

Accelerating GMM-based patch priors for image restoration: Three ingredients for a 100× speed-up

Shibin Parameswaran¹, Charles-Alban Deledalle^{1,2}, Loïc Denis³ and Truong Q. Nguyen¹

¹UC San Diego, CA, USA, ²CNRS, Univ. Bordeaux, France, ³Univ Lyon, UJM-Saint-Etienne, France

Objectives

Our goal is to develop a fast and efficient image restoration algorithm utilizing Gaussian Mixture Model (GMM) patch priors. To this end, we perform:

- 1 detailed complexity analysis of Expected Patch Log-Likelihood (EPLL) algorithm
- 2 introduce innovative approximations to accelerate EPLL by a factor of 100

Introduction

We consider the problem of estimating an image $\mathbf{x} \in \mathbb{R}^N$ (N is the number of pixels) from noisy linear observations:

$$\mathbf{y} = \mathcal{A}\mathbf{x} + \mathbf{w}$$

where:

$\mathcal{A} : \mathbb{R}^N \rightarrow \mathbb{R}^M$ is a linear operator

$\mathbf{w} \in \mathbb{R}^M$ is i.i.d noise component from $\mathcal{N}(0, \sigma^2)$

The EPLL algorithm is a image restoration method that uses a Gaussian mixture model (GMM) prior on natural image patches.

