EFFICIENT FINE-TUNING OF NEURAL NETWORKS FOR ARTIFACT REMOVAL IN DEEP LEARNING FOR INVERSE IMAGING PROBLEMS

Alice Lucas[1], Santiago Lopez-Tapia[2], Rafael Molina[2], Aggelos K. Katsaggelos[1]
[1]Department of Electrical and Computer Engineering, Northwestern University, Evanston, IL, USA
[2]Computer Science and Artificial Intelligence Department, Universidad de Granada, Spain

Introduction

Deep Neural Networks (DNN) for super-resolution (SR) are prone to generating artifacts in their produced images. These artifacts may arise from the use of an image formation model at test time that differs from the one assumed in training, or from challenging training dynamics associated with Generative Adversarial Networks (GANs).

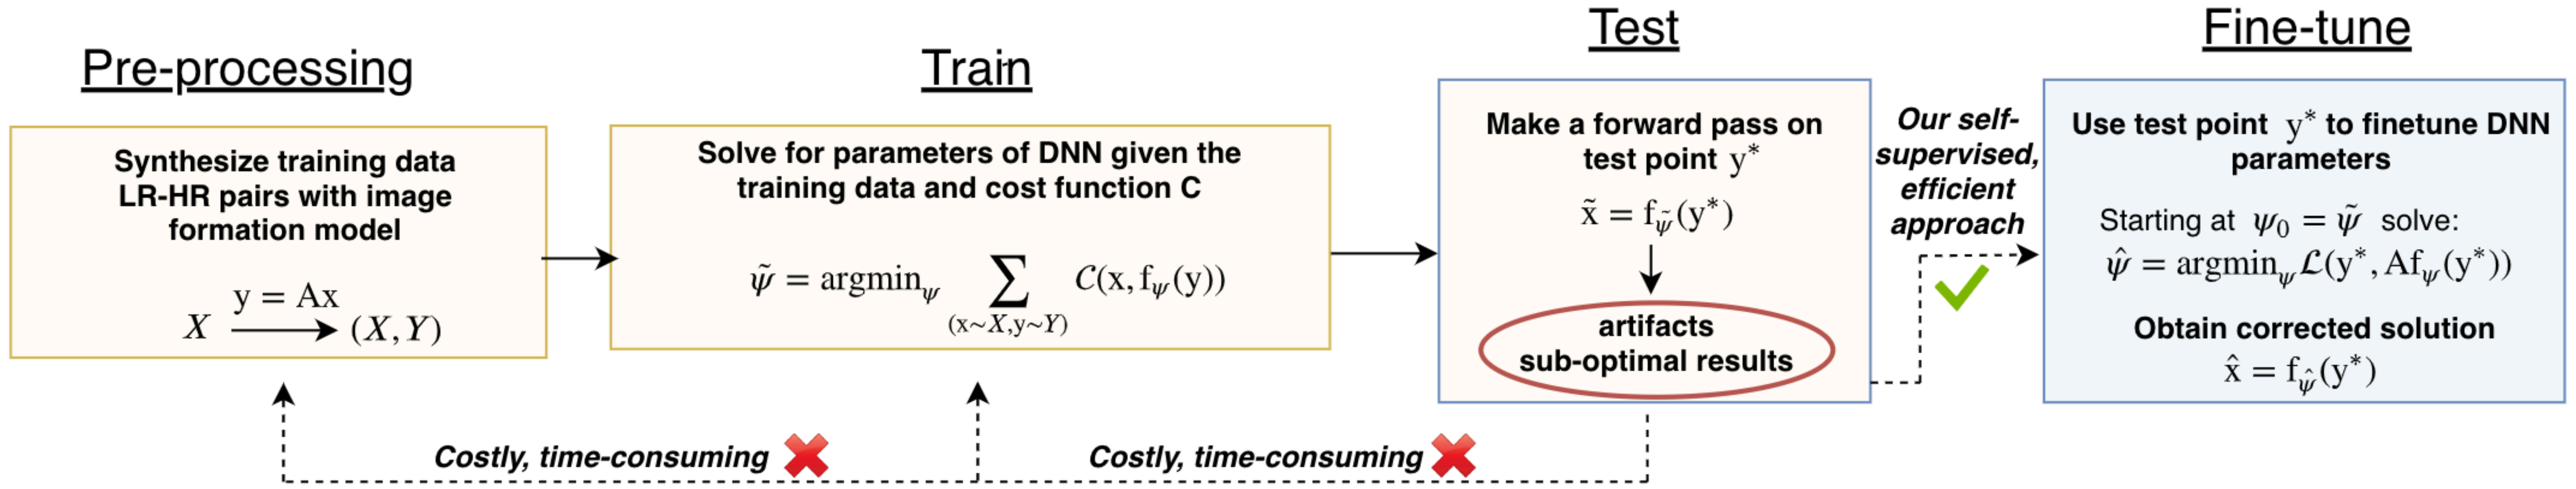

Project Aims: In order to correct for artifacts produced by DNNs in inverse imaging problems, we present a fine-tuning framework that uses the data consistency constraint to enhance the obtained solution at test time. Our method is agnostic to the type of artifacts and does not require re-training the DNN.

Methods

We assume access to a SR neural network $f_{\tilde{\psi}}(\cdot)$ with converged parameters $\tilde{\psi}$ trained on degradation operator A. When tested on a test image y^*, the DNN introduces artifacts in the super-resolved solution. Our fine-tuning method employs the data consistency constraint $Af_\psi(y^*) = y^*$ to fine-tune our network and obtain a new set of optimal parameters $\hat{\psi}$:

$$\hat{\psi} = \arg\min_\psi \mathcal{L}\left(y^*, Af_\psi(y^*)\right) = \arg\min_\psi ||Af_\psi(y^*) - y^*||_2^2$$

We use gradient descent to solve the optimization problem and set $\psi_0 = \tilde{\psi}$ at the first iteration. We use automatic differentiations offered by DL libraries to compute the necessary gradients. The fine-tuned artifacts-free solution becomes $\hat{x} = f_{\hat{\psi}}(y^*)$.

