



A Modular System for Enhanced Robustness of Multimedia Understanding Networks via Deep Parametric Estimation

Francesco Barbato*
francesco.barbato@unipd.it
University of Padova
Samsung Research UK
Italy / United Kingdom

Umberto Michieli
u.michieli@samsung.com
Samsung Research UK
United Kingdom

Mehmet Kerim Yucel
m.yucel@samsung.com
Samsung Research UK
United Kingdom

Pietro Zanuttigh
zanuttigh@dei.unipd.it
University of Padova
Italy

Mete Ozay
m.ozay@samsung.com
Samsung Research UK
United Kingdom

ABSTRACT

Performance degradation caused by corrupted multimedia samples is a critical challenge for machine learning models. Previously, three groups of approaches have been proposed to tackle this issue: i) enhancer and denoiser modules to improve the quality of the noisy data, ii) data augmentation approaches, and iii) domain adaptation strategies. All have drawbacks limiting applicability; the first requires paired clean-corrupted data for training and has a high computational cost, while the others can only be used on the same task they were trained on. In this paper, we propose SyMPIE to solve these shortcomings, designing a small, modular, and efficient system to enhance input data for robust downstream multimedia understanding with minimal computational cost. Our SyMPIE is pre-trained on an upstream task/network that should not match the downstream ones and does not need paired clean-corrupted samples. Our key insight is that most input corruptions found in real-world tasks can be modeled through global operations on color channels of images or spatial filters with small kernels. We validate our approach on multiple datasets and tasks, such as image classification (on ImageNetC, ImageNetC-Bar, VizWiz, and a newly proposed mixed corruption benchmark named ImageNetC-mixed) and semantic segmentation (on Cityscapes, ACDC, and DarkZurich) with consistent improvements of about 5% relative accuracy gain across the board¹.

CCS CONCEPTS

• **Computing methodologies** → **Scene understanding; Vision for robotics; Image segmentation; Reconstruction.**

*This work has been produced during an internship at Samsung Research UK.

¹Our codebase and the new ImageNetC-mixed benchmark are available at <https://github.com/SamsungLabs/SyMPIE>.



This work is licensed under a Creative Commons Attribution International 4.0 License.

MMSys '24, April 15–18, 2024, Bari, Italy

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0412-3/24/04

<https://doi.org/10.1145/3625468.3647623>

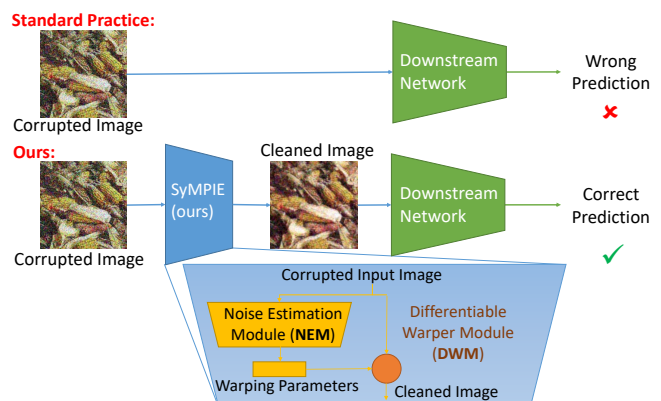


Figure 1: An illustration of our modular system (SyMPIE) for efficient image enhancement targeting increased model robustness to corruptions in different multimedia tasks. SyMPIE contains two modules, namely, NEM and DWM.

KEYWORDS

Model Robustness, Content Enhancement, Denoising, Image Classification, Image Segmentation

ACM Reference Format:

Francesco Barbato, Umberto Michieli, Mehmet Kerim Yucel, Pietro Zanuttigh, and Mete Ozay. 2024. A Modular System for Enhanced Robustness of Multimedia Understanding Networks via Deep Parametric Estimation. In *ACM Multimedia Systems Conference 2024 (MMSys '24)*, April 15–18, 2024, Bari, Italy. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3625468.3647623>

1 INTRODUCTION

Deep learning models have been widely used in several multimedia systems and applications. Nonetheless, recent studies have explored their robustness to corrupted input data [21, 22], showing critical performance degradation affecting the quality of experience of many downstream applications² such as extended reality [2] and video streaming [62]. Indeed, the data-hungry nature of deep neural

²We use *upstream task* to identify the one used to pre-train our system (e.g., classification) and *downstream task* to identify the deployment one.

network (DNN) models, as well as the ever-growing complexity of network architectures, make the generated models sensitive to even mild distribution shifts causing a severe degradation in performance [46, 47, 55]. These eventualities are often naturally met in many real-world applications, where data may unavoidably encounter natural alterations or corruptions [22, 38], such as sensor degradation (e.g., *shot noise*, *defocus blur*), compression and packet loss artifacts (e.g., *block artifacts*, *corrupted colors*), or framing issues (e.g., *low contrast*, *rotations*, *images too bright*), to mention a few.

Considering the increasing adoption of deep models, this issue has become of paramount relevance. Therefore, a new research field emerged to attempt to make models more robust under different perspectives. Existing approaches to improve model robustness can be categorized into three main branches: (i) data augmentation approaches, (ii) test-time adaptation approaches, and (iii) enhancing and denoising approaches, such as autoencoder-based, generative, and adversarial denoisers (e.g., GANs, diffusion models, etc.).

Data augmentation techniques [6, 24, 67] either re-train AI models from scratch or fine-tune them by applying a large set of general-purpose augmentations (often synthetic) that mimic common corruptions encountered in the real world on data. In other words, this paradigm aims to produce a model that is robust against corrupted images and that can be very effective in improving model abilities to generalize under data distribution shift. The main drawbacks of this approach are that the trained model can only perform the task that has been trained with supervision on; the trained model tends to perform well on images with training-like distortions; the types of augmentations should be decided *a priori*, and they cannot be easily tuned to a scenario with variable corruptions (unless via expensive retraining).

To improve model robustness to variable-type corruptions at inference time compared to training time, Test-Time Adaptation (TTA) methods have been proposed [27, 28, 69]. TTA methods focus on resolving data distribution shifts directly at test-time via dynamic updates of pre-trained models based on the specific characteristics of the target test data [15, 44, 59]. TTA methods enable a model to improve its performance when encountering variations, unseen examples, or changing conditions at test time. Unlike data augmentation techniques, TTA methods do not involve expensive training procedures, but they are still bounded by the same task to be used during the pre-training and testing/adaptation phases.

Other popular techniques are enhancement and denoising, *i.e.*, the process of improving input samples before they are fed to the downstream network [56]. The most promising techniques are nowadays based on deep learning models. For instance, autoencoders are used to encode compact representations of input samples and then decode them to minimize a reconstruction loss with respect to ground truth samples [1, 34, 68]. Generative and/or adversarial models (e.g., diffusion models [11, 58, 65] or GANs [8, 31, 42, 43, 60]) modify input samples by iterative processes or style transfer techniques. The main drawbacks of these methods are the need for expensive inference time and for paired clean-corrupted samples at training time while being able to handle in-domain corruptions only (*i.e.*, particular corruption types seen at training time).

Overall, current techniques all come with some key limitations that we briefly discussed. In this paper, we propose a System for

Modular Parametric Image Enhancement (SyMPIE), to overcome some of these limitations. SyMPIE implements a hybrid strategy combining the best of data augmentation and enhancement approaches while being computationally efficient and fully compatible with any model and any downstream task.

We propose to tackle the issue from a completely different direction, designing a modular system that predicts parameters used by an ad-hoc module to enhance the received input (see Fig. 1). We utilize generic data augmentation pipelines and common upstream task/network (e.g., image classification with ResNet50) to pre-train our model (System for Modular Parametric Image Enhancement, SyMPIE). Our key insight is that a large portion of input corruptions found in the real world can be modeled by applying either global operations on color channels (for example, a night scene can be approximated by a darkened daytime scene), or spatial noise which can be filtered out by a fast convolution operation with small kernel size (for example, a Laplacian filter or a blurring operation can be used to sharpen an image or to reduce noise on the image). Additionally, our model is pre-trained minimizing a standard classification task loss, alleviating the need for paired clean-corrupted samples during pre-training. In other words, the downstream task and model do not necessarily need to match the upstream task and model used for pre-training our model. For example, we may train our model on a corrupted ImageNet [23, 48] benchmark for image classification using a Convolutional Neural Network (CNN), and deploy it on semantic segmentation benchmarks [51] with adverse conditions to support a transformer architecture.