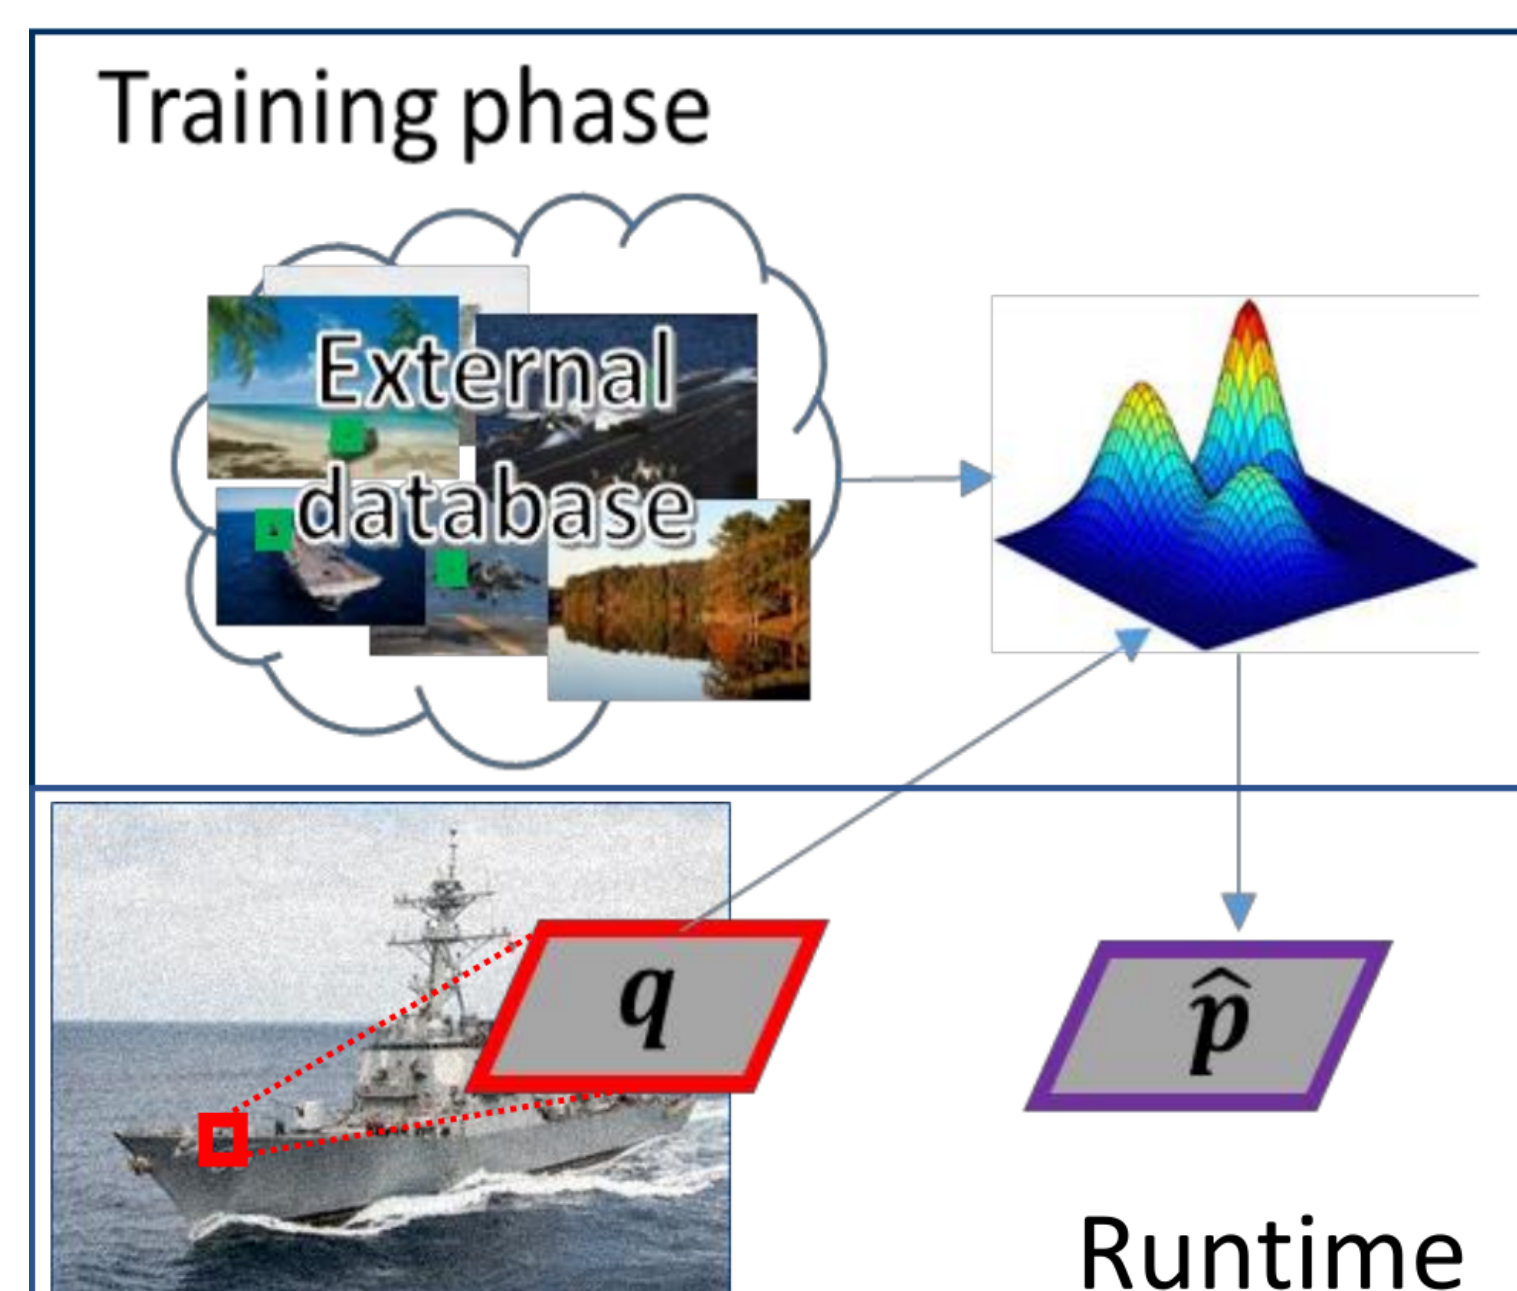
$$\operatorname{argmin}_{\mathbf{x}} \frac{P}{2\sigma^2} \|\mathcal{A}\mathbf{x} - \mathbf{y}\|^2 - \sum_{i \in \mathcal{I}} \log p(\mathcal{P}_i \mathbf{x})$$

where

$\mathcal{I} = \{1, \dots, N\}$ is the set of pixel indices

$\mathcal{P}_i : \mathbb{R}^N \rightarrow \mathbb{R}^P$ is the patch extractor at pixel i .

