Adapting the System to Users Based on Implicit Data: Ethical Risks and Possible Solutions

Anna Spagnolli[1(✉)], Mauro Conti[2], Giorgia Guerra[3], Jonathan Freeman[4], David Kirsh[5], and Aimee van Wynsberghe[6]

[1] Human Inspired Technologies Research Centre and Department of General Psychology, Padua University, Padua, Italy
anna.spagnolli@unipd.it
[2] Human Inspired Technologies Research Centre and Department of Mathematics, Padua University, Padua, Italy
mauro.conti@unipd.it
[3] Departmentof Political Science, Law and International Studies, Padua University, Padua, Italy
giorgia.guerra@unipd.it
[4] Goldsmiths College, University of London, London, UK
J.Freeman@gold.ac.uk
[5] University of California at San Diego, San Diego, USA
kirsh@ucsd.edu
[6] Unversity of Twente, Enschede, Netherlands
a.l.vanwynsberghe@utwente.nl

Abstract. Symbiotic systems are systems that gather personal data *implicitly* provided by the user, derive a *profile/model* of the user from such data and *adjust* their output/service according to their notion of what would be desirable to the user thus modeled. Because of these three characteristics, symbiotic systems represent a step forward towards facilitated, simplified, user-friendly digital devices, or do they? Here we propose three cases describing realistic applications of symbiotic systems that potentially encapsulate some serious risk to their users. Experts of five different domains (i.e., ethics, security, law, human-computer interaction and psychology) dissect each case to identify the risks to the users and derive some possible minimization strategies. This panel aims at contributing to a beneficial development of symbiotic systems as it can be achieved by increasing users' discernment and awareness of their consequences for society and everyday life.

Keywords: Symbiotic system · Implicit data · Ethics · Security · Design · Risks · Information leakage · Privacy · Awareness

1 Introduction

The increasing pervasiveness and refinement of systems that can acquire personal data from users based on the users' own actions on the system - or actions in the environment where the system sensors are positioned - marks a qualitative, discrete change

L. Gamberini et al. (Eds.): Symbiotic 2016, LNCS 9961, pp. 5–22, 2017.
DOI: 10.1007/978-3-319-57753-1_2

in the way in which the human interaction with technologies articulates. The label 'symbiotic systems' is meant to mark such a change, to make it visible, to single it out from a continuous line of technological development in physiological computing, machine learning, and sensing devices.

Symbiotic technologies establish a **symbiotic relationship** with the human users. They extract information about the users' state, behaviors or preferences, and adapt their services to the users' profile thereby created. Figure 1 illustrates the rationale of this process: a human who interprets his/her relation with the environment and acts consequently is tapped onto by a prosthetic system that can interpret the users' state and response and can plan interventions on the environment based on a set of built-in criteria. This model highlights three main characteristics of symbiotic technologies:

- symbiotic technologies **acquire implicit information about the user**, namely data that users might not be aware of giving out and whose release and content they might not be able to control;
- the implicit information is elaborated so as to create a **model of the user(s)**;
- the system takes over some of the users' **decisions**, according to the user model and to what it is programmed to consider as the best way to serve that model of user.

This scheme can foresee several variations: the symbiotic system can get information from many human beings instead of just one, thanks to current network technology and sensors; and the users of the system can be different from the humans giving out implicit information. These variations notwithstanding, the core of the system remains centered on the capturing/modeling/deciding process that makes symbiotic systems able to bypass the user to which they are connected.

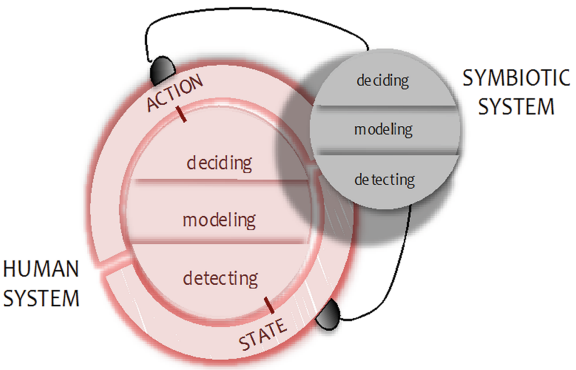

Fig. 1. A schematic representation of the symbiotic system tapping into a human system to by-pass its detecting/modeling/deciding processes.

On part of the users feeding the system with personal information at low or no cost, the symbiotic relation is a transaction out of which direct or indirect benefits are expected, mostly in the form of having their tasks simplified, their peculiarities acknowledged and their preferences indulged by the system. From this point of view, symbiotic systems bring to the next level the notion of a user-centered system, where the system tailors to the user'state without the user even needing to ask. At the same time, however, the exact notion of what makes a system 'user-centred' is questioned by symbiotic systems, given the possibility that such systems - if incautiously designed - create serious security, ethic, legal and psychological risks to the user. While much debate on the risks of technologies merging with humans focuses on the bodily level of such merge (for example, humans with x-ray sight, robots with realistic emotional responses or networked machines taking control of human society [2]), the merge is already occurring pervasively at a more functional, **information** level. Recently, for instance, concerns have been raised about the ethics risks of merging Virtual Reality analytics and social networks for persuasive purposes [11] or to gather face recognition data in VR peripherals to allegedly customize the users' experience [5]. This chapter reports the reflections on such possible risks made by an interdisciplinary panel of experts gathered during Symbiotic 2016 in Padua. The panel tried to identify concrete ethic risks of current pervasive symbiotic systems from an interdisciplinary perspective, and to have such risks if not solved - at least effectively framed.