Therefore, our SyMPIE is similar to enhancement and denoising approaches, with the key advantage of containing much smaller modules that are pre-trained without paired clean and corrupted ground truth.

Our contributions can be summarized as follows.

- We propose a lightweight modular image enhancement system, named SyMPIE, which predicts parameters of *ad-hoc* operators that are applied to input samples to improve the content understanding.
- In the experimental analyses, SyMPIE consistently improves the accuracy of downstream models for the corresponding task with minimal impact on the number of compute operations (about 2GFLOPs on top of 343GFLOPs of a classical ResNet50 architecture for a full HD resolution input) and can be seamlessly applied on top of any competing approach improving their accuracy (up to 8% relative gain) without re-training. Unlike data augmentation methods, SyMPIE does not require long training procedures.
- Unlike existing denoising approaches, SyMPIE is fast, does not require paired clean-corrupted samples during pre-training, and can be reused in multiple setups without losing its efficacy. In particular, SyMPIE proves to be fully compatible with any convolutional and transformer-based architecture, including recent foundation models (e.g., CLIP). SyMPIE enhances the accuracy of the aforementioned models on several downstream tasks (e.g., image classification, semantic segmentation) in the presence of corrupted test data improving the quality of experience of the final end users.

2 RELATED WORKS

Robust generalization. The robustness of computer vision models against distribution shifts is of vital importance in multimedia applications. Starting with the discovery of adversarial examples [54], many different robustness venues have been explored by the community, such as common image corruptions [23], images with conflicting shapes/textures [14], and style variations [32]. The core idea behind robust generalization is to make multimedia models increasingly invariant to such shifts; *e.g.* a model trained on a training set should *generalize* well to samples with unseen styles, corruptions and perturbations. Arguably, the most common effects found on real data are image corruptions, which have been standardized and categorized in special benchmark datasets [23] for the evaluation of deep learning architectures. This is the scenario we focus on in our paper, due to the practical importance. Note that we consider a wide variety of possible corruptions, while the (*Video*) *Coding for Machines* field focuses only on compression-related artifacts where very interesting results can be achieved [12, 17].

There are various methods that aim to provide robustness to computer vision models. They can be categorized into three groups; i) data augmentation methods, ii) test-time adaptation methods, and iii) preprocessing (*e.g.*, denoising and enhancement) methods. **Data augmentation** methods are built on a simple premise; one can achieve robustness by retraining or fine-tuning a model on a training set updated with the samples that the original model failed on (*e.g.*, unseen distributions). Adversarial training [36] has been one of the most prominent examples of such methods, where adversarially corrupted examples are included in the training set during retraining of the model. Since then, the same core idea has been used by a plethora of methods that addressed robustness from a data augmentation perspective. AugMix and its variants [24, 61] proposed to use randomly sampled augmentations in a cascaded manner to diversify the training distribution. DeepAugment proposed to use image translation models to produce new training images [20]. Learned augmentation policies tailored for improved robust generalization are proposed in [10]. The use of fractals and feature maps to create new samples [25], frequency-specific expert model ensembles [49], max-entropy augmentations [41] and frequency-spectrum perturbations [6, 67] are some examples of data augmentation methods which have proven themselves successful against combating the issues caused by distribution shifts.

Test-time adaptation. Although they currently hold the state-of-the-art on many benchmarks, data augmentation methods have a key disadvantage, where they require retraining or finetuning every time they encounter (and fail against) new distributions, which may or may not be unknown in practical scenarios. Assuming one can perform these expensive updates periodically, model capacity issues as well as catastrophic forgetting [30] are likely to be new accuracy bottlenecks, which will inevitably lead to practical issues in deployment scenarios. In order to avoid such problems, another branch of methods focuses on analyzing input images to detect if they are from *unseen* distribution. Using network prediction consistency [13, 35], specialized detector models [16] and sample statistics [63] have proven to be successful for adversarial examples, and similar ideas have been extended to common image corruptions [15, 27, 28, 44, 59, 69].

Going further than just detecting critical samples, more recent methods aim to improve robustness in test-time. Tent [59] minimizes the test-time model prediction entropy via optimizing channel-wise affine parameters per-batch. MixNorm [27] points out the inherent assumption of having samples from a single distribution in a batch, and its detriment on models with batch-norm layers. They then propose MixNorm layers, a replacement for batch-norm layers during test-time that adapts to query image statistics by using novel views of the same image. MEMO [69] minimizes the entropy of predictions of the input sample point and its augmented views, which improves robustness even with a single image. NOTE [15] proposes instance-aware batch-norm layers to improve robustness with a single image without the overhead of additional forward passes (*e.g.*, for augmented views of the input sample) required for other methods. SAR [44] proposes a new entropy minimization method that takes into account samples with large gradients and avoids collapsing to trivial solutions, thus improving overall robustness performance.

Preprocessing. Test-time adaptation methods largely alleviate the issues inherent to data augmentation approaches; they perform partial updates at most, therefore they are fast and cheap. However, they still require a form of training, which requires periodic updates in deployment scenarios where data shift is ever-present.

Another branch of methods that aims to achieve robust generalization can be roughly categorized as preprocessing methods. These approaches do not tackle the problem as a normalization issue, unlike test-time adaptation methods. Instead, they attempt to bridge the distribution gap between the query sample and the model training set, by trying to *enhance* and *denoise* the query image. Some naive preprocessing strategies such as label smoothing, spatial smoothing, and total-variance minimization are shown to help improve robustness [63, 66]. Thinking of the corruption over the image as noise, conventional auto-encoding mechanisms can be used to purge the input image of its corruption-induced elements [1]. Essentially, these methods aim to learn an approximation of a transformation over the query image, which makes the underlying issue an image-to-image translation problem. There are many image translation methods using GANs [8, 31, 43, 60] or more powerful diffusion models [11, 58] that can be either used off-the-shelf or updated to work well with our scenario. Although these methods do not necessarily update the actual downstream task model, they still require per-distribution updates for the model, which makes it expensive. Furthermore, the enhancement or denoising model adds an expensive step that might not be feasible for certain deployment scenarios.

The proposed approach aims to melt the best characteristics of these methods in a single pot: high accuracy, fast execution time, cross-distribution and cross-task generalization, and removing the need for paired data for training.

3 A MODULAR SYSTEM FOR IMPROVING MODEL ROBUSTNESS

In this section, we introduce the main contributions of this work. We detail our two-module network architecture in Sec. 3.1. Its training procedure exploits exponentially smoothed versions of our

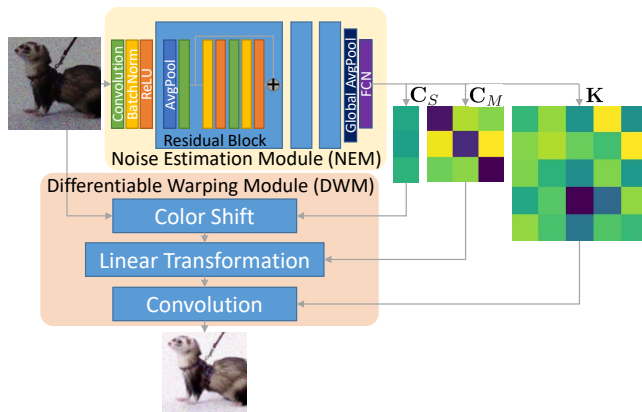


Figure 2: A detailed scheme of our modules working together to enhance the content of an image. The Noise Estimation Module (NEM) receives a corrupted input and predicts a triple of parameters (C_S , C_M , K). These parameters are used by the Differentiable Warping Module (DWM) to enhance the image using parametric operators.

modules for improved generalization and is discussed in Sec. 3.2. Finally, the inference process is described in Sec. 3.3.

3.1 A System for Modular Parametric Image Enhancement (SyMPIE)

Our System for Modular Parametric Image Enhancement (SyMPIE) is made of two modules. The first module (Sec. 3.1.1) is the Noise Estimation Module (NEM), E , which is implemented by a small CNN to estimate a set of parameters used by the second module from training samples. The second module (Sec. 3.1.2) is the Differentiable Warping Module (DWM), D , and is used to process input samples in order to remove distortions from them. To obtain the final SyMPIE architecture, we combine the two modules in a single block by $U = D \circ E$ where \circ denotes the module composition.

