

362
363

Deep Optimal Transport: A Practical Algorithm for Photo-realistic Image Restoration - Supplementary Material

364 A Background and extensions

365 A.1 Numerical Example

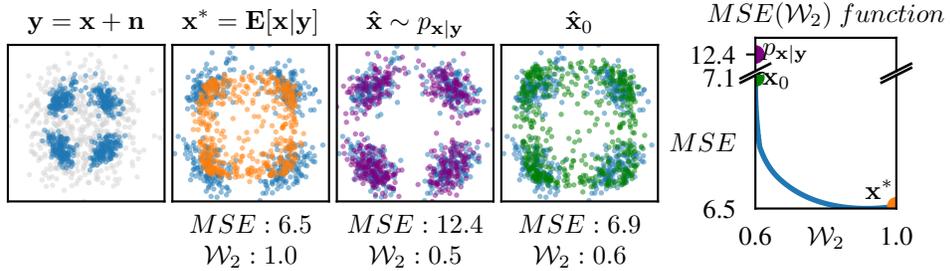


Figure 7: 2D Gaussian mixture denoising. Source samples are shown in blue. The *MMSE* estimator (\mathbf{x}^* , orange) attains the best *MSE* but the worst perceptual index \mathcal{W}_2 . The posterior samples ($\mathbf{x}|y$, purple) attain the best perceptual index but half of the optimal *MSE* performance. The \mathbf{D}_{\max} estimator ($\hat{\mathbf{x}}_0$, green) maintains the *MSE* of \mathbf{x}^* while attaining a perceptual quality close to $\mathbf{x}|y$. The DP curve is obtained by interpolating $\hat{\mathbf{x}}_0$ and \mathbf{x}^* using eq. (4).

366 To guide the reader in understanding the MMSE transport paradigm, we showcase our method on a
367 2-dimensional denoising problem. To avoid a too trivial uni-modal example, we draw the clean signal
368 from a 4-components Gaussian mixture with non-trivial covariances. We derive linear MMSE and
369 posterior estimators from [38] and proceed by applying the closed-form transport operator introduced
370 in eq. (3).

371 Note that to avoid deviating from our actual method, we refrain from using more advanced transport
372 operators better suited for multi-modal data. Indeed, those are not a practical solution for real-world
373 image datasets, as they require much more samples than actually available.

374 We summarize the experiment results in fig. 7. We observe that we obtain the best perceptual quality
375 by sampling from the posterior distribution. However, we witness a significant decrease in *MSE*
376 performance as predicted by [2]. In contrast, the \mathbf{D}_{\max} estimator enjoys a good perceptual index
377 while maintaining a close-to-optimal distortion performance.

378 A.2 Stochastic transport operator

379 Throughout our experiments, we found out that increasing the patch-size p can result in numerical
380 instabilities. Recall that the linear transport operator presented in eq. (3) uses the inverse square root
381 of the source covariance matrix $\Sigma_{\mathbf{x}_1}$. When p is large, (typically $p \geq 7$), we obtain ill-conditioned
382 covariance matrices. When the smallest singular value is still positive, we add a small stability
383 constant to the matrix diagonal to ensure it is strictly positive definite. However, the numerical errors
384 sometimes adds up to negative eigenvalues.¹ In this case, we clamp the negative eigenvalues to zero
385 and use the stochastic (one-to-many) transport operator proposed by [11],

$$\mathbf{T}_{p_{\mathbf{x}_1} \rightarrow p_{\mathbf{x}_2}}^{\text{stochastic}}(x_1) = \Sigma_{\mathbf{x}_2}^{\frac{1}{2}} \left(\Sigma_{\mathbf{x}_2}^{\frac{1}{2}} \Sigma_{\mathbf{x}_1} \Sigma_{\mathbf{x}_2}^{\frac{1}{2}} \right)^{\frac{1}{2}} \Sigma_{\mathbf{x}_2}^{-\frac{1}{2}} \Sigma_{\mathbf{x}_1}^{\dagger} (x_1 - \mu_{\mathbf{x}_1}) + \mu_{\mathbf{x}_2} + w, \quad (5)$$

386 when $\Sigma_{\mathbf{x}_1}^{\dagger}$ denotes the pseudo-inverse of $\Sigma_{\mathbf{x}_1}$ (after negative eigenvalues where clamped) and
387 $w \sim \mathcal{N}(0, \Sigma_{\mathbf{x}_2}^{\frac{1}{2}} (I - \Sigma_{\mathbf{x}_2}^{\frac{1}{2}} \mathbf{T}^* \Sigma_{\mathbf{x}_1}^{\dagger} \mathbf{T}^* \Sigma_{\mathbf{x}_2}^{\frac{1}{2}})^{\frac{1}{2}} \Sigma_{\mathbf{x}_2}^{\frac{1}{2}})$, with $\mathbf{T}^* = \Sigma_{\mathbf{x}_2}^{-\frac{1}{2}} \left(\Sigma_{\mathbf{x}_2}^{\frac{1}{2}} \Sigma_{\mathbf{x}_1} \Sigma_{\mathbf{x}_2}^{\frac{1}{2}} \right)^{\frac{1}{2}} \Sigma_{\mathbf{x}_2}^{-\frac{1}{2}}$.