To facilitate discussion, three likely everyday life applications of symbiotic systems were presented to the panel's attention, which bore some problematic implications to users and which focused respectively on detecting, modeling and deciding. Each case was examined from five different perspectives: **ethics** (Aimee van Wynsberghe, President of the Foundation for Responsible Robotics and Assistant Professor at the University of Twente, the Netherlands), **information security** (Mauro Conti, University of Padua, Italy), **law** (Giorgia Guerra, University of Padua, Italy), **human-computer interaction** (David Kirsh, University of California, San Diego, US) and **psychology** (Jonathan Freeman, University of London, Goldsmith College, UK perspective). Anna Spagnolli organized and chaired the panel. The examination of each of the three cases is reported in the following sections. The main points are summarized in the final section of this chapter.

2 Case 1, Detecting Users: The Unusual Job Interview

The first case describes the use of symbiotic system where the detection of personal data for uses that are counterproductive to the unaware data owner. Indications about the verisimilarity of such a case are provided by [6, 8, 12, 14] showing that the technology and the preconditions for much of the components of the scenario are already in place.

An international firm is hiring personnel in the accounting department. At the selection interview, the candidates are asked to use a pc and fill in an electronic form, collecting basic demographic information but also data about previous jobs and a description of the reasons why they think they suit the vacant position. Each candidate is also asked to sign an informed consent allowing the firm to acquire the responses to the questionnaire, which will be used for archival reasons as well as to direct the subsequent interview with the human resources manager. The form specifies that the hiring company will use the written responses to the questionnaire as well as data on the pc usage while typing the responses (time, pressure and trajectory of the mouse movements). The questionnaire is mandatory and the firm commits to keep all collected data confidential. None of the candidates is aware that data on typing behavior can be used not just to identify the user but also to detect the probability that s/he is lying.

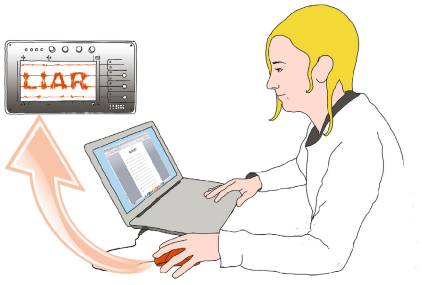

2.1 Ethics

RISK: There are multiple risks involved in this case, e.g. lack of informed consent, lack of protection of best interests of the user, risk of deceit, and lack of transparency regarding data collection and data use. I suggest the area of risk involved in case 1 derives from a combination of the above mentioned risks, i.e. the system collects information about a user without informed consent or transparency for a purpose that is detrimental to the user while the benefits from such a collection belong only to the company collecting data through the system. More specifically, this represents a case of treating direct users as a **'means to an end'**, as mere tools to achieve the company's goal, and this would be an ethically questionable position in which to put a user. Once we can justify the use of individuals for this kind of practice – as a means to an end – the list of ethically questionable activities that may be condoned increases exponentially.

SOLUTION: A possible solution could be to mimic the current practices in academic institutes whereby an ethics committee is established to monitor and approve research practices. In so doing, ethical approval has to be obtained for studies to be conducted and to provide guidance on how to do so as well as in order to protect the rights and interests of the participants involved. Outside the academy this could happen on the form of an **Institutional Review Board for companies or an advocacy group**

that transcends the company as a body independent from it (i.e., not in the form of a department in the company itself). Independence is important as a way of striving for objective decisions and avoiding persuasion of ethics decisions.

2.2 Security

RISK: From a security perspective, the problem in case 1 is that the user provides more information than s/he means to. S/he only means to input text in the form, but in fact s/he provides a lot more information while doing so. This process, i.e., trying to use data to infer some other information from it, is known as **information leakage**. There are plenty of other examples of information leakage. For example, analyzing the incoming and outgoing network traffic, the energy consumption patterns and the movement of the phone recorded when using a mobile phone, without having access to content stored on the phone, is sufficient to infer information about the user such as: the application installed on the phone, sex, age range, preferred language, text typed etc. Information leakage is a very real technical possibility, and represents a security risk because such information can be exploited in a malicious way.

SOLUTION: Despite the transaction described in this case represents such an infringement to several policies and users' rights that in some domains it would never been accepted (e.g., in academics), it is nonetheless a very pervasive kind of transaction nowadays. The use of data for extracting more information than it intended to convey by the owner of that data is technically possible and this must make us suspicious that such possibility is actually exploited. It would then be good to find technical means to prevent such leakage to occur in the first place. The best protection in this case would be to **use one's own interface** to use when typing and then avoid typing on other parties' interfaces. And before that I consider it necessary to increase the **awareness** in the user that information leakage is a likely occurrence, so they can ask for using their own interfaces in cases such as the one described in the scenario (as well as for policies, etc.).