A detailed overview of our modules is shown in Fig. 2: SyMPIE can be inserted between the input data and any generic downstream model. Three key features distinguish our approach from existing image enhancement approaches, namely: (i) the removal of paired clean-corrupted images as a requirement for training; (ii) the prediction of enhancing parameters rather than of the cleaned images directly; (iii) the ability to generalize across several upstream and downstream tasks and networks.

Existing works, indeed, train enhancer models as a regression task on the clean images, effectively optimizing the reconstruction PSNR (Peak Signal-to-Noise Ratio). We tackle the issue from a completely different point of view, focusing on efficiency and downstream accuracy, instead. Incidentally, this allows our model to be trained and fine-tuned on any available dataset, especially on real-world datasets that contain naturally corrupted images with no clean counterpart ground truth.

3.1.1 Noise Estimation Module (NEM). To estimate the parameters needed for enhancement, we designed a small residual CNN, E ,

which estimates explicitly the set of parameters used by the differential warping module D . Motivated by the observation that most of the distortions of the input samples affect the color space only [37, 40, 70], we identify a set of filters and linear transformations to be applied to the input samples for their improvement.

In our setup, we aim at estimating the parameters of i) a filter kernel $\mathbf{K} \in \mathbb{R}^{K \times K}$ (where kernel size K used in the experiments is given in Sec. 4.1); ii) a linear transformation of color channels, represented by the matrix $C_M \in \mathbb{R}^{d \times d}$; and iii) a global color shift $C_S \in \mathbb{R}^d$ where d is the number of input channels (e.g., $d = 3$ for RGB images). This set of parameters allows our architecture to model several naturally occurring corruptions, such as under/over-exposed images, sensor noise (e.g., Gaussian, impulse, shot noise), unbalanced white point, etc. without the need to employ expensive deep learning models for enhancement.

To estimate the above parameters and keep the additional footprint minimal, we implement the NEM as a small residual CNN (see the top half of Fig. 2). This allows the NEM module to efficiently extract global information from an input image, and predict suitable parameters for the subsequent module. More in detail, $E : \mathcal{X} \ni \mathbf{X} \mapsto (\mathbf{K}, C_M, C_S) \in \mathbb{R}^A$, where \mathbf{X} is a corrupted input, \mathcal{X} denotes the space of input images, and (\mathbf{K}, C_M, C_S) are the parameters used for the cleaning process. More implementation details are given in Sec. 4.1.

Moreover, in the vanilla deep learning-based image enhancement formulation [1, 34], the network weights used to process the images are fixed. Therefore, the models learn to partition their own parameter spaces to deal with different situations, reducing the overall efficacy and parameter efficiency. Our model, instead, explicitly predicts the parameters depending on the input, therefore, our model can seamlessly adapt to multiple types of corruptions without the need to partition its parameter space.

3.1.2 Differentiable Warping Module (DWM). To enhance input images using predictions obtained from the NEM, we use our Differentiable Warping Module (DWM), D , depicted in the bottom half of Fig. 2. The module receives as input a corrupted image and the parameters estimated by the NEM and applies the parametric operators to the former to enhance the image. Formally, we have $D : (\mathbf{X}, (\mathbf{K}, C_M, C_S)) \mapsto \bar{\mathbf{X}}$, where $\bar{\mathbf{X}} \in \mathcal{X}$ is the improved image.

The operations of D are completely differentiable, allowing back-propagation of gradients through the DWM, and optimizing the predictions of the NEM in an unsupervised manner. This enables our method to be trained without any paired data, since gradients propagated through the DWM come directly from predictions over a frozen upstream model. In other words, the NEM is trained to enhance samples in order to maximize the performance of a frozen upstream network. This is a great advantage compared to existing works which either cannot enhance the inputs to improve the performance of downstream models (e.g., diffusion models) or cannot effectively leverage the gradient flow to produce cleaner inputs (e.g., adversarial approaches). Furthermore, when modifying input samples directly, several approaches [39] introduce distortions that appear random and unnoticeable to the human eye but can completely change the output prediction of the model. In our setup, we avoid this by directly estimating the parameters of transformations whose kernels are translation-invariant (that is, given a translation

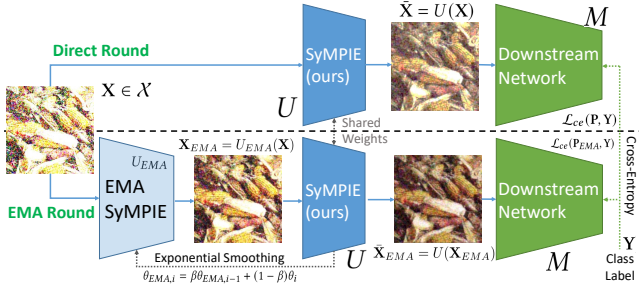


Figure 3: An overview of the training procedure of our modular system.

Algorithm 1 Training step of SyMPIE.

Require: $X \in \mathcal{X}$ is an input image, and $Y \in \mathcal{Y}$ is a vector of the corresponding ground truth category label.

Require: M is the pre-trained upstream task model, U is our SyMPIE to be optimized and U_{EMA} is its smoothed version.

Require: β is the exponential smoothing rate, and λ_{EMA} is the weight of the regularization loss.

- 1: $X_{EMA} \leftarrow U_{EMA}(X)$.
 - 2: $\tilde{X} \leftarrow U(X)$.
 - 3: $\tilde{X}_{EMA} \leftarrow U(X_{EMA})$.
 - 4: $P = M(\tilde{X})$.
 - 5: $P_{EMA} = M(\tilde{X}_{EMA})$.
 - 6: $l_1 = \mathcal{L}_{ce}(P, Y)$.
 - 7: $l_2 = \mathcal{L}_{ce}(P_{EMA}, Y)$.
 - 8: Backpropagate the aggregated loss $l = l_1 + \lambda_{EMA}l_2$.
 - 9: Update the parameters θ of U .
 - 10: Update the parameters of U_{EMA} by $\theta_{EMA} \leftarrow \beta\theta_{EMA} + (1-\beta)\theta$.
 - 11: Repeat Steps 1 to 10 until convergence.
-

T , the parameters are generated by a multidimensional function $\tilde{\pi}(x) = \tilde{\pi}(T(x))$ for all pixel locations x). In particular, our modules apply transformations to the whole image at the same time rather than multiple transformations on local subsets of pixels. Therefore, our modules enforce a stronger and visible change in the input image to allow a change in the downstream prediction. In general, our approach is optimized to enhance the content of an input and improve the performance of the upstream task, not just the perceived image quality, unlike the other approaches. The use of global operations also stabilizes the gradients received by the NEM, enabling its training with fewer samples than other approaches (e.g., when compared to augmentation strategies, our approach requires up to 60 times fewer training samples).

3.2 Training Process

A key benefit of our module is its end-to-end training on any given upstream task (we train it on image classification tasks in our experiments) for its deployment on any downstream task (e.g., image classification, semantic segmentation, etc.) without further need of fine-tuning it. In the following, we describe the training procedure

used to optimize our model (summarized in Fig. 3) and the data augmentation strategy used during the tuning. A detailed description is provided as pseudo-code in Algorithm 1.

We use a frozen model M pre-trained on a source upstream task (e.g., classification) to provide class predictions from the cleaned input samples ($P = M(\tilde{X})$), and we compare them to the labels $Y \in \mathcal{Y}$ using the cross-entropy loss \mathcal{L}_{ce} , to obtain $l_1 = \mathcal{L}_{ce}(P, Y)$.

As will be shown in Sec. 4.4.2, a common failing point of denoisers is modal collapse upon their iterative application to the same image. To avoid this, we introduce an additional regularization term to the training loop, which employs the exponentially smoothed version of our modules (we call this U_{EMA}). The parameters of U_{EMA} at each training iteration $i > 0$ are computed by $\theta_{EMA,i} = \beta\theta_{EMA,i-1} + (1-\beta)\theta_i$, where β is the exponential smoothing rate and $\theta_{EMA,0} = \theta_0$ with θ_i being the parameters of U at iteration i . During the optimization, U_{EMA} is used to generate an intermediate image $X_{EMA} = U_{EMA}(X)$, which will show different visual cues than X , aiding in the generalization.