$p(\cdot) = \sum_{k=1}^K w_k \mathcal{N}(0, \Sigma_k)$ is zero-mean GMM prior



EPLL Solution


Use Half Quadratic Splitting strategy

$$\operatorname{argmin}_{\mathbf{x}, \mathbf{z}_1, \dots, \mathbf{z}_N} \frac{P}{2\sigma^2} \|\mathcal{A}\mathbf{x} - \mathbf{y}\|^2 + \frac{\beta}{2} \sum_{i \in \mathcal{I}} \|\mathcal{P}_i \mathbf{x} - \mathbf{z}_i\|^2 - \sum_{i \in \mathcal{I}} \log p(\mathbf{z}_i) \quad (1)$$

Solve (1) by an alternating optimization scheme:

$$\begin{cases} \hat{\mathbf{z}}_i \leftarrow \operatorname{argmin}_{\mathbf{z}_i} \frac{\beta}{2} \|\mathcal{P}_i \hat{\mathbf{x}} - \mathbf{z}_i\|^2 - \log p(\mathbf{z}_i) \\ \hat{\mathbf{x}} \leftarrow \operatorname{argmin}_{\mathbf{x}} \frac{P}{2\sigma^2} \|\mathcal{A}\mathbf{x} - \mathbf{y}\|^2 + \frac{\beta}{2} \sum_{i \in \mathcal{I}} \|\mathcal{P}_i \mathbf{x} - \hat{\mathbf{z}}_i\|^2 \end{cases}$$

Main steps of EPLL shown in Algorithm 1.




More than 100× speed-up obtained due to the proposed accelerations

Algorithm 1 The five steps of an EPLL iteration	Without accelerations	With the proposed accelerations
for all $i \in \mathcal{I}$		
$\hat{\mathbf{z}}_i \leftarrow \mathcal{P}_i \mathbf{x}$ (Patch extraction)	0.46s 1 %	0.03s 7 %
$k_i^* \leftarrow \operatorname{argmin}_{1 \leq k_i \leq K} \log w_{k_i}^{-2} + \log \Sigma_{k_i} + \frac{1}{\beta} \operatorname{Id}_P + \hat{\mathbf{z}}_i^T (\Sigma_{k_i} + \frac{1}{\beta} \operatorname{Id}_P)^{-1} \hat{\mathbf{z}}_i$ (Gaussian selection)	43.53s 95 %	0.23s 66 %
$\hat{\mathbf{z}}_i \leftarrow (\Sigma_{k_i^*} + \frac{1}{\beta} \operatorname{Id}_P)^{-1} \Sigma_{k_i^*} \hat{\mathbf{z}}_i$ (Patch estimation)	0.95s 2 %	0.05s 13 %
$\tilde{\mathbf{x}} \leftarrow (\sum_{i \in \mathcal{I}} \mathcal{P}_i^T \mathcal{P}_i)^{-1} \sum_{i \in \mathcal{I}} \mathcal{P}_i^T \hat{\mathbf{z}}_i$ (Patch reprojection)	0.23s 1 %	0.01s 4 %
$\hat{\mathbf{x}} \leftarrow (\mathcal{A}^T \mathcal{A} + \beta \sigma^2 \operatorname{Id}_N)^{-1} (\mathcal{A}^T \mathbf{y} + \beta \sigma^2 \tilde{\mathbf{x}})$ (Others)	0.52s 1 %	0.03s 10 %
Total	45.69s	0.35s