2.3 Law

RISK: The data collected here is used for a purpose different from the one for which it was collected; it can be considered as a **deceitful use of data**, and in this specific case this kind of use could create a risk of discriminatory behavior during the recruitment process. The informed consent is insufficient, being not transparent about the use of data. The additional complication is that the collected information, in order to be used, has to be **interpreted**. The reliability information extracted from pc usage is also questionable and still currently under debate. Consequently, this procedure could limit the users' self-determination, since candidates would have behaved differently had they known the actual purpose of the data collection.

SOLUTION: I think the burden for reducing the risk is on the **company**, which is required (especially in the US legal system) to have a more and more "proactive" role

by providing more information in the consent form thereby increasing the **transparency** of the process. In particular, the user should be informed that - based on evidence law where "a brick is not a wall" - the interpretation of the information about future use collected in this way (movement of mouse etc.) might not be positively used by the company to inform its recruitment decisions: the reliability of interpreted information has always a margin of questionability that the user needs to be aware of (and the company alike).

2.4 HCI

RISK: From a human-computer interaction perspective, this case represents: (1) **asymmetric value**, where one side stands to gain more than the other side; (2) **asymmetric risk** where one side stands to lose more than the other side; and (3) **asymmetric knowledge** where one side knows what is going on while the other side does not. In a symbiotic relation the two sides are called symbiont and host. The symbiont is the company capturing and analyzing the information and the host is the user unwittingly providing the information. The alleged benefit for the candidate (the 'host') would be to have a chance to obtain the vacant position. Since giving data is mandatory for the application to be considered, it can be construed as forced by the company and of value to them. For the candidate, however, it is not clear that there is any direct benefit for giving out the information. Despite this asymmetry of value this still qualifies as a symbiotic interaction. In biology, symbiotic relations do not need to be mutually beneficial to the parties involved. One can be a parasite on the other.

SOLUTION: One solution is to **change the values in the asymmetry**. This could be achieved by having a policy according to which people have rights over personal information. This requires **defining personal information** and defining the difference between public and private domain. Personal information could be assumed to have a tier structure. When it is clear to the candidate what s/he is giving up – as determined by the tier level of private information – the company could be required to acknowledge that storing and analyzing that information is worth a certain amount. The exchange of information is then seen as a transaction which they have an obligation to pay for in some way. The candidate then would be in a position to make an informed decision whether to give up the information. If the company does not reveal the value or does not offer an appropriate amount in exchange then the candidate has grounds for a lawsuit, since whether s/he knows it or not there was a transaction and it violates the law or recognized policy concerning transactions in such cases.

A society might also decide to set a **cut-off** in such asymmetry, establishing for instance that no parasitic asymmetry is acceptable regardless of cost. That would become a law. Another thing to note is that in addition to the right to know a person might have intrinsic rights over future use. Since the value of a piece of information changes once it is aggregated with the information collected from several individuals it has one value in isolation and a rather greater value when part of the aggregation. If the information is now assigned a value as a function of the value of the aggregation

the asymmetry of value is potentially reversed. It might even prove to be so counterproductive to the individual company to purchase information for recruitment purposes since now they would have to pay many individuals, when all they wanted was to select a few candidates. On the other hand, if there is general knowledge to be gained from this sort of data capture then a new business might spring up that provides analytic tools to companies to help them made better hiring decisions. In that case, individual companies would no longer need to capture this data and store. They could make judgments about candidates using existing analyses (sold by the new company) and throw away the data they collect themselves right after they make their decision.

2.5 Psychology

RISK: What in this case strikes from a psychology perspective first of all is a deficiency in informed consent. The case also points to the question of the **timescale** of the use of data: companies collect tons and tons of data which today are not useful, but might turn out to have a high value in the future. Logically the user cannot provide informed consent for the use of their data for a scenario or use case which today is unknown.

SOLUTION: The only approach to solving this risk is to aim for maximum **transparency**: informing users of how their data might be used today (the basis of all informed consent), and also illustrating how it might be used in future. Where future uses of the data are markedly different from those envisaged at the time users consent to their data being collected, it is clear that the data collector should be required to obtain consent for the new use of the data. There are some interesting developments in this space right now, with a number of UK companies developing 'games' as part of selection and recruitment, where behaviors which are monitored during selection process game-play reveal underlying psychological characteristics of the applicant. Whilst these characteristics may be analyzed with the informed consent of the appli-cant today, it is easy to imagine that other relationships between the game-play behavior and other psychological characteristics are discovered in the future. Without seeking new informed consent from the applicant at this point, user's test results should not be analyzed for the newly discovered characteristics.

3 Case 2, Modeling the User: The Tailored Persuasive Message

The second case illustrates a scenario where a symbiotic system is used to Profiling for persuasive purposes. Indications about the realism of such a case are provided by [1, 4, 9], showing how current users' analytics can be exploited for persuasive or manipulative purposes.