The intermediate sample is then fed to the modules currently being optimized (i.e., U) to obtain $\tilde{X}_{EMA} = U(X_{EMA})$, which is used in the computation of the regularization loss term. The two-step sample is then processed as normal, obtaining a prediction from the classifier model $P_{EMA} = M(\tilde{X}_{EMA})$ to compute a new loss term $l_2 = \mathcal{L}_{ce}(P_{EMA}, Y)$, weighted by λ_{EMA} .

Our modules attempt to enhance input samples which could not otherwise be handled effectively by popular downstream networks. For this purpose, we design a set of augmentations that serves as a proxy for the distortions experienced at deployment time, as we describe next.

3.2.1 Data Augmentation Pipeline. To train our modules, we employ a variegated data augmentation pipeline, serving two main purposes: (i) to encourage generalization of our architecture over a vast array of possible input corruptions, (ii) to reflect conditions that are seen in the real world during deployment. In particular, we trained our architecture on the ImageNet-1k [48] dataset augmented using the corruptions proposed in [23] together with four additional ones to mimic adverse weather conditions at variable degree of severity.

The four additional corruptions are the following:

- (1) The *darken* corruption mimics the effect of under-exposure of the scene by reducing the intensities of all pixels in a consistent fashion. For example, this situation may happen after encountering glare that forces cameras to reduce the exposure before re-adjusting. A visual example of the resulting images can be seen in the first row of Fig. 4.
- (2) The *horizon* corruption shifts the white-point of the input image to mimic haze on sunset/sunrise scenes. The new white points (RGB) are sampled uniformly in the range $[255, 192 + \delta, 192 - \delta]$, $\delta \sim \mathcal{U}[-32, 32]$, as shown in Fig. 5a. An example is reported in the second row of Fig. 4.
- (3) The *night* corruption simulates nighttime acquisition. We tackle this problem in two ways: (i) changing the white point to a dark blue (uniformly sampled in the color range $[32, 32, 64 + \delta]$, $\delta \sim \mathcal{U}[-24, 24]$, as shown in Fig. 5b) and (ii) darkening the brightest pixels of the scene (which tend to




Figure 4: Qualitative examples of our new input corruptions.



Figure 5: Color ranges used to sample new white points when applying horizon and night corruptions.

be found in the sky region of an outdoor image). The effect of this procedure is shown in the third row of Fig. 4.

- (4) The *white-point* corruption mimics different behaviors of cameras’ white-point balancing procedure, which could lead to images with unrealistic colors even if their relative chromaticity is consistent. In this case, we select a random white point using a uniform distribution $\mathcal{U}[223, 287]$ (i.e., a variation of an eighth) for each of the 3 RGB components and then re-balance the image pixel values accordingly. A visual example is reported in the last row of Fig. 4.

3.3 Inference Process

The goal of our system is to be as easy to use and modular as possible. Therefore, we strove to keep the inference process footprint minimal. In Fig. 1, we reported a comparison between standard inference practice on a downstream task (top half), and the modified pipeline using our modules (bottom half). Notice how our modules can be seamlessly embedded in any architecture to improve its final accuracy.

In detail, to use our system on a given input sample, we first normalize and standardize it (as detailed in Sec. 4.1). Then, for faster processing, we downsize a copy of the input image to have the smallest dimension equal to 232 pixels and extract a central square crop of 224 pixels from it. This low-resolution cropped version is fed to the NEM, which predicts the parameters to enhance the image. Finally, the predicted parameters and the full-resolution image are fed to the DWM to clean the image and enhance its content.

The resulting output can either be used as-is, if the final objective was attaining a cleaned image, or fed to any downstream model to attain a more accurate prediction on the considered task. Overall, the additional computational overhead is minimal, with our modules supporting up to about 300fps at full HD resolution (see Sec. 4.2.4).

Layer	Operations	Feature Size
in	conv 3x3, 64	224x224x64
block1	avg. pool, 5x5, stride 4 2 x (conv 3x3, 128)	56x56x128
block2	avg. pool, 5x5, stride 4 2 x (conv 3x3, 256)	14x14x256
block3	avg. pool, 5x5, stride 4 2 x (conv 3x3, 512)	4x4x512
out	global avg. pool fcn, 37	37

Table 1: Breakdown of our NEM architecture. BatchNorm layers and ReLU activations are added after each convolution.

4 RESULTS AND DISCUSSION

To validate the generalization capability of our modular system, we evaluate its performance on two main tasks: image classification (Sec. 4.2), and semantic segmentation (Sec. 4.3). In the following, we will report the quantitative and qualitative results attained in various tasks by discussing them and drawing comparisons with competing strategies.

4.1 Implementation Details

Experimental Setup. To highlight the generalization capability of our approach, we trained our modules only once on the upstream classification task on the ImageNet dataset. The optimization lasted for 50k steps, using batch size 384 and Adam [29] optimizer with learning rate 10^{-3} scheduled according to a cosine annealing strategy. We remark that when optimizing SyMPIE, the upstream module can be completely frozen, reducing the computational complexity of training. This is the scenario we consider in our work. Furthermore, the training is done once and the same pre-trained weights are used for all the downstream models and tasks.

Our NEM is implemented using a 3-block architecture with strong downsampling (i.e., with stride 4) between layers. A detailed breakdown is reported in Tab. 1. In total, our module uses seven 2D convolutions and a single fully connected layer to project the down-sampled features into the space of parameters $(\mathbf{K}, C_M, C_S) \in \mathbb{R}^A$. We considered a kernel size $K = 5$, which yields $A := d(d+1) + K^2 = 37$ in our setup. During training, the exponential moving average rate was defined as $\beta = 0.9$, and the weight factor for the regularization loss was set to $\lambda_{EMA} = 0.5$. Without loss of generality, the normalization of the input images at inference-time must match the normalization seen during training. In our case, \mathbf{X} is normalized in the range $[0, 1]$ and standardized using mean $\mu = [0.485, 0.456, 0.406]$ and standard deviation $\sigma = [0.229, 0.224, 0.225]$.

Datasets. We employed our system on several real and synthetic datasets to verify its effectiveness in different scenarios for the image classification task (Sec. 4.2). In particular, we used the ImageNetC [23] and ImageNetC-Bar [40] datasets having synthetic corruptions (both based on ImageNet [48])³ and the real dataset VizWiz [3] having natural corruptions. We also introduced

³Both datasets provide corruptions in five intensity levels.

Arch.	Pre-train	Ours	Blur				Digital				Noise			Weather				Clean	Corr.	Avg.	Δ	(% Δ)
			Defocus	Glass	Motion	Zoom	Contrast	Elastic	JPEG	Pixelate	Gaussian	Impulse	Shot	Brightness	Fog	Frost	Snow					
RN50	V2 [57]	✗	43.6	29.7	43.2	42.9	60.0	48.8	61.4	53.0	42.2	37.2	40.5	72.6	59.5	47.0	41.2	80.7	48.2	-	-	
		✓	44.1	36.0	61.2	46.8	61.8	51.4	61.7	61.1	43.6	39.0	42.0	72.0	62.1	49.7	51.0	80.5	52.2	4.0	8.3	
	V1 [57]	✗	37.8	27.7	38.4	35.9	37.1	44.2	55.1	49.0	35.4	31.3	33.2	66.1	43.4	36.8	30.7	75.7	40.1	-	-	
		✓	36.1	30.6	52.5	38.2	33.4	45.0	55.4	56.5	35.9	32.0	33.7	64.4	38.4	38.7	40.4	75.4	42.1	2.0	5.0	
	HA [67]	✗	60.0	47.8	60.3	48.6	55.2	52.9	58.9	70.3	61.2	61.3	61.2	70.9	56.6	54.4	49.6	75.1	57.9	-	-	
		✓	60.2	50.9	68.9	50.0	55.5	54.5	60.3	71.3	61.7	61.4	61.7	70.7	54.8	58.0	58.1	75.1	59.9	2.0	3.5	
	PRIME [41]	✗	46.9	38.9	47.4	45.3	56.6	56.1	61.9	60.5	55.2	55.2	55.3	70.1	52.3	48.9	44.2	76.9	53.0	-	-	
		✓	46.6	42.7	61.9	47.0	58.8	56.5	61.9	64.3	56.4	56.1	56.1	69.2	54.1	50.0	50.3	76.5	55.5	2.5	4.7	
	PIXMIX [26]	✗	43.9	30.5	43.3	40.8	59.1	47.1	61.5	54.6	47.9	46.4	46.4	70.4	55.9	49.6	47.9	77.8	49.7	-	-	
		✓	44.1	35.5	60.7	44.5	60.4	49.1	61.5	59.3	49.6	48.0	48.4	69.6	56.8	51.0	52.5	77.7	52.7	3.0	6.0	
VGG16	V1 [57]	✗	25.2	20.0	28.0	28.3	28.8	35.7	43.0	34.0	21.9	17.7	19.8	58.9	35.2	26.8	22.4	71.4	29.7	-	-	
		✓	22.3	21.6	39.5	28.6	24.4	36.2	43.4	46.0	23.0	19.8	21.2	56.9	28.2	29.2	31.2	71.2	31.4	1.7	5.7	
Swin-T	V1 [57]	✗	45.2	35.5	52.4	42.4	64.1	53.2	64.1	58.2	52.2	48.9	49.0	73.5	61.8	56.7	50.9	81.3	53.9	-	-	
		✓	44.9	38.6	62.1	43.6	63.7	55.1	63.9	63.8	52.8	49.5	49.6	72.8	62.3	57.2	56.4	80.7	55.7	1.8	3.3	
CLIP	V1 [45]	✗	22.4	14.4	21.9	21.1	28.4	25.6	34.5	29.7	18.3	12.9	16.7	45.2	33.4	23.4	18.9	55.2	24.5	-	-	
		✓	18.6	17.3	30.8	21.4	25.4	26.1	34.9	35.4	18.4	12.9	17.2	43.7	29.9	24.2	24.3	55.1	25.4	0.9	3.7	