Complexity Reduction

$$\mathcal{O}(NP^2K) \rightarrow \mathcal{O}(NP\bar{r} \log K/s^2)$$

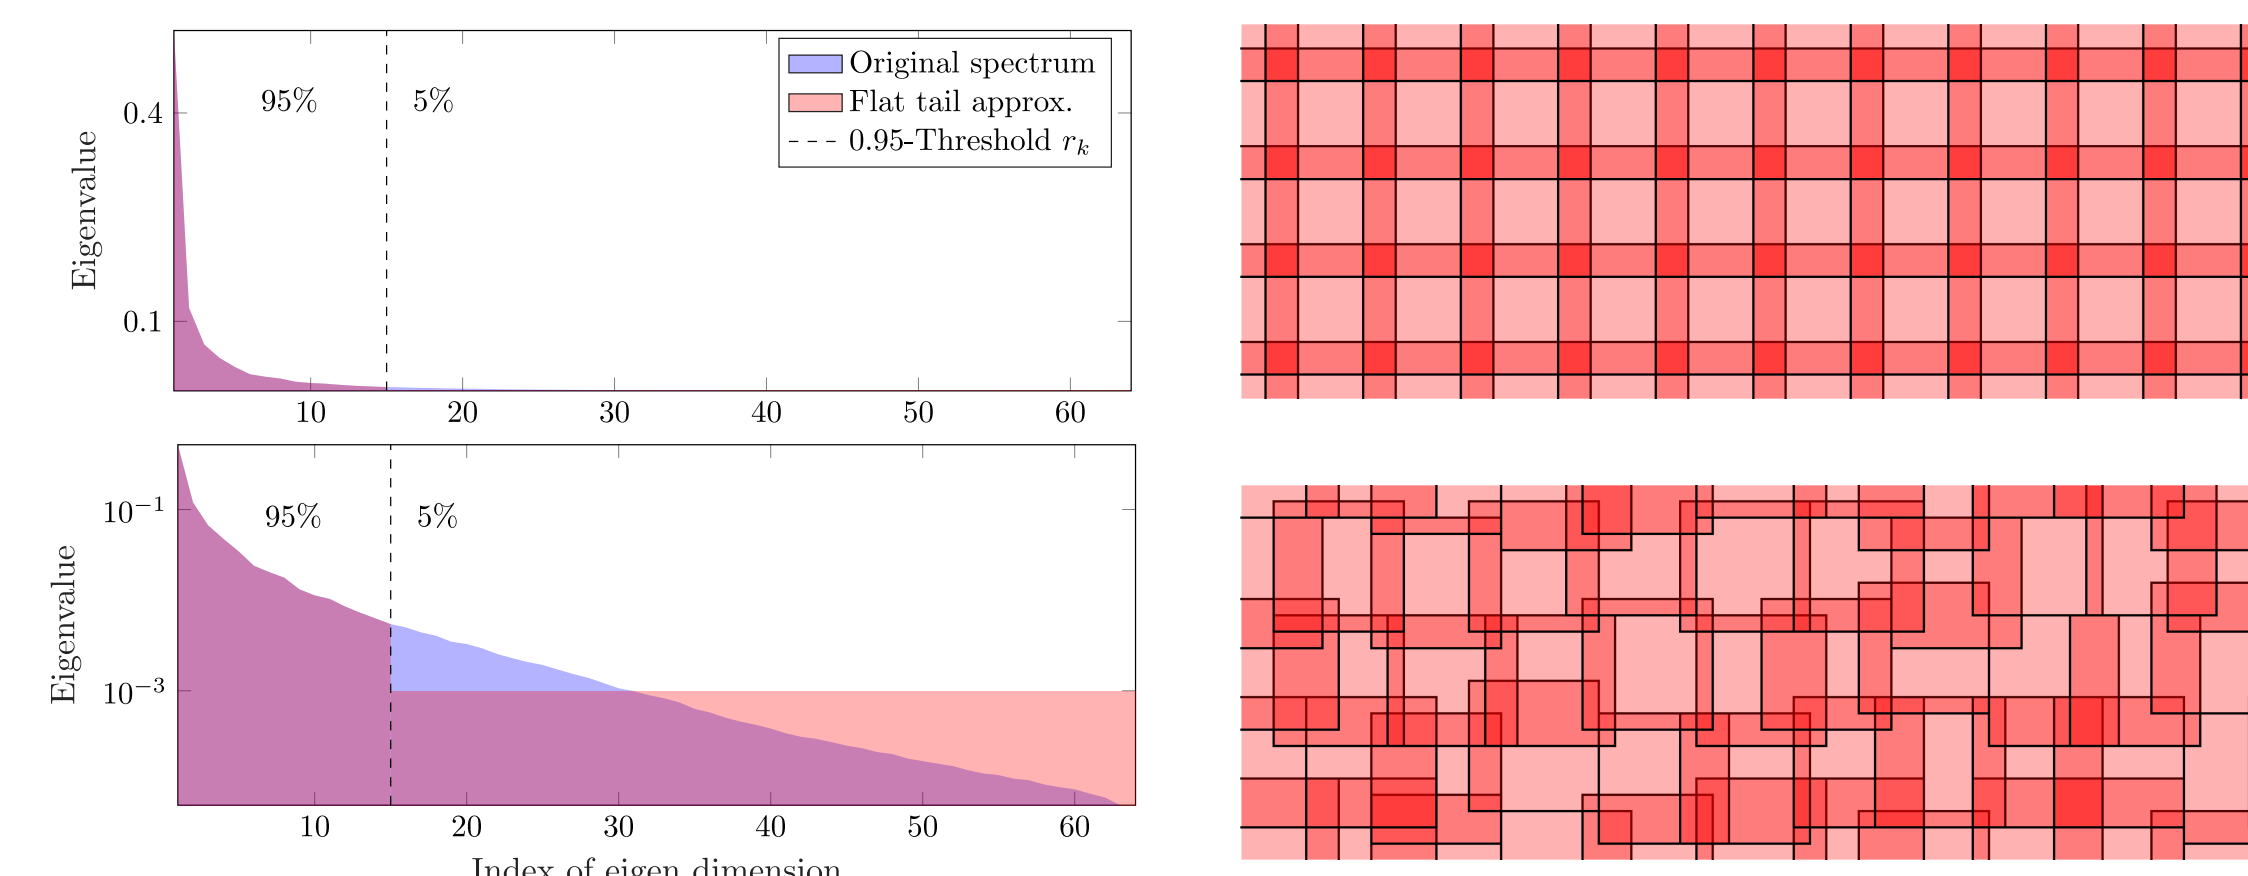
Acceleration strategies



Binary search tree
 Gaussian selection: $\mathcal{O}(NKP^2)$
 Patch estimation: $\mathcal{O}(NP^2)$
 Patch extraction: $\mathcal{O}(NP)$
 Patch reprojection: $\mathcal{O}(NP)$

Flat trail approximation
 Gaussian selection: $\mathcal{O}(NKP^2)$
 Patch estimation: $\mathcal{O}(NP^2)$
 Patch extraction: $\mathcal{O}(NP)$
 Patch reprojection: $\mathcal{O}(NP)$

Random subsampling
 Gaussian selection: $\mathcal{O}(NKP^2)$
 Patch estimation: $\mathcal{O}(NP^2)$
 Patch extraction: $\mathcal{O}(NP)$
 Patch reprojection: $\mathcal{O}(NP)$