"Users of a popular search engine log in to a set of free services connected to the search engine; once they are logged-in, the computer keeps track of all webpages that are visited from the browser where the log in occurs. The search engine stores a great amount of data about users and elaborates profiles about choices and preferences connected to the user and to similar users. In particular, it runs a free game showing several scenes and characters, recording the users positive or negative response to those characters. The search engine also runs and advertising service, which is used during a major election in one European city. The election candidate, who purchases the advertising service for the occasion, uses the stored preferences to automatically personalize the campaign ads and to make his/her portrait as a person looking as similar as possible to the individual elector. It also sends favorable voting forecasts collected from lay citizens whose profile matches the profile of the user, in terms of tastes, sports, family situation etc. Voters are positively impressed by such declarations and feel that the candidate is very close to them. The advertising firm revenue from this campaign is huge, but no money was paid to the search engine users for disclosing their information in the first place."

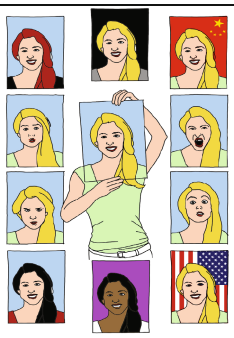

3.1 Ethics

RISK: Despite the similarities this case shares with the first one, it represents more serious threats for several reasons. In the first case the possibility that the system (of data collection and use) will be exploited for purposes that are negative for the user was questionable; in the present case such negative purposes are certain. Moreover, the negative consequences are far reaching and threaten the values at the very core of a democratic society – protection of human rights as well as undermining the democracy process. I would cluster two different groups of risks: 1. privacy violations and, 2. Human rights violations. For the first cluster, these violations would include: collection of data without informed consent, secondary use of data without consent, use of data for a purpose other than the one specified, possibility to de-anonymize, and lack of transparency. In the philosophical and ethical literature the definition of privacy used to focus on the control one had over one's data; deciding who has access, what it is used for, when it is used, and how much of it is used. Now, the definition of privacy has started to evolve into a concept in which the formation of one's **identity** is the central focus; being able to establish one's identity without having one created for you or becoming just another number. For the second cluster, the threat to human rights exists in the potential to manipulate emotions by targeting preferences and habits of users. This poses a threat to the values of **autonomy and dignity** of individuals. These serious negative consequences change the nature of the scenario; the potential infringement on human rights adds to the threat of the democratic process.

In deciding what to do about competing conceptions of the good life one may focus on the consequences of an action (the **consequentialist approach**) or the duties and principles on which the action is based (the **deontological approach**). Of course it is not so easy to isolate consequences from duties and many ethicists nowadays would go so far as to say that the line dividing one ethical theory from another are blurring; however, it is important in this case to be sure to point that in this case if one can attempts to justify the threat to human rights like autonomy and dignity (by undermining the democratic process) by saying that it could be a "good political candidate that is chosen", then the potential to engage in similar practices in the future with terrible outcomes (e.g. voting for a candidate or policy that is not good) becomes quite real.

SOLUTION: This is a difficult situation to find a solution for as it requires that companies and politicians are honest about their research practices even if it means they lose money. From an ethics perspective it is important to **empower** people to find their own voice to base their decisions on; this is why I would support education to allow people to inform their decision based on deeper knowledge of the issue. This education may come through the media or through an institution. As a solution to preventing these things from happening I would recommend establishing an **advocacy group** or ethics committees that work together to monitor and find solutions that make the symbiotic system process at stake ethical.

3.2 Security

RISK: From an information security point of view this case represents an example of **user profiling.** Currently, most Internet services and social networks collect data from users and try to profile them via the so-called user profiling in order to create target groups which they can target with other services (e.g., advertising) based on their characteristics. When they notify that they might sell such information to third parties, we are talking of buyers in the order of hundreds of companies what will get and use our data. While this information is commonly used in the aggregate form, it could also be exploited maliciously to try and define the profile of a specific user, his/her preferences, location and behavior. In addition the actual term of our decisions about such information is often ambiguous, as is the case when flagging as private an information on certain social media, which will never be private in the same way as we mean it to be.

SOLUTION: Increasing people's **awareness** of the risks involved in releasing data is nowadays necessary since security risks are chiefly underestimated. Such awareness and ability to **imagine possible consequences** should also be projected in the future possible use of the data which is light-mindedly disclosed now. To appreciate the importance of prospective thinking we can consider DNA information sent out to Internet services in exchange for knowing something about ourselves; this equals to disclosing core identity information, information that cannot be changed and that regards not just the individual releasing such information but all his/her relatives and progeny. It is easy to imagine that such information might be used in the future for genetic research applications or discriminations we are not even aware of nowadays and that we should me more jealous of such information than other we protect with much more alacrity.

Regarding possible technical solutions, there are several and they should be better known and more pervasively implemented (the anonymous internet browser system TOR, for example). Awareness, however, is the key solutions since it will not only convince people to prefer safer solutions but also to motivate users, regulators, technicians to ask for such solutions when they are not offered.