Table 2: Results for the image classification task on the ImageNetC dataset (higher is better).

ImageNetC-mixed⁴, a new benchmark to investigate the reliability of models when presented with multiple corruptions at once. We built it by randomly applying 1-to-3 corruptions to the same image, sampled from the joint pool of augmentations of our data augmentation pipeline and those proposed in ImageNetC.

For what concerns semantic segmentation (Sec. 4.3), we analyzed the models using driving scenes tackling the domain adaptation problem for clear-to-adverse weather conditions. In this case, we employed three real world datasets: i) Cityscapes [9] as the training (source) domain; ii) ACDC [51] and iii) DarkZurich [50] as testing downstream (target) domains.

Metrics. In image classification, we report the per-corruption accuracies (Acc , \uparrow), their mean ($Corr. Avg.$, \uparrow), and the accuracy on corruption-free data ($Clean$, \uparrow). In semantic segmentation, we report the results as per-class IoU (Intersection over Union) or its mean (mIoU, \uparrow). In all cases, Δ (% Δ) refers to the absolute (relative) gain with respect to the considered baseline.

4.2 Results for Image Classification

In the image classification task, we investigate three main scenarios: i) single synthetic corruptions (using ImageNetC [23]) in Sec. 4.2.1, ii) multiple synthetic corruptions (using ImageNetC-mixed) in Sec. 4.2.2, and iii) real-world corruptions (using VizWiz [3]) in Sec. 4.2.3. Finally, we analyze the computational cost in Sec. 4.2.4.

4.2.1 Results on Single Synthetic Corruptions (ImageNetC).

For this discussion, we refer to Table 2, where we report the final accuracy attained by our approach when mounted on multiple different backbones. In particular, we use a ResNet50 [18] with different pre-training weights (TorchVision [57] V2 and V1, HA [67], PRIME [41], and PIXMIX [26]), a VGG16 [52], a Swin-Tiny [33], and finally the CLIP [45] foundation model. Notably, we train our SyMPIE model only once using ResNet50 with TorchVision’s V2 weights⁵, and we use the trained model as-is on all other experiments, highlighting the generalization capability of our approach.

⁴The split will be released for reproducibility.

⁵This design choice was motivated by the superior performance on clean data with respect to other convolutional networks.

Our SyMPIE consistently improves the average accuracy across the considered models with an average gain of 2.2% in absolute terms and an average relative gain of 5.0%. As anticipated, a remarkable feature of SyMPIE is the capability to improve the performance even for architectures designed to work well in the ImageNetC task, such as HA [67], PRIME [41], and PIXMIX [26]. These are three strong state-of-the-art data augmentation approaches and, therefore, can handle corrupted samples from ImageNetC better than weaker data augmentation approaches. Keeping this in mind, the improvement brought by our approach when used jointly with these architectures is even more striking, signifying that the effect of our module is complementary to existing state-of-the-art approaches, and using both strategies together can improve the absolute accuracy significantly. Numerically, when using our system together with HA, we are able to improve the performance over TorchVision’s V1 baseline by almost 20% in absolute terms (49.4% relative).

Moreover, the gain is well-spread across the various corruptions and, in cases particularly suited to be tackled by our constrained model, we can appreciate some higher gains. For example, in *Motion Blur* and *Snow*, we get significant performance gains (up to 9% absolute points) even on already-robust backbones like those pre-trained via HA. We believe that the *motion blur* kernel can be estimated accurately with our learned filter, whereas our affine color transforms handle the low contrast brought by *snow*, regardless of the localized bright spots. We experience slight performance drops on the *fog* corruption, due to changes to the frequency distribution our system may introduce via the global filter. Methods such as HA are designed with frequency-spectra changes in mind. Therefore, they expect a distribution of input images with specific frequency characteristics. Another performance drop is observed on the *Brightness* corruption. This corruption is modeled as a non-linear change in the HSV color space, and our affine color transforms could not approximate the underlying function accurately. Note that these limitations introduce only marginal drops in accuracy, and do not change the overall improvements brought by our system.

Lastly, we discuss the results when employing our approach with downstream architectures that are different than the one used for pre-training our modules. We verify the usefulness of our method

Pre-training	Base	Ours	Δ	% Δ
V2 [57]	38.2	41.0	2.8	7.3
V1 [57]	31.2	32.9	1.7	5.5
HA [67]	47.2	48.9	1.7	3.6
PRIME [41]	43.4	44.9	1.5	3.5
PIXMIX [26]	41.1	43.2	2.1	5.1

Table 3: Accuracy on ImageNetC-mixed via ResNet50.

Ours	Blur	Bright	Frame	Rotation	Obscured	Dark	Clean	Corr. Avg.
\times	44.4	40.9	42.7	38.5	37.9	44.3	48.6	41.5
\checkmark	44.6	40.9	42.6	39.2	39.7	45.3	48.8	42.0

Table 4: Results for the VizWiz with ResNet50 V2.

on the widely used VGG16 convolutional architecture and on a transformer-based architecture (e.g., Swin-Tiny). Then, we considered the CLIP foundation model [45] that was pre-trained to align image-text embeddings to the same semantic value, rather than explicitly recognizing the input image category. Our modules bring improvement even in this case, highlighting how the approach is able not only to change the graphical appearance but also to highlight the semantic content, making it easier for the downstream network.

4.2.2 Results on Mixed Synthetic Corruptions (ImageNetC-mixed). Following the tests on a single corruption at a time, we consider the more challenging setting of multiple corruptions together using the hereby proposed ImageNetC-mixed dataset. (refer to the datasets section in 4.1 for more information). The results of these analyses are shown in Table 3. Despite the existing data augmentation approaches losing about 10% accuracy compared with the single corruption case, our approach maintains a stable gain of 2.0% in absolute terms, corresponding to 5.0% relative improvement. This proves the ability of our approach to handle the composition of input corruptions, as it is often encountered in practice.

4.2.3 Results on Real-World Corruptions (VizWiz). In the previous sections, we have investigated the performance of our modules on synthetic corruptions, following the standard evaluation pipeline [23, 24]. However, it is important to verify that the performance improvement is maintained even on real data. To this end, we used the corrupted real-world VizWiz dataset [3], which, to the best of our knowledge, is the only real-world corrupted images dataset for the considered task currently available. The numerical results are reported in Table 4, where we observe a 1.2% relative gain on corrupted images, while also improving the accuracy on clean images by relative 0.4%.

4.2.4 Computational Cost. We report the complexity of our approach in FLOPs (Floating Point Operations) in Table 5, comparing it to other architectures that can be used as (or converted into) input-level denoisers. We computed the cost of our approach using the PTFlops library [53] at a resolution of $1920 \times 1080\text{px}^2$.