¹We tried to avoid overflow when summing over the images by using 64 bit precision

388 B Practical choices and considerations in our algorithm

389 B.1 Working in latent space

390 We adopt the latent transport approach where the images are embedded into the latent space of
391 a pre-trained auto-encoder. Let $\mathbf{E}(\cdot)$, $\mathbf{D}(\cdot)$ denote the encoder and decoder, respectively. Even if
392 $\mathbf{D}(\mathbf{E}(t)) = t$, it is likely that $\mathbf{E}(\cdot)$ “deforms” the space, i.e., $\|\mathbf{E}(s) - \mathbf{E}(t)\| \neq \|s - t\|$, which means
393 that the optimal transport plan in the latent space could be *different* than the plan we seek in the pixel
394 space (the cost function in eq. (3) has changed). We can address this by modifying the latent cost
395 function to account for the deformation via the following change of variables

$$\mathbb{E} \left[\|\hat{\mathbf{x}} - \mathbf{x}\|^2 \right] = \mathbb{E} \left[\frac{\|\mathbf{E}(\hat{\mathbf{x}}) - \mathbf{E}(\mathbf{x})\|^2}{\left| \frac{\partial \mathbf{E}(\mathbf{x})}{\partial \mathbf{x}} \right| \cdot \left| \frac{\partial \mathbf{E}(\hat{\mathbf{x}})}{\partial \hat{\mathbf{x}}} \right|} \right], \quad (6)$$

396 where $\left| \frac{\partial \mathbf{E}(\mathbf{x})}{\partial \mathbf{x}} \right|$ is the determinant of the Jacobian matrix of $\mathbf{E}(\cdot)$. However it is not a practical solution
397 since we lose access to the closed-form solution eq. (2). Note that the latent MSE approximation is
398 usually desirable when dealing with natural images (e.g. to elaborate image quality measure [39],
399 perceptual quality metrics [25]). It is also true in our case but it means we can no longer claim we
400 obtain the \mathbf{D}_{\max} estimator.

401 With that, we argue that switching to a latent cost is actually a strength rather than a weakness of
402 our method. Indeed, using the MSE between deep latent variables has shown to be a better fit to
403 compare natural images than directly working in the pixel space [32]. The authors of [7] trained
404 their VAE (which is used in our experiments) to remove “imperceptible details” from the latent
405 representation, in order to better focus on higher level image semantics. In section 5.1 we validate
406 this claim by showing that our algorithm maintains the “perceptual” discrepancy performance of the
407 original estimator (e.g., LPIPS).

408 B.2 Overlapping patches extraction strategy

409 For Convolutional Neural Network (CNN) encoders², let (c, H_e, W_e) denote the shape of the latent
410 representation (CNN encoders produce 3-dimensional encoded tensors), where H_e, W_e the spatial
411 extent and c is the number of channels (i.e., the number of convolution kernels in the last convolution
412 layer). The covariance matrices $\Sigma_{\hat{\mathbf{x}}_e}, \Sigma_{\mathbf{x}_e}$ contain $\frac{(c, H_e, W_e)^2}{2}$ parameters, which may require a
413 large amount of samples for large latent images with $H_e, W_e \gg 1$. To mitigate the quadratic
414 dependency on $H_e \cdot W_e$, we assume that the latent pixels depend only on the pixels in their close
415 neighborhood. In practice, we unfold the latent representation, extracting all overlapping patches of
416 shape (c, p, p) . A similar approximation exists in the style-transfer literature [8, 9], where instead of
417 patches, only the pixels are considered (i.e., this is a private case of our approach with $p = 1$). In
418 section 5.3 we empirically show that increasing p improves the perceptual quality at the expense of
419 MSE performance, given that enough training samples are available.

420 B.3 Shared distribution

421 When dealing with natural image scenes, it is beneficial to suppose that overlapping patches share
422 common statistical attributes [39, 40]. In the case of a CNN encoded image, this approximation
423 remains satisfying because we ultimately look at filter activations which are spatial-invariant with each
424 latent patch having the same receptive field. Therefore, we assume that the overlapping patches are
425 all samples from the same distribution. This approach dramatically reduces the number of estimated
426 parameters, and also multiplies the number of samples at our disposal by $H_e \cdot W_e$, which alleviates
427 the curse of dimensionality. We demonstrate these practical benefits in section 5.3. In practice, given
428 N images, we “flatten” all the extracted patches to vectors $\underline{v}_{cp^2 \times 1}$ which we stack into a sample
429 matrix $\underline{X}_{NH_e W_e \times cp^2}$. We then aggregate the samples to compute the MVG statistics: $\mu = \underline{X}^T \mathbf{1}$,
430 $\Sigma = \frac{NH_e W_e}{NH_e W_e - 1} (\underline{X} - \mu)(\underline{X} - \mu)^T$. As $NH_e W_e$ may be very large, we perform all computations in
431 double precision. When training, this process is done twice; once for the natural image samples, and
432 once for the restored samples we wish to transport.