3.3 Law

RISK: From a legal point of view this case is about the indirect and unaware use of personal data. More technically, it is an **information security law case**, meaning a distinguished concept from concepts such as privacy and data security. [Many laws that purport to encourage cybersecurity are, in fact, designed with a focus on protecting privacy or encouraging data security]. Unlike privacy and data security, **cybersecurity** is focused not only on the information, but the entire system and network. For this reason, laws that focus only on privacy and data security may not consider all factors necessary to promote cybersecurity.

SOLUTION: **Transparency** and adoption of **best practices**.

On the one hand, the popular search engine where users log in to set free services should inform users of future potential uses of their personal data through the connected advertising service.

On the other hand, there are also some best practices for users that can help to reduce the risk of violation of privacy and unforeseeable use of data:

1. Anonymization: don't collect personal data if you don't need it. Work with anonymous or de-identified data if you can.
2. Disclose only the data you need and required, especially try to minimize the disclosure of sensitive personal data.
3. Encrypt sensitive personal data during transit
4. Check your contract with customers to ensure that you are not agreeing to unreasonable security practice in place.

Also we must consider that different kinds of information would have a different consequence in legal terms: the legal "weight" of human values and rights (e.g. religion; sexual orientation etc.) is different from the legal "weight" of choices and preferences (e.g. kind of theatre preferences; food preferences; sympathy or not for domestic animals etc.). The unaware storage of consumers' information about those different aspects (human rights and human preferences) has, obviously a different legal protection in case of breach. Case number 2 is important because it lets stakeholders thinking about these differences.

3.4 HCI

RISK: I would frame this case similarly to the first case, namely as one of **asymmetry in risk, value and knowledge**. Unlike that case, though, this one adds a complication in terms of human-computer interaction, namely that there is a third party involved, and

this party is not the one directly involved in the symbiotic relation. That is, the politician purchasing the service is the third party that benefits from the personal data. Presumably there was a transaction between the symbiont – the party gathering the information and the one who should have paid for it – and the politician.

SOLUTION: The system that manages collection and all transactions – the symbiont and its owners – has an obligation to reveal the value of the information that is being collected. This value concerns the transaction occurring between the system and the user disclosing data but especially between the system owners and the customer using the data collected. One thing is to use the collected information internally and another is to sell it. Can this information be sent to any country whatsoever? For any purpose? Its value depends on what others pay for it. So regulation should shift focus, in this case, from the transaction with the system users to the transactions between the system owner and its customers who buy the collected data. Therefore, the solution would be to develop policy to **regulate transactions whenever collected data are sold**.

Purely from a HCI perspective, transparency can be increased by **re-design**, improving the comprehensibility by clever visualization of the meaning of terms of usage that otherwise is specified in 20 pages of text. The owner of the original information – the information host – ought to also have the right to expect to **re-negotiate** or renew the terms of the transaction once it is apparent the value of the information has changed. Information has a continuously changing value, it is not used once and once forever. Once new value becomes apparent there might be appropriate conditions to ask for a renegotiation. At the same time, acknowledging that some information might not be valuable once it is collected but prove valuable afterwards, it might be foreseen that its value should be paid via **fee-for-use**, namely only when users' data is actually used for some profit or benefit to the system owner.

3.5 Psychology

RISK: I think that in psychological terms this should be understood in terms of **identity and individual differences**, but I wonder how this case - being enabled by a symbiotic technology - differs from typical political campaigning, where politicians commonly adapt their propositions to the audience or the interlocutor.

SOLUTION: In this case, as in the first one, what is critical is to increase transparency. I think that **transparency** in transactions - even commercial ones - should be a leading principle, allowing a better understanding of the value of the transaction. Even more so in a political domain, where democracy is at stake and there is arguably a need for even higher standards of fairness in transactions. And to achieve such transparency, I think that **education** of the citizen and the consumer have a huge role to play (transparency being not possibly in the absence of a user understanding what data a system is acquiring). Also relevant here is that there is a natural pressure for systems to be easy to use (to maximize likelihood of engagement) and it is likely that the user will not realize the importance of transparency of a system until s/he experiences some negative consequence of a lack of transparency.

4 Case 3, Deciding for the User: The AR-Guided Surgical Intervention

The third case illustrates a case where a symbiotic system might lead to the user loosing agency due to automatic responses directly provided by the system. Indications about the realism of such a case are provided by [3, 7, 10].

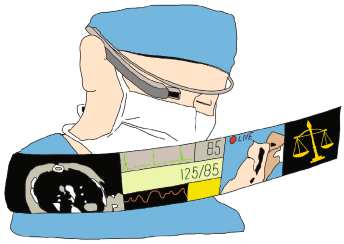

"A hospital has adopted a laparoscopic surgery equipment which is connected to the apparatus acquiring vital signals from the patient under surgery. During the surgical procedure, the physician wears a pair of augmented reality eyeglasses to receive information about vital signals along with data about success probability for each procedure applied to the patient during the surgery, in order to better inform his/her ongoing decisions. This data is recorded for archival reasons by the hospital and can be used in case of a lawsuit against the hospital after surgeries that do not succeed. It is the hospital policy to avoid conflicts between the choice of the physician during the surgical procedure and the data displayed by the machine, the evidence of the patient status recorded by the equipment. Thus in case of the patient reaching a critical condition, the physician is recommended to only rely on the standard procedure and quit any other attempts which is less likely to succeed. Therefore, to a certain extent, the decisions are embedded in the intelligence of the computer elaborating data and making recommendations. Somehow a moral decision is incorporated into the machine."

