Supplementary Material for the Manuscript: Large Inpainting of Face Images with Trainlets

Jeremias Sulam, *Student Member, IEEE,* and Michael Elad *Fellow, IEEE*

I. Introduction

This report provides supplementary material for the accompanying paper, expanding on the details of our problem formulation and providing further experimental results of face image inpainting.

II. Effect of regularizing parameter λ

Recalling the formulation of our inpainting problem, as detailed in the paper, the reconstruction of the missing image content is driven by the following minimization objective:

$$\min_{\mathbf{x}} \ \|\mathbf{y} - \mathbf{M}\mathbf{D}\mathbf{x}\|_2 + \lambda\|\mathbf{x}\|_1, \tag{1}$$

where $\mathbf{y}$ is the input image (corrupted by the mask $\mathbf{M}$), $\mathbf{D}$ is the dictionary and $\mathbf{x}$ is the (sparse) representation vector of the image. The parameter λ is a penalty (or regularization) parameter, providing a trade-off between the ℓ_2 term and the sparse enforcing term (ℓ_1). In the particular case of image inpainting, λ influences the quality of the final reconstruction: a very low value of this parameter provides a low reconstruction error (on the available image region), but promotes dense solutions. These solutions are not likely under the trained model and therefore yield unnaturally looking results, often with severe artifacts. A high value of lambda promotes very sparse solutions, but at the expense of a higher reconstruction error. These cases provide results with a poor definition of face features.

In what follows, we intend to provide some empirical evidence of this effect. For each example (in the left column), we run our inpainting formulation employing FISTA [3], given the trained dictionary $\mathbf{D}$ obtained with Trainlets [4]. We employ 5 increasing values of λ, from left to right, in the range $[0.05, 50]$. As can be seen, a too small penalty term induces artifacts, while high values of this parameter yields blurry results.

Fig. 1. Inpainting of the image on the left column, for increasing values of λ (from left to right) in the range [0.05,50], with Trainlets.

III. More Experimental Results

In this section, we provide further experimental results, comparing our approach based on Trainlets, with the patch-propagation algorithm of [1] and two global models: PCA and the Separable Dictionary Learning of [2]. In light of the previous section, and as commented in the paper, we have tuned the λ parameter for each competing algorithm and image separately for the sake of fairness.

Fig. 2. Inpainting results. From left to right: masked image, patch propagation [1], PCA, SEDIL [2], Trainlets [4], and the original image.

Fig. 3. Inpainting results. From left to right: masked image, patch propagation [1], PCA, SEDIL [2], Trainlets [4], and the original image.

REFERENCES

[1] Z. Xu and J. Sun, "Image inpainting by patch propagation using patch sparsity," *Image Processing, IEEE Transactions on*, vol. 19, no. 5, pp. 1153–1165, 2010. 1, 2, 3

[2] S. Hawe, M. Seibert, and M. Kleinsteuber, "Separable dictionary learning," *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 438–445, 2013. 1, 2, 3

[3] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM journal on imaging sciences*, vol. 2, no. 1, pp. 183–202, 2009. 1

[4] J. Sulam, B. Ophir, M. Zibulevsky, and M. Elad, "Trainlets: Dictionary learning in high dimensions," *IEEE Transactions on Signal Processing*, vol. 64, no. 12, pp. 3180–3193, 2016. 1, 2, 3