References & Acknowledgements

The VSR neural networks used for the experiments shown in this poster were borrowed from: Lucas, A., Lopez-Tapia, S., Molina, R., & Katsaggelos, A. K. (2019). Generative adversarial networks and perceptual losses for video super-resolution. *IEEE Transactions on Image Processing*, *28*(7), 3312-3327.
This work was supported in part by the Sony 2016 Research Award Program Research Project. The work of SLT and RM was supported by the the Spanish Ministry of Economy and Competitiveness through project DPI2016-77869-C2-2-R and the Visiting Scholar program at the University of Granada. SLT received financial support through the Spanish FPU program.

Results

Experiment #1: change in scale factor at test time. When a network is trained for a specific scale factor but provided an image downsampled by another factor at test time, distortions arise in the resulting super-resolved image. We wish to remove these distortions while keeping the increase in resolution. The figures below show the results of fine-tuning a VSR network on a LR sequence downsampled by 3 when it was originally trained for scale factor 4. The PSNR and SSIM are displayed below each figure.

(a) Bicubic interpolation
16.90/0.6630

(b) Network output
23.51/0.7230

(c) Fine-tuned output
25.60/0.7497

Experiment #2: artifacts produced by GANs. While GANs have attained new performance in SR restoration, their produced images are often accompanied with generated artifacts. We aim to remove these undesired artifacts.

(a) Bicubic interpolation
17.67/0.5411

(b) Network output
21.62/0.6892

(c) Fine-tuned output
22.04/0.7183

Discussion

Our fine-tuning method is successful at removing artifacts at test time. Two major advantages are (1) its efficiency as it does not require pre-processing a dataset or re-training and (2) being agnostic to the type of inverse problem of artifacts.