4.1 Ethics

RISK: To me this represents a prototypical symbiotic relation since parts of the decisional powers are externalized by the physician to the machine based on data it receives. It is a whole system including patients' data, physician's decision and systems' elaboration/recommendations. The most important part of this scenario is the fact that the surgeon is responsible for deciding whether or not he/she will take the advice of the machine. The moment this changes and the surgeon must do what the machine tells it (whether this is an explicit formal policy at the hospital or an implicit one) the scenario changes and the relation is no longer ethically acceptable or desirable. The second scenario is one in which the surgeon's freedom to choose has been limited. This limitation threatens the professionalization of medicine as one would have to be concerned about who is taking decisions and who is liable in case of problems, i.e. if the surgeon does what the machine tells it to do then will the machine be liable if someone goes wrong? Further, will we sue or fire the machine for damages? But responsibility from an ethics point of view is more than liability; it requires a moral

agent with intentions who is able to reason through the consequences, understanding the consequences of an action; therefore this cannot be delegated to a machine.

Another potential risk is **deskilling** if the surgeon learns to rely on the technology and its elaboration more than s/he does on his/her own judgment. Moreover, with the use of this new technology the surgeon may not have an instance in which he/she is able to train using conventional methods.

SOLUTION: Using military terms, the surgeon must remain *in* the loop instead of being put *on* the loop. This means that the surgeon should be in control of giving commands and making choices. Part of the **training** with the systems should be to understand how the machine reasons and how to manage disagreements with the machine especially during emergencies. Another part of the training should be to make sure surgeons know how to perform the surgery if/when the machine breaks and the surgeon must rely on his/her own skills without the technology [13].

4.2 Security

RISK: From information security perspective this scenario points on the one hand at the aspects related to data storage in hospitals, and on the other at the possibility that the system is programmed to make decisions, be controlled remotely or being hacked for **malevolent purposes**.

SOLUTION: Solutions from information security perspective for this scenario are common good practices for storage of confidential information and to make computer systems secure: avoiding unauthorized parties to take control of such systems.

4.3 Law

RISK: This case regards the use of extra-clinical tools to support clinical decisions.

In general the gradual shift towards the use of extra-clinical tools to supplement the informed consent process and support clinical decisionsì could present the risk to consider the tool not simply as a decision-support tool but a decision-replacement tool (instead of the patient-physician's decision).

A second important risk is of blurring the principal role of the physician and interfering with his/her **freedom of therapeutic choice** whose responsibility (not simply legal liability) is shared with the patient.

Physicians have long faced tort liability for breach of informed consent if a patient is harmed as a result of the physician's failure to provide the information needed to make an informed medical decision. However, with increased reliance on extra-clinical tools, informed consent mechanisms incur an increased risk of malpractice liability.

The physician who simply relies on eyeglasses without reasoning on the bases of his/her knowledge, under even the most traditional tort principles, will be liable for malpractice. (Failure to engage fully in the informed consent process, even if decision support tools are made available, is a clear breach of the standard of care).

What if the pair of augmented reality eyeglasses gave wrong information? Similar problems where already present in the field of medical guidelines application.

Several scenarios could be traced:

1. the physician follows the eyeglasses indications and is personally persuaded by this choice;
2. the physician follows the eyeglasses indications but personally would have made another choice;
3. the physician does not follow the eyeglasses indications because on the base of his/her knowledge would have made another choice and indeed decides to make such choice.

SOLUTION: It is clear that medical providers who prescribe or use decision support tools may face tort liability if they misuse the tools or provide negligent counseling. This is a simple and relatively uncontroversial expansion of traditional malpractice liability. But the use of decision support tools also poses a secondary problem - namely, that patients may be harmed if the decision aids they use are faulty, misleading, or biased. If the regulatory or certification process aimed at ensuring the quality of decision aids fails, injured patients will look to tort law to provide a remedy. And since current tort doctrine makes it extremely difficult for such claims to succeed, it is time for policymakers and legal scholars to evaluate the costs and benefits of expanding tort liability in these cases.

The risk could be minimized with **a good training and instructions** by the producer on real opportunities offered by the eyeglasses. All information about the real help technology will offer to the patient should be exactly represented to the patient before the surgery in order to share "the potential scenarios".

Apart from this, it is still an open legal question in this field to determine whose responsibility it is to minimize this risk. Producers' will have an important role as well as physicians in transfering to the patient the useful information and sharing potential scenarios prior of the intervention (learned intermediary hand role).

Within the personalized medicine era, these eyeglasses have to be seen as a functional instrument of help in critical situation.