Our module is one to four orders of magnitude faster than the diffusion model competitors. Moreover, its computational complexity does not scale significantly with the input resolution - since a resizing stage is done before parameter estimation and the only

	Model	GFLOPs
Diffusion	SyMPIE (ours)	2.0
	LDM-8 [58]	37.1
	Frido [58]	37.3
	Dido-gating [58]	39.7
	Diff 64 [11]	658.0
	Diff 128 [11]	1860.0
	Diff 256 [11]	6700.0
Auto-Encoder	UNet [1]	103.0
	GoConv-UNet [1]	67.7
	OrgOct-UNet [1]	65.6
	MSUDCAE [68]	37.5
	MSDCAE [68]	23.5
GANs	CBAM_MSUDCAE [68]	37.4
	CycleGAN [60]	52.9
	U-GAT-IT-light [8]	105.0
	NICE-GAN [8]	67.6
	N2D-GAN [31]	43.6
	SPN2D-GAN [31]	94.0
STGAN [43]	519.0	

Table 5: Computational complexity of our method compared to other input-level image enhancement strategies.

operations applied on the full-resolution image are the spatial convolution and a matrix multiplication over the channels (see Sec. 3.3). This means that we can handle images of arbitrary resolution, contrary to what happens for, e.g., the fixed-size diffusion models. Similar considerations also hold for auto-encoders and GANs. This confirms the efficiency of our strategy.

Finally, we computed the throughput of our module on an NVIDIA GTX 1080Ti GPU, obtaining a speed of 289.4 images/second; corresponding to a per-image inference time of about 3ms, meaning that our approach can be easily employed in real-time applications with minimal impact on the inference time of downstream architectures.

4.3 Results for Semantic Segmentation

The analyses up to this point have been in the same task as training (i.e., image classification) with some changes in the downstream architecture. However, to truly show the generalization capabilities of our approach, we change the downstream task completely and verify that the accuracy improvement is preserved. For this investigation, we choose the semantic segmentation task and, in particular, domain adaptation to adverse weather conditions. We train the baseline architecture used in [4, 5] (DeepLabV2 [7] with ResNet50 backbone [19]) on the Cityscapes [9] dataset, and deploy it on two adverse weather datasets, namely, ACDC [51] and DarkZurich [50].

The quantitative results of this study are reported in Table 6, where our architecture improves the mIoU score by 4.1% and 3.8% relative gain on ACDC and DarkZurich, respectively. We report some qualitative results on ACDC in Fig. 6, showing the original RGB image, the image warped by our architecture, and the corresponding predictions made by the segmentation model. In the first image (snow scenario), our model improves the overall contrast. This aids in distinguishing the sidewalk on the right of the

Model	Dataset	Road	Sidewalk	Building	Wall	Fence	Pole	T. Light	T. Sign	Vegetation	Terrain	Sky	Person	Rider	Car	Truck	Bus	Train	Motorcycle	Bicycle	mIoU	Δ	$\% \Delta$
Source Only	ACDC	73.5	33.2	62.1	21.3	20.6	23.3	43.4	34.1	69.6	25.1	75.4	38.9	15.9	65.9	29.7	30.2	36.6	16.5	23.0	38.9	-	-
	DarkZurich	63.2	12.0	49.9	9.1	12.2	9.7	15.1	7.7	40.3	12.8	2.5	16.0	7.9	31.1	-	-	0.0	8.4	17.0	18.6	-	-
Ours	ACDC	68.6	36.2	60.1	22.7	22.2	25.5	46.0	35.6	72.3	25.9	72.4	42.7	18.3	68.4	35.0	30.2	39.5	19.6	27.7	40.5	1.6	4.1
	DarkZurich	53.0	15.9	46.6	10.7	13.3	10.7	21.4	8.4	45.7	12.1	1.0	16.8	7.8	34.2	-	-	0.0	10.0	21.0	19.3	0.7	3.8

Table 6: Quantitative results for semantic segmentation using a DeepLabV2 [7] architecture with the ResNet50 backbone.




Figure 6: Qualitative results on the ACDC semantic segmentation benchmark.

Pre-training	Ours	Blue	Brown	Caustic	Checker	C. Sine	Inv. Sparkle	Perlin	Plasma	S. Sine	Sparkle	Clean	Corr. Avg.	Δ	$\% \Delta$
V2 [57]	\times	61.1	66.5	59.3	57.2	19.3	28.7	67.4	36.4	30.0	69.0	80.7	48.5	-	-
	\checkmark	52.2	65.8	58.5	56.1	20.1	29.4	66.7	37.1	33.3	68.2	80.5	48.7	0.2	0.4
HA [67]	\times	64.3	64.0	58.7	48.7	40.1	25.1	68.8	31.5	57.4	64.8	75.1	52.3	-	-
	\checkmark	64.7	63.5	58.6	48.0	41.1	26.9	68.9	31.6	59.1	64.7	75.1	52.6	0.3	0.6

Table 7: Quantitative results on the ImageNetC-Bar dataset with ResNet50.

scene (light blue box), which was missed almost completely by the segmentation system when shown the original data. Moreover, it helps to recognize more accurately the building (light green box). In the second image (*night scene*), instead, our method brightens the scene, allowing us to identify the person on the side (light green box), which was not recognized by the source only architecture. We are also able to slightly improve the precision of the sidewalk on the left (light blue box).

4.4 Ablation Studies

We report some ablation studies, namely: an analysis of corruptions that cannot be modeled by our strategy (ImageNetC-Bar, Sec. 4.4.1); a study on the iterative application of our module, which is a common failure point for denoisers and image enhancers (Sec. 4.4.2); a proof of concept for applicability on video compression (Sec. 4.4.3).

4.4.1 Results on corruptions that cannot be modeled by SyMPIE by construction (ImageNetC-Bar). We further analyze the reliability of our method to corruptions that cannot be modeled by our modules. For this purpose, we test its performance on the ImageNetC-Bar dataset [40], which provides a corruption set complementary to ImageNetC. In general, the corruptions present in the dataset cannot be suitably modeled by our approach, since most of them contain non-spatially-uniform distortions.

Nonetheless, we observe small accuracy gains when using our modules. This means that, even if we are not able to model the corruptions in the dataset, SyMPIE can modify the inputs to make the classification task easier. Looking at the results, we can identify three corruptions where our method brings consistent and noticeable gains: *Inverse Sparkle*, *Plasma*, and *Single Sine*. In the first corruption, our method improves the accuracy by an average of 2.5% in absolute terms (4.8% relative gain), in the second we improve by 0.4% (1.1%), while in the third we improve by 2.5% (7.0%). These three corruptions, in some way, are similar to those present in ImageNet-C and can be dealt with more effectively by our approach. In particular, the *Plasma* noise is a colored version of *Fog*; *Inverse Sparkle* is similar to *Brightness* in some instances, and the effect of *Single Sine* is similar to *Contrast* and *Fog*.

4.4.2 Iterative Application. Furthermore, we investigate a common fallacy of denoising and enhancing models, especially those based on autoencoder architectures [64]: when fed their own outputs iteratively, their prediction may collapse into undesirable modes, or destroy completely the input content. To evaluate qualitatively the susceptibility of our approach to this scenario, we fed the predictions provided by our system to itself four times and collected the results in Fig. 7. The first row contains the image corrupted with the fifteen strategies provided by the dataset with




Figure 7: Qualitative results on iterated application of our method on an input image from the ImageNetC validation set. The first row represents the input, while the others correspond to four iterative applications of our modules.

Compression Rate	MSE ✗	MSE ✓	Δ	$\% \Delta$
45	0.087	0.085	-0.002	-2.6
50	0.117	0.113	-0.005	-4.0
55	0.117	0.112	-0.005	-4.0
60	0.117	0.113	-0.004	-3.8

Table 8: Results for feature extraction from compressed video data.

a mid-severity level, while the other rows show the outputs of iterative applications of our approach.

From the figure, we see that our architecture does not suffer from the mentioned problem, and all outputs are consistent with the original input. Moreover, in many cases, one can see a clear improvement following the subsequent application of our module. The most striking cases are the *motion blur*, *frost*, and *snow* corruptions highlighting the potentiality for deployment of our architecture. In the first case, our approach is able to completely remove the blurring artifacts, while in the other two, it removes the haze from the image and increases contrast. Also, our method enhances significantly the overall image content and colors, while removing the blur at the same time for the *Glass Blur* (second column).