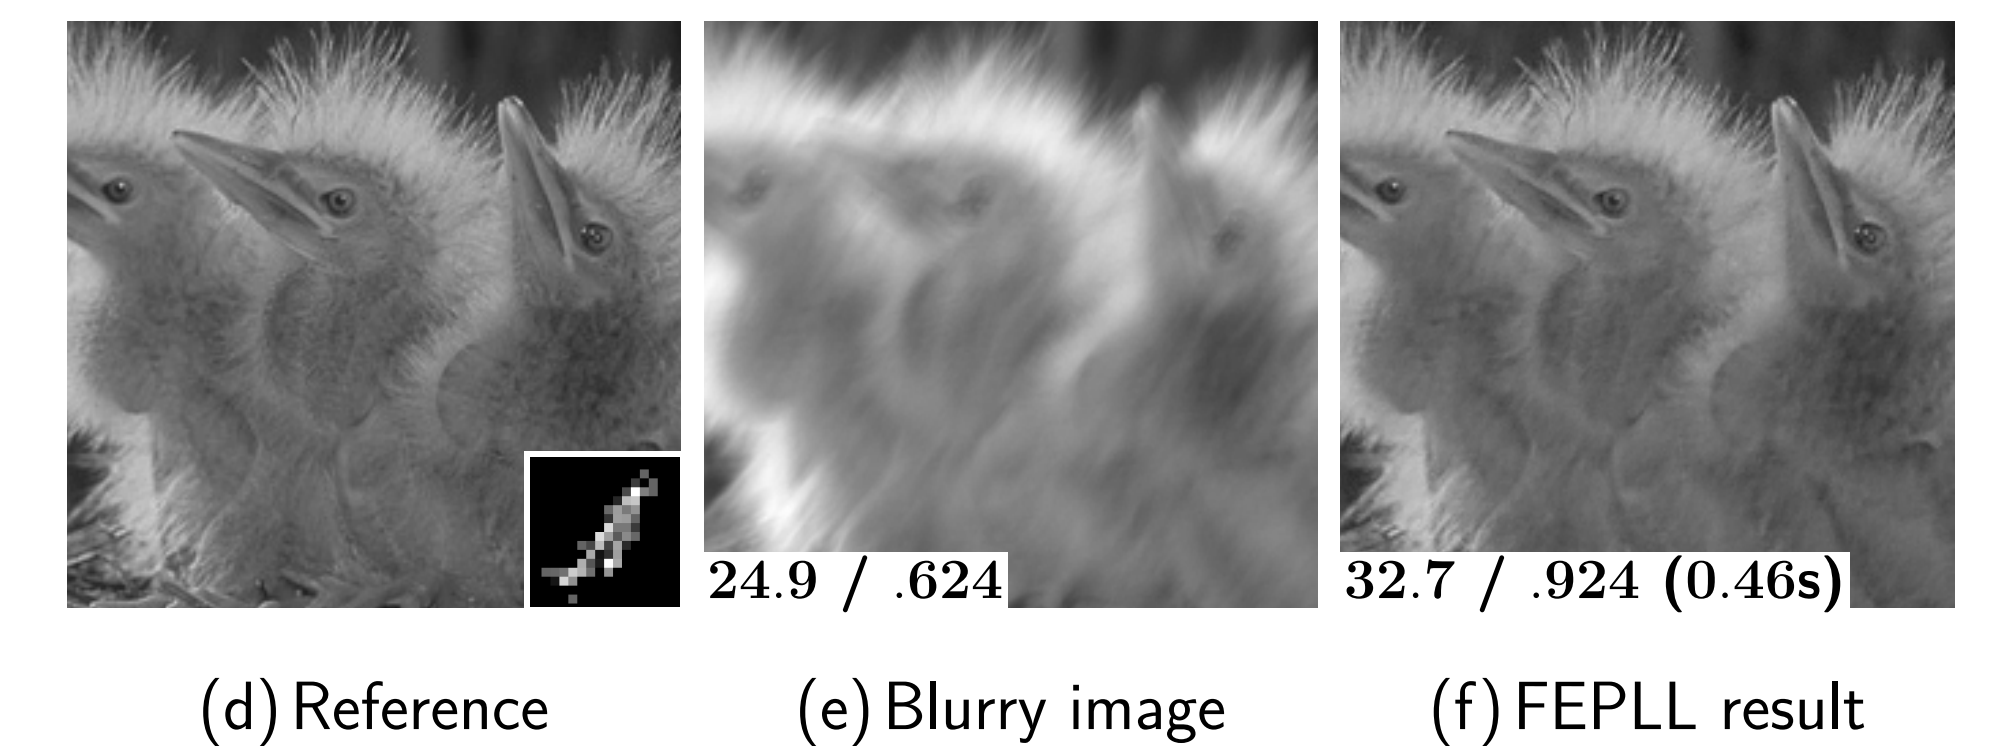


Results

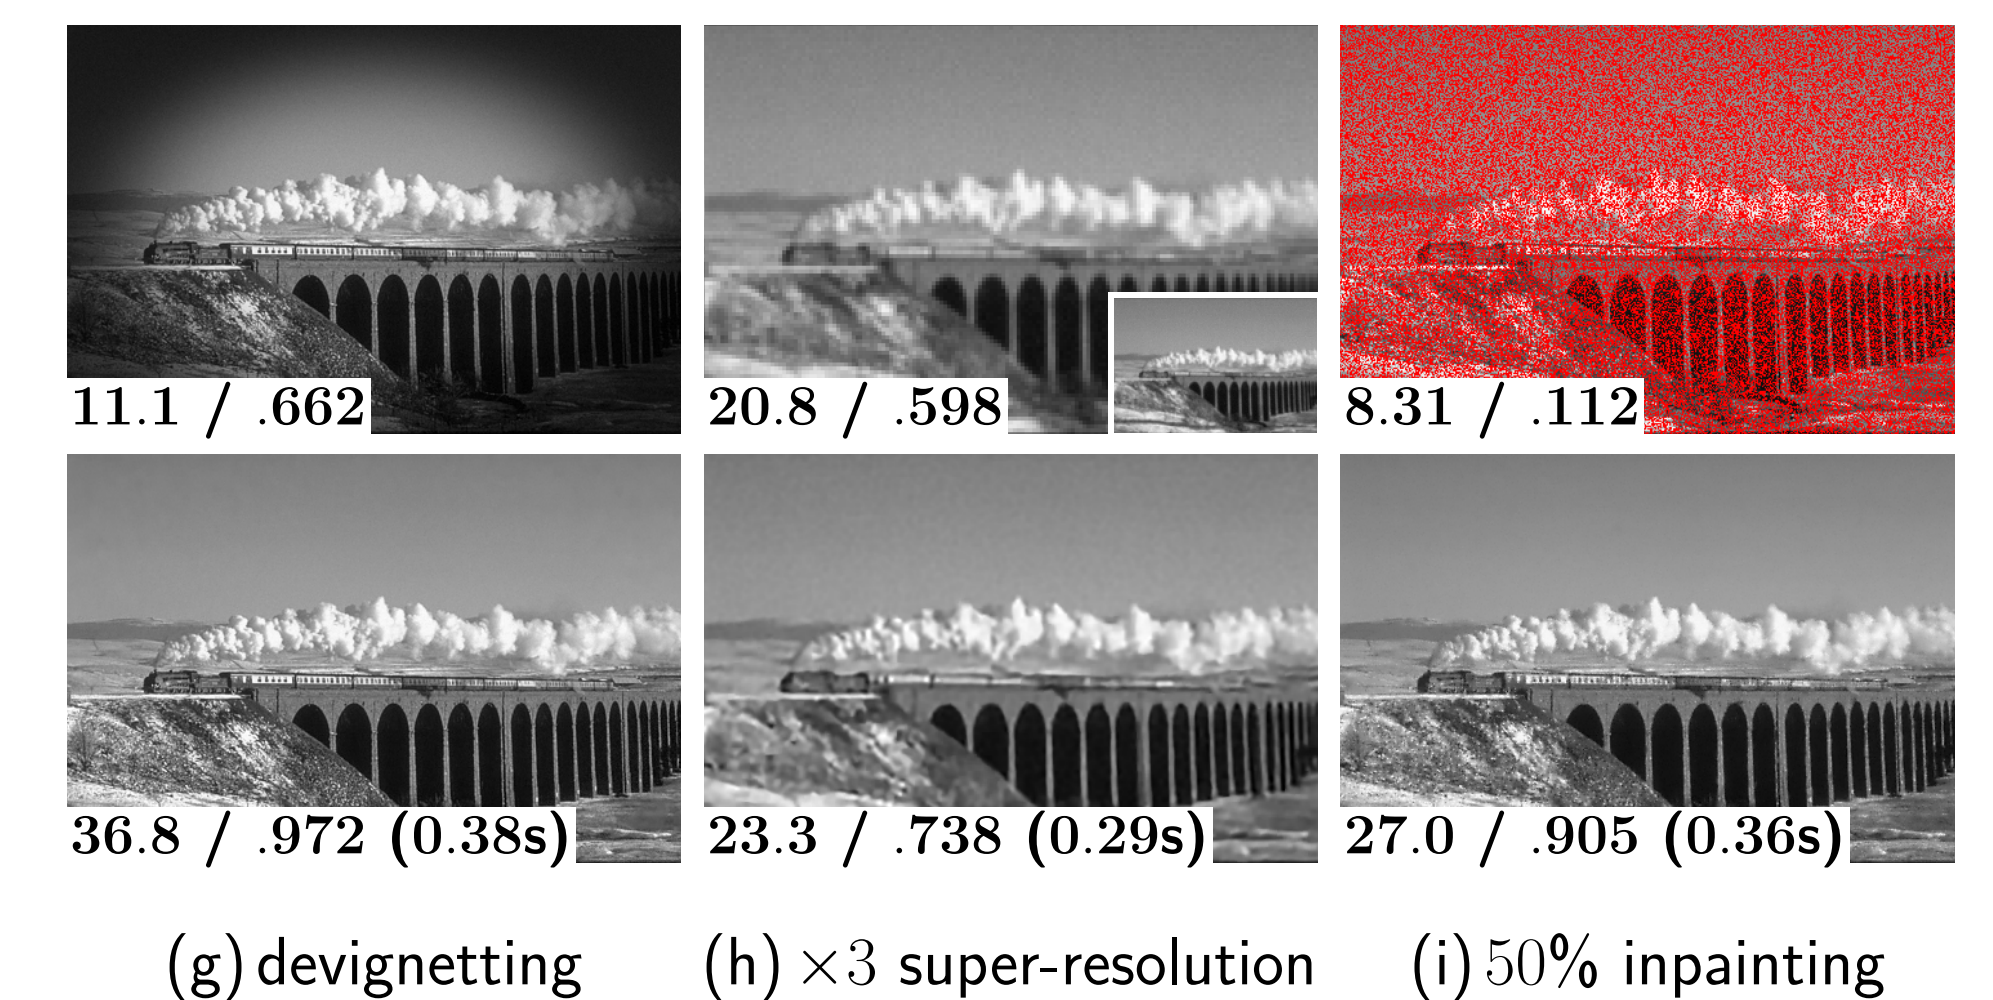
Denoising:



Deblurring:



Other inverse problems:



Conclusion

We accelerate EPLL by a factor greater than 100 with negligible loss of image quality ($<0.5\text{dB}$). The speed-up is achieved solely by reducing the algorithmic complexity of EPLL. The genericity of our acceleration strategies makes the algorithm applicable to different inverse problems without re-training.

References

- [1] Zoran, Daniel and Weiss, Yair. From learning models of natural image patches to whole image restoration. In *International Conference on Computer Vision*, pages 479–486. IEEE, November 2011.