The risk of restricting physician's freedom of choice is inherent and is not avoidable. Perhaps, every physician will have to be aware that the liability for the final decision is due to his/her own choice, so it would be important for him/her to know from the producer the risk of error margin of the high tech product.

It has to be underlined that, because transitioning elements of the consent process into extra-clinical arenas is a dramatic change in the practices of protection medical freedom of choice and informed consent, it necessitates a new kind of conversation about liability. First, although **product liability** law sometimes subjects creators of faulty products to strict liability (that is, liability regardless of fault), decision support tools do not fall within the legal **definition of a "product"** and so are not subject to strict liability. Their **inherent autonomy** is currently under analyses.

Second, in the future it will be crucial to re-analyze issues of **vicarious liability** of the hospital and other involved subjects.

It should be also underlined that if a hospital system requires physicians to use decision aids for particular conditions it will also have a role in the allocation of liabilities, but this element will not be a "safe harbour" for physician who decided.

4.4 HCI

RISK: This case shows that the parties in the symbiotic relation have roles that depend not only on their **knowledge** but also on **external** practices such as the legal attribution of responsibility. The same case appeared years ago concerning expert systems for blood diagnosis, which were about 95% as good as a doctor on a good day. That is way better than most doctors on most days. Yet still hospitals ended up not using them because of the risk of legal suits. In normal cases gross failure leads to a law suit of the doctor. But who do you sue when it is an expert system? And what are the standards that one applies? The risk was that the responsibility for imperfection would have been laid at the foot of the programmer. And that risk might be too high given that the same program would be used in many places. The trouble is that when you think like this you give up reliable expertise (the expert system) to defend a general principle of morality or law. And yet the system is often the best way to proceed.

SOLUTION: Responsibility in this case should be allocated as a function of accountability and ultimately of knowledge. But we want humans in the loop. For instance, if an expert system left the final decision to the physician but also had a facility that would allow the physician to: (a) **delegate the decision** to the system, on a case-by-case basis; or (b) ask the system for its reasons for its suggestion or decisions and to take issue when the reasons are not clear enough, then we manage to keep humans in the loop. The final decision now lies with the physician. And there is the same mechanism used among teams of humans – they talk it out by asking each other for their reasons. The system and doctor now would be a learning team.

4.5 Psychology

RISK: I agree with the other panelists' responses, this case reduces the surgeon's **autonomy** and decreases the surgeon's skills. This reminds of the same issue currently at stake with self-driving cars, where the driver must be able to deactivate the cars' automatic behavior to get in control of the situation.

SOLUTIONS: The system should be transparent, **explaining** itself and then allowing the surgeon to make decisions including the decision to delegate decisions.

5 Synthesis and Conclusions

As mentioned in the introduction, and illustrated in the three cases and related discussion, symbiotic systems bring about a set of information transactions and intervene on the human beings and the environment in ways that are not inherently good or bad, but that at any rate mobilize some interests and values. Symbiotic systems, by which we mean here systems that have the three characteristics mentioned in the introduction (i.e., detecting users' implicit information, modeling the user's state based on such information and making decisions based on such model) can represent a resource to empower the human agent but can also end up serving private goals, manipulate and - ultimately - alienate human beings from their rights. In particular, the risks - as they have been identified during the panel - are of information leakage and malicious user

profiling; deceitful use of data and threat to information security; decrease in self-determination, privacy and dignity; deskilling; asymmetry in value, risks and knowledge. The actual position that a given symbiotic system holds in the continuum between beneficial and harmful applications, between protecting the many and serving the few, depends on the **society's ability to set criteria of acceptability for their design and usage and to compel the application of such criteria**. Given the scale and the subdued level at which symbiotic systems operate, it seems urgent that society decides what is a legitimate symbiotic system and finds mechanisms to secure the respect of such decisions. The discussion during the panel highlighted some of these criteria and provisions, which are tentatively summarized in Fig. 2.

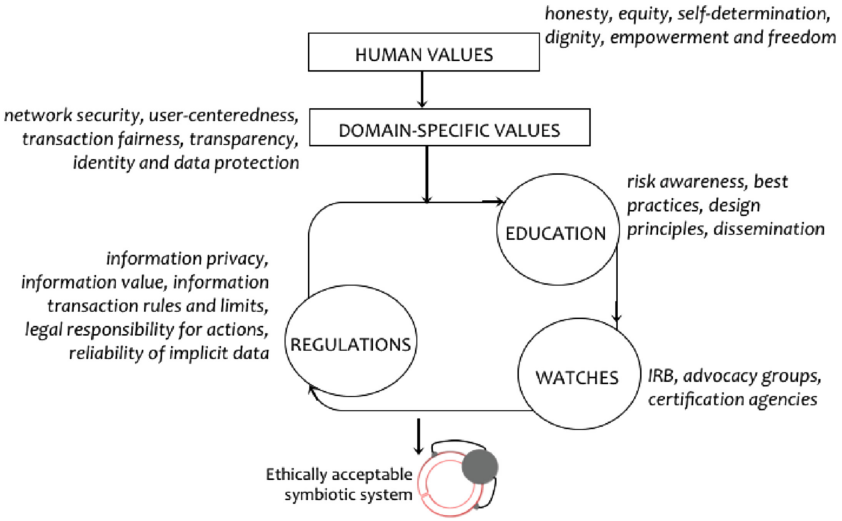

Fig. 2. A framework to collectively take charge of ethically acceptable symbiotic systems