4.4.3 Video Compression. Finally, we provide a proof-of-concept investigation of our proposed system for live-streaming applications (e.g., for scene understanding applications on broadcasted data). We downloaded a freely available video from YouTube⁶ and we compressed it multiple times using the VLC Media Player software⁷ at different compression rates, before extracting the frames from the videos for their analysis.

We take each frame obtained from the original, compressed, and processed videos to a VGG19 [52] architecture, and compare the similarity of the features extracted by the model (via mean squared error, MSE, ↓). This metric is representative of the overall

reconstruction fidelity as well as the potential applicability to an end task, since the extracted features are strongly correlated to the final predictions, regardless of the end task employed. The results given in Table 8 show that our method brings a 4.0% (2.6%) relative improvement in the high- (low-) compression regime in the linear scale.

5 CONCLUSIONS

In this paper, we introduced a novel, modular, and efficient system that predicts explicit parameters for content enhancement of an input image targeting improved accuracy on downstream tasks. The computational footprint of the modules is minimal, being more than 10x faster than competing approaches (2GFLOPs in total), and enjoys a throughput of about 300 images per second. A key feature of our approach is the capability of training without any paired (clean/corrupted) input samples, but rather learning automatically the most suitable transformations of an input image using the supervision on an upstream task. Even more remarkably, once the system has been trained on a given upstream task, it can generalize to arbitrary downstream tasks without fine-tuning. To confirm the generalization claim, we validate our approach on three classification datasets and two semantic segmentation datasets, achieving noticeable improvement in all of them.

In future work, we plan to introduce new parameterizations and operators to capture a richer representation of the input and estimate more accurately the enhancing parameters (e.g., features of frequency components, or non-spatially-uniform operators), while maintaining high efficiency and full differentiability, which are fundamental for our approach. Another avenue of research will target compressed and streaming video data, investigating the effect of re-using parameters across multiple frames with the goal of further improving the computational efficiency, while also evaluating our approach from the quality-of-experience point of view. A final plan involves the investigation of a multi-task training scenario, allowing the model to learn more general transformations.

⁶<https://www.youtube.com/watch?v=nTGXmmBTmzo>, accessed on 26 Nov 2023

⁷<https://www.videolan.org/>, accessed on 26 Nov 2023

REFERENCES

- [1] AKBARI, M., LIANG, J., HAN, J., AND TU, C. Generalized octave convolutions for learned multi-frequency image compression. *arXiv preprint arXiv:2002.10032* (2020).
- [2] ALBREIKI, F., GHAZAL, S. A., LAHOUD, J., ANWER, R., CHOLAKKAL, H., AND KHAN, F. On the robustness of 3d object detectors. In *Proceedings of the 4th ACM International Conference on Multimedia in Asia* (2022), pp. 1–7.
- [3] BAFGHI, R. A., AND GURARI, D. A new dataset based on images taken by blind people for testing the robustness of image classification models trained for imagenet categories. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 16261–16270.
- [4] BARBATO, F., MICHELI, U., TOLDO, M., AND ZANUTTIGH, P. Road scenes segmentation across different domains by disentangling latent representations. *arXiv preprint arXiv:2108.03021* (2021).
- [5] BARBATO, F., TOLDO, M., MICHELI, U., AND ZANUTTIGH, P. Latent space regularization for unsupervised domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops* (June 2021), pp. 2835–2845.
- [6] CHEN, G., PENG, P., MA, L., LI, J., DU, L., AND TIAN, Y. Amplitude-phase recombination: Rethinking robustness of convolutional neural networks in frequency domain. In *Proceedings of International Conference on Computer Vision (ICCV)* (2021).
- [7] CHEN, L.-C., PAPANDREOU, G., KOKKINOS, I., MURPHY, K., AND YUILLE, A. L. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs, 2017.
- [8] CHEN, R., HUANG, W., HUANG, B., SUN, F., AND FANG, B. Reusing discriminators for encoding: Towards unsupervised image-to-image translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (2020), pp. 8168–8177.
- [9] CORDTS, M., OMRAN, M., RAMOS, S., REHFELD, T., ENZWEILER, M., BENENSON, R., FRANKE, U., ROTH, S., AND SCHIELE, B. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016).
- [10] CUBUK, E. D., ZOPH, B., MANE, D., VASUDEVAN, V., AND LE, Q. V. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501* (2018).
- [11] DHARWAL, P., AND NICHOL, A. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* 34 (2021), 8780–8794.
- [12] DUAN, L., LIU, J., YANG, W., HUANG, T., AND GAO, W. Video coding for machines: A paradigm of collaborative compression and intelligent analytics. *IEEE Transactions on Image Processing* 29 (2020), 8680–8695.
- [13] FEINMAN, R., CURTIN, R. R., SHINTRE, S., AND GARDNER, A. B. Detecting adversarial samples from artifacts. *arXiv preprint arXiv:1703.00410* (2017).
- [14] GEIRHOS, R., RUBISCH, P., MICHAELIS, C., BETHGE, M., WICHMANN, F. A., AND BRENDEL, W. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. *arXiv preprint arXiv:1811.12231* (2018).
- [15] GONG, T., JEONG, J., KIM, T., KIM, Y., SHIN, J., AND LEE, S.-J. NOTE: Robust continual test-time adaptation against temporal correlation. In *Advances in Neural Information Processing Systems (NeurIPS)* (2022).
- [16] GROSSE, K., MANOHARAN, P., PAPERNOT, N., BACKES, M., AND MCDANIEL, P. On the (statistical) detection of adversarial examples. *arXiv preprint arXiv:1702.06280* (2017).
- [17] GÜNDÜZ, D., QIN, Z., AGUERRI, I. E., DHILLON, H. S., YANG, Z., YENER, A., WONG, K. K., AND CHAE, C.-B. Beyond transmitting bits: Context, semantics, and task-oriented communications. *IEEE Journal on Selected Areas in Communications* 41, 1 (2023), 5–41.
- [18] HE, K., ZHANG, X., REN, S., AND SUN, J. Deep residual learning for image recognition, 2015.
- [19] HE, K., ZHANG, X., REN, S., AND SUN, J. Deep residual learning for image recognition. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016).
- [20] HENDRYCKS, D., BASART, S., MU, N., KADAVATH, S., WANG, F., DORUNDO, E., DESAI, R., ZHU, T., PARAJULI, S., GUO, M., ET AL. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2021), pp. 8340–8349.
- [21] HENDRYCKS, D., BASART, S., MU, N., KADAVATH, S., WANG, F., DORUNDO, E., DESAI, R., ZHU, T., PARAJULI, S., GUO, M., SONG, D., STEINHARDT, J., AND GILMER, J. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of International Conference on Computer Vision (ICCV)* (2021).
- [22] HENDRYCKS, D., AND DIETTERICH, T. Benchmarking neural network robustness to common corruptions and perturbations. *International Conference on Learning Representations (ICLR)* (2019).
- [23] HENDRYCKS, D., AND DIETTERICH, T. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261* (2019).
- [24] HENDRYCKS, D., MU, N., CUBUK, E. D., ZOPH, B., GILMER, J., AND LAKSHMINARAYANAN, B. AugMix: A simple data processing method to improve robustness and uncertainty. *International Conference on Learning Representations (ICLR)* (2020).
- [25] HENDRYCKS, D., ZOU, A., MAZEIKA, M., TANG, L., LI, B., SONG, D., AND STEINHARDT, J. Pixmix: Dreamlike pictures comprehensively improve safety measures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 16783–16792.
- [26] HENDRYCKS, D., ZOU, A., MAZEIKA, M., TANG, L., LI, B., SONG, D. X., AND STEINHARDT, J. Pixmix: Dreamlike pictures comprehensively improve safety measures. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2021).
- [27] HU, X., UZUNBAS, M. G., CHEN, S., WANG, R., SHAH, A., NEVATIA, R., AND LIM, S.-N. Mixnorm: Test-time adaptation through online normalization estimation. *International Conference on Learning Representations (ICLR)* (2022).
- [28] KHURANA, A., PAUL, S., RAI, P., BISWAS, S., AND AGGARWAL, G. SITA: Single Image Test-time Adaptation. *ArXiv:2112.02355* (2021).
- [29] KINGMA, D. P., AND BA, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [30] KIRKPATRICK, J., PASCANU, R., RABINOWITZ, N., VENESS, J., DESJARDINS, G., RUSU, A. A., MILAN, K., QUAN, J., RAMALHO, T., GRABSKA-BARWINSKA, A., ET AL. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences* 114, 13 (2017), 3521–3526.
- [31] LI, X., AND GUO, X. Spn2d-gan: Semantic prior based night-to-day image-to-image translation. *IEEE Transactions on Multimedia* (2022).
- [32] LIN, H., VAN ZUIJLEN, M., PONT, S. C., WIJNTJES, M. W., AND BALA, K. What can style transfer and paintings do for model robustness? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 11028–11037.
- [33] LIU, Z., LIN, Y., CAO, Y., HU, H., WEI, Y., ZHANG, Z., LIN, S., AND GUO, B. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision* (2021), pp. 10012–10022.
- [34] LORE, K. G., AKINTAYO, A., AND SARKAR, S. Llnet: A deep autoencoder approach to natural low-light image enhancement. *Pattern Recognition* 61 (2017), 650–662.
- [35] LU, J., ISSARANON, T., AND FORSYTH, D. Safetynet: Detecting and rejecting adversarial examples robustly. In *Proceedings of the IEEE international conference on computer vision* (2017), pp. 446–454.
- [36] MADRY, A., MAKELOV, A., SCHMIDT, L., TSIPRAS, D., AND VLADU, A. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083* (2017).
- [37] MASOOD, S., HUSSAIN, A., JAFFAR, M. A., AND CHOI, T.-S. Color differences based fuzzy filter for extremely corrupted color images. *Applied Soft Computing* 21 (2014), 107–118.
- [38] MICHELI, U., AND OZAY, M. Online continual learning for robust indoor object recognition. In *Proceedings of IEEE International Conference on Robots and Systems (IROS)* (2023).
- [39] MILLER, D. J., XIANG, Z., AND KESIDIS, G. Adversarial learning targeting deep neural network classification: A comprehensive review of defenses against attacks. *Proceedings of the IEEE* 108, 3 (2020), 402–433.
- [40] MINTUN, E., KIRILLOV, A., AND XIE, S. On interaction between augmentations and corruptions in natural corruption robustness. *Advances in Neural Information Processing Systems* 34 (2021), 3571–3583.
- [41] MODAS, A., RADE, R., ORTIZ-JIMÉNEZ, G., MOOSAVI-DEZFOOLI, S.-M., AND FROSSARD, P. Prime: A few primitives can boost robustness to common corruptions. In *European Conference on Computer Vision* (2022), Springer, pp. 623–640.
- [42] NI, Z., YANG, W., WANG, S., MA, L., AND KWONG, S. Towards unsupervised deep image enhancement with generative adversarial network. *IEEE Transactions on Image Processing* 29 (2020), 9140–9151.
- [43] NIE, X., JIA, J., DING, H., AND WONG, E. K. Gigan: Gate in gan, could gate mechanism filter the features in image-to-image translation? *Neurocomputing* 462 (2021), 376–388.
- [44] NIU, S., WU, J., ZHANG, Y., WEN, Z., CHEN, Y., ZHAO, P., AND TAN, M. Towards stable test-time adaptation in dynamic wild world. In *International Conference on Learning Representations (ICLR)* (2023).
- [45] RADFORD, A., KIM, J. W., HALLACY, C., RAMESH, A., GOH, G., AGARWAL, S., SASTRY, G., ASKELL, A., MISHKIN, P., CLARK, J., KRUEGER, G., AND SUTSKEVER, I. Learning transferable visual models from natural language supervision. *ArXiv:2103.00020* (2021).
- [46] RECHT, B., ROELOFS, R., SCHMIDT, L., AND SHANKAR, V. Do imagenet classifiers generalize to imagenet? *ArXiv:1902.10811* (2019).
- [47] RIZZOLI, G., BARBATO, F., CALIGIURI, M., AND ZANUTTIGH, P. Syndrome-multi-modal uav dataset for urban scenarios. In *Proceedings of International Conference on Computer Vision Workshops (ICCVW)* (2023), pp. 2210–2220.
- [48] RUSSAKOVSKY, O., DENG, J., SU, H., KRAUSE, J., SATHEESH, S., MA, S., HUANG, Z., KARPATY, A., KHOSLA, A., BERNSTEIN, M., ET AL. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)* (2015).
- [49] SAIKIA, T., SCHMID, C., AND BROX, T. Improving robustness against common corruptions with frequency biased models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2021), pp. 10211–10220.
- [50] SAKARIDIS, C., DAI, D., AND GOOL, L. V. Guided curriculum model adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2019),

