2686, 2016.

Zanella A., Bui N., Castellani A., Vangelista L., and Zorzi M. Internet of things for smart cities. *IEEE Internet of Things journal*, 1(1):22–32, 2014.

Zhang Y., Wang W., Du W., Qian C., and Yang H. Coloured petri net-based active sensing system of real-time and multi-source manufacturing information for smart factory. *The International Journal of Advanced Manufacturing Technology*, 94 (9):3427–3439, 2018.

Zhang Y., Gu Y., Vlatkovic V., and Wang X. Progress of smart sensor and smart sensor networks. In *Fifth World Congress on Intelligent Control and Automation (IEEE Cat. No. 04EX788)*, volume 4, pages 3600–3606. IEEE, 2004.

Zhou K. *Active sensing with applications to mobile robotics.* PhD thesis, University of Minnesota, 2012.

Zhuang Y., Yang J., Li Y., Qi L., and El-Sheimy N. Smartphone-based indoor localization with bluetooth low energy beacons. *Sensors*, 16(5):596, 2016.