Both the ultimate goals of the system and the ways in which it operates must be legitimate, making sure that it directly pursues, respects or at least does not compromise **society values** such as honesty, equity, self-determination, dignity and freedom. At a less abstract level, these overarching ethic values are achieved by way of **domain-specific values** that decline them into more concrete objectives, such as network security, user-centeredness, transaction fairness, transparency, identity and data protection. These objectives are achieved by way of education, regulation and agents. **Education** is the process to understand the risks of such systems and to learn risk minimization procedures (such as informed consents, enabling of own interfaces to input data in third party's systems, reporting lack of comprehension, design principles). **Regulation** includes the definition of what is a private and public information, how and when the value of a given piece of information is to be (re)defined, what transactions are legitimate, who are the beneficiaries, who is responsible for a decision based on symbiotic systems, and what is the legal value of information derived from implicit

data. **Watches** are agencies such as institutional review boards or advocacy groups that specifically monitor ethical risks, certify procedures and represent users' rights by voicing them publicly and promoting regulations and education initiatives.

In conclusion, ethically acceptable symbiotic systems are the results of a collective effort, where the burden of dealing with the risks of releasing personal information is shared with all relevant stakeholders and publicly debated.

Acknowledgments. The present panel was partially funded by the EU project MindSee (grant agreement n. 611570). Individual panelists are responsible for their interventions as they are written here, each one for his/her own area of expertise; the panel organizer is responsible for the introduction and conclusions. We are grateful to Piero Turra for the illustrations of the three cases.

References

1. Acar, G., Eubank, C., Englehardt, S., Juarez, M., Narayanan, A., Diaz, C.: The web never forgets: Persistent tracking mechanisms in the wild. In: Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, pp. 674–689. ACM (2014)
2. Barfield, W.: The future to merge with AI machines. In: Cyber-Humans, pp. 267–284. Springer, Cham (2015)
3. Baum, S.: Advancing smart glasses as augmented reality tool for surgeons drives Vital Medicals' capital raise. MedCityNews (2015). http://medcitynews.com/2015/02/advancing-smart-glasses-augmented-reality-tool-surgeons-drives-vitals-medicals-capital-raise/
4. Bessi, A., Petroni, F., Del Vicario, M., Zollo, F., Anagnostopoulos, A., Scala, A., Zollo, F., Quattrociocchi, W.: Viral misinformation: the role of homophily and polarization. In: Proceedings of the 24th International Conference on World Wide Web, pp. 355–356. ACM (2015)
5. Brandom, R.: Someone's trying to gut America's strongest biometric privacy law. The verge (2016). http://www.theverge.com/2016/5/27/11794512/facial-recognition-law-illinois-facebook-google-snapchat
6. Burgoon, J.K., Blair, J.P., Qin, T., Nunamaker, J.F.: Detecting deception through linguistic analysis. In: Chen, H., Miranda, R., Zeng, D.D., Demchak, C., Schroeder, J., Madhusudan, T. (eds.) ISI 2003. LNCS, vol. 2665, pp. 91–101. Springer, Heidelberg (2003). doi:10.1007/3-540-44853-5_7
7. Doryab, A., Bardram, J.E.: Designing activity-aware recommender systems for operating rooms. In: Proceedings of the 2011 Workshop on Context-awareness in Retrieval and Recommendation, pp. 43–46. ACM (2011)
8. Granhag, P.A., Vrij, A., Verschuere, B.: Detecting Deception: Current Challenges and Cognitive Approaches. Wiley, Chichester (2015)
9. Kramer, A.D., Guillory, J.E., Hancock, J.T.: Experimental evidence of massive-scale emotional contagion through social networks. Proc. Nat. Acad. Sci. **111**(24), 8788–8790 (2014)
10. Musen, M.A., Middleton, B., Greenes, R.A.: Clinical decision-support systems. In: Shortliffe, E.H., Cimino, J.J. (eds.) Biomedical Informatics, pp. 643–674. Springer, London (2014)

11. O'Brolcháin, F., Jacquemard, T., Monaghan, D., O'Connor, N., Novitzky, P., Gordijn, B.: The convergence of virtual reality and social networks: threats to privacy and autonomy. Sci. Eng. Ethics **22**(1), 1–29 (2016)
12. Sartori, G., Orru, G., Monaro, M.: Detecting deception through kinematic analysis of hand movement. Int. J. Psychophysiology **108**, 16 (2016)
13. van Wynsberghe, A., Gastmans, C.: Telesurgery: an ethical appraisal. J. Med. Ethics **34**(10), e22 (2008)
14. Won, A.S., Perone, B., Friend, M., Bailenson, J.N.: Identifying anxiety through tracked head movements in a virtual classroom. Cyberpsychology Behav. Soc. Networking **19**(6), 380–387 (2016)