²This methodology can easily be extrapolated to other encoder architectures.

433 **B.4 Size of the latent representation**

434 When increasing the capacity of models with a fixed encoding rate, deepening is preferable than
 435 widening. Indeed, increasing c makes the covariance estimation dramatically harder while increasing
 436 H_e, W_e enlarges the sample pool. Therefore, the VAE from [7] with $c = 4$ and $H_e, W_e \gg 1$ is a
 437 particularly good candidate for our method. For $p = 3$ for instance, the covariance matrix admits only
 438 1296 parameters while each 512^2 image contributes 4096 samples to its estimation. As we see next,
 439 this greatly contributes to reducing the number of training samples needed to estimate the covariance
 440 matrices and allows to compute the transport operator in a few-shot manner.

441 **B.5 Transport**

442 In a single pass on a data set of natural images and a (possibly different) data set of restored
 443 samples, we compute $T_{p_{\hat{x}_e} \rightarrow p_{x_e}}^{\text{MVG}}$ (see eq. (2)). Note that each latent distribution could sometimes be
 444 degenerate, especially for severe degradations. Fortunately, the classical MVG transport operator can
 445 be generalized to ill-posed settings where $\Sigma_{\hat{x}}$ is a singular matrix (see appendix A.2).

446 **B.6 Decoding**

447 Since the transported patches overlap, we “fold” them back into a latent image $\hat{x}_{0,latent}$ by averaging.
 448 The latent image is then decoded back to the pixel space, i.e. $\hat{x}_0 = \mathbf{D}(\hat{x}_{0,latent})$. Since $\mathbf{E}(\cdot)$ is not
 449 invertible, the decoder $\mathbf{D}(\cdot)$ is used as a convenient approximation in the training domain of the
 450 auto-encoder. A corollary of this approximation is that the auto-encoder should in theory be trained
 451 on the image distribution we aim to transport, which weakens our claim to a fully blind algorithm.

452 All the steps described above are summarized in fig. 4.

453 **B.7 Transporting the degraded measurement**

454 We tried applying our algorithm on the degraded measurement directly. Indeed we observe quali-
 455 tatively and quantitatively that transporting the degraded measurement y amplifies the degradation
 (refer to fig. 8).

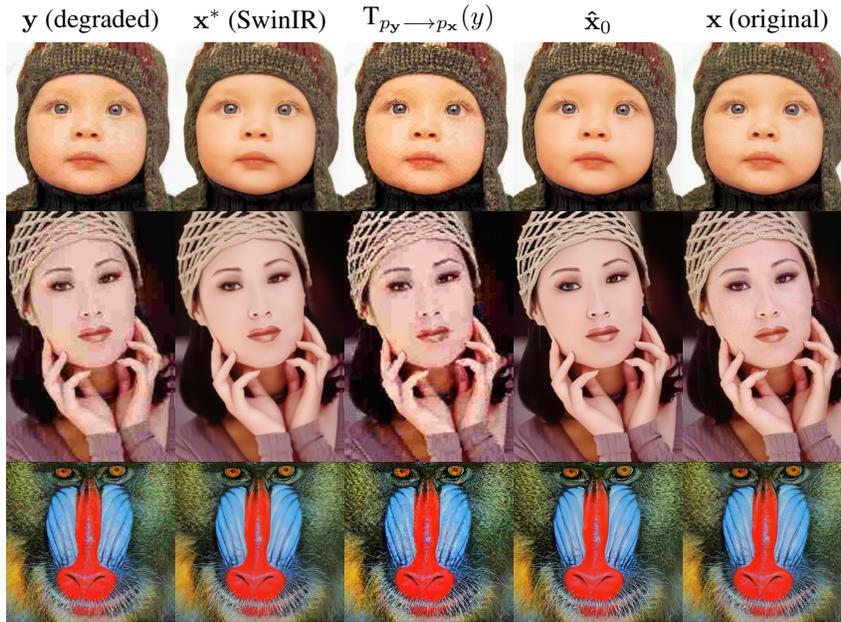


Figure 8: Transporting the degraded measurement (JPEG_{q=10}) directly is not enough to restore the image. It can sometimes even exacerbate the degradation. Quantitatively, the degraded sample y has better PSNR and FID than its transported version (respectively 27.26 dB and 13.88 FID v.s. 23.69 dB and 15.88 FID).