- pp. 7374–7383.
- [51] SAKARIDIS, C., DAI, D., AND VAN GOOL, L. ACDC: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *Proceedings of International Conference on Computer Vision (ICCV)* (2021).
 - [52] SIMONYAN, K., AND ZISSERMAN, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
 - [53] SOVRASOV, V. ptflops: a flops counting tool for neural networks in pytorch framework, 2018-2023.
 - [54] SZEGEDY, C., ZAREMBA, W., SUTSKEVER, I., BRUNA, J., ERHAN, D., GOODFELLOW, I., AND FERGUS, R. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199* (2013).
 - [55] TESTOLINA, P., BARBATO, F., MICHELI, U., GIORDANI, M., ZANUTTIGH, P., AND ZORZI, M. Selma: Semantic large-scale multimodal acquisitions in variable weather, daytime and viewpoints. *IEEE Transactions on Intelligent Transportation Systems* (2023).
 - [56] TIAN, C., FEI, L., ZHENG, W., XU, Y., ZUO, W., AND LIN, C.-W. Deep learning on image denoising: An overview. *Neural Networks* 131 (2020), 251–275.
 - [57] TORCHVISION MAINTAINERS AND CONTRIBUTORS. TorchVision: PyTorch's Computer Vision library. <https://github.com/pytorch/vision>, 2016.
 - [58] ULHAQ, A., AKHTAR, N., AND POGREBNA, G. Efficient diffusion models for vision: A survey, 2022.
 - [59] WANG, D., SHELHAMER, E., LIU, S., OLSHAUSEN, B., AND DARRELL, T. Tent: Fully test-time adaptation by entropy minimization. In *International Conference on Learning Representations (ICLR)* (2021).
 - [60] WANG, H., GUI, S., YANG, H., LIU, J., AND WANG, Z. Gan slimming: All-in-one gan compression by a unified optimization framework. In *European Conference on Computer Vision* (2020), Springer, pp. 54–73.
 - [61] WANG, H., XIAO, C., KOSSAIFI, J., YU, Z., ANANDKUMAR, A., AND WANG, Z. Augmax: Adversarial composition of random augmentations for robust training. In *Advances in Neural Information Processing Systems (NeurIPS)* (2021).
 - [62] WANG, R., SUN, Y., GANDELSMAN, Y., CHEN, X., EFROS, A. A., AND WANG, X. Test-time training on video streams. *arXiv:2307.05014* (2023).
 - [63] XU, W., EVANS, D., AND QI, Y. Feature squeezing: Detecting adversarial examples in deep neural networks. *arXiv preprint arXiv:1704.01155* (2017).
 - [64] YACOBY, Y., PAN, W., AND DOSHI-VELEZ, F. Failure modes of variational autoencoders and their effects on downstream tasks. *arXiv preprint arXiv:2007.07124* (2020).
 - [65] YI, X., XU, H., ZHANG, H., TANG, L., AND MA, J. Diff-retinex: Rethinking low-light image enhancement with a generative diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2023), pp. 12302–12311.
 - [66] YUCEL, M. K., CINBIS, R. G., AND DUYGULU, P. How robust are discriminatively trained zero-shot learning models? *Image and Vision Computing* 119 (2022), 104392.
 - [67] YUCEL, M. K., CINBIS, R. G., AND DUYGULU, P. HybridAugment++: Unified frequency spectra perturbations for model robustness. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023), pp. 5718–5728.
 - [68] ZHANG, H., LIU, S., WANG, C., LU, S., AND XIONG, W. Color-patterned fabric defect detection algorithm based on triplet attention multi-scale u-shape denoising convolutional auto-encoder. *The Journal of Supercomputing* (2023), 1–26.
 - [69] ZHANG, M., LEVINE, S., AND FINN, C. Memo: Test time robustness via adaptation and augmentation. *Advances in Neural Information Processing Systems (NeurIPS)* (2022).
 - [70] ZHENG, J., VALAVANIS, K. P., AND GAUCH, J. M. Noise removal from color images. *Journal of Intelligent and Robotic Systems* 7 (1993), 257–285.