AN ITERATIVE DENOISING AND BACKWARD PROJECTIONS METHOD

AND ITS ADVANTAGES FOR

- INPAINTING
- DEBLURRING
- BLIND DEBLURRING
- …

Tom Tirer & Raja Giryes

ICIP 2018

Method’s highlights

- Inspired by the Plug-and-Play (P&P) framework
  - Solving general imaging inverse problems using existing denoising methods
  - The image prior is specified implicitly by the chosen denoiser
- Not based on variable splitting and ADMM/HQS like P&P
- Less parameters than P&P
  + Automatic tuning mechanism!
- Similar computational cost per iteration to P&P (often requires less iterations)

Venkatakrishnan et al., “Plug-and-play priors for model based reconstruction”, 2013

The model

- The problem of image restoration can be generally formulated by

$$\boldsymbol{y} = \boldsymbol{H}\boldsymbol{x} + \boldsymbol{e}$$

$\boldsymbol{x} \in \mathbb{R}^n$ represents the unknown original image

$\boldsymbol{y} \in \mathbb{R}^m$ represents the observations

$\boldsymbol{H}$ is an $m \times n$ degradation matrix (known)

$\boldsymbol{e} \sim \mathcal{N}(\mathbf{0}, \sigma_e^2 \boldsymbol{I}_m)$ is the measurement noise

The model

- The problem of image restoration can be generally formulated by

$$y = Hx + e$$

- $H = I_n$: denoising

- H is a selection of m rows of I_n : inpainting

- H is a blurring operator : deblurring

The model

- The typical cost function

$$f(\tilde{\boldsymbol{x}}) = \frac{1}{2\sigma_e^2}\|\boldsymbol{y} - \boldsymbol{H}\tilde{\boldsymbol{x}}\|_2^2 + s(\tilde{\boldsymbol{x}})$$

- $s(\boldsymbol{x})$ is a prior image model, required in order to successfully estimate $\boldsymbol{x}$ from the observations $\boldsymbol{y}$

Iterative Denoising & Backward Projections (IDBP)

- Typical cost function (assume m<n):

$$f(\tilde{\boldsymbol{x}}) = \frac{1}{2\sigma_e^2}\|\boldsymbol{y} - \boldsymbol{H}\tilde{\boldsymbol{x}}\|_2^2 + s(\tilde{\boldsymbol{x}})$$

- Equivalent:

$$\begin{aligned} f(\tilde{\boldsymbol{x}}) &= \frac{1}{2\sigma_e^2}\|\boldsymbol{H}(\boldsymbol{H}^\dagger \boldsymbol{y} - \tilde{\boldsymbol{x}})\|_2^2 + s(\tilde{\boldsymbol{x}}) \\ &= \frac{1}{2\sigma_e^2}\|\boldsymbol{H}^\dagger \boldsymbol{y} - \tilde{\boldsymbol{x}}\|_{\boldsymbol{H}^T\boldsymbol{H}}^2 + s(\tilde{\boldsymbol{x}}) \end{aligned}$$

$$\boldsymbol{H}^\dagger \triangleq \boldsymbol{H}^T(\boldsymbol{H}\boldsymbol{H}^T)^{-1}$$

$$\|\boldsymbol{u}\|_{\boldsymbol{H}^T\boldsymbol{H}}^2 \triangleq \boldsymbol{u}^T\boldsymbol{H}^T\boldsymbol{H}\boldsymbol{u}.$$

Iterative Denoising & Backward Projections (IDBP)

- Original optimization problem: $\min_{\tilde{\boldsymbol{x}}} f(\tilde{\boldsymbol{x}})$

  can be written as:

$$\min_{\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{y}}} \frac{1}{2\sigma_e^2} \|\tilde{\boldsymbol{y}} - \tilde{\boldsymbol{x}}\|^2_{\boldsymbol{H}^T \boldsymbol{H}} + s(\tilde{\boldsymbol{x}}) \quad \text{s.t.} \quad \tilde{\boldsymbol{y}} = \boldsymbol{H}^\dagger \boldsymbol{y}$$

  Note that due to the degenerate constraint:

  $\tilde{\boldsymbol{y}} = \boldsymbol{H}^\dagger \boldsymbol{y}$ (fixed)

Iterative Denoising & Backward Projections (IDBP)

- Basic idea: Loosen the variable $\tilde{y}$ in a restricted manner, which can facilitate the estimation of x.

  How?

Iterative Denoising & Backward Projections (IDBP)

- 1. replace constraint $\tilde{\boldsymbol{y}} = \boldsymbol{H}^{\dagger}\boldsymbol{y}$ with $\boldsymbol{H}\tilde{\boldsymbol{y}} = \boldsymbol{y}$

  ⇨ degrees of freedom to $\tilde{\boldsymbol{y}}$

  We get

$$\min_{\tilde{\boldsymbol{x}},\tilde{\boldsymbol{y}}} \frac{1}{2\sigma_e^2}\|\tilde{\boldsymbol{y}} - \tilde{\boldsymbol{x}}\|^2_{\boldsymbol{H}^T\boldsymbol{H}} + s(\tilde{\boldsymbol{x}}) \quad \text{s.t.} \quad \boldsymbol{H}\tilde{\boldsymbol{y}} = \boldsymbol{y}$$

- However: components of $\tilde{\boldsymbol{y}}$ in the null space of $\boldsymbol{H}$ are not controlled (unbounded)!
  May complicate the optimization with respect to $\tilde{\boldsymbol{x}}$

Iterative Denoising & Backward Projections (IDBP)

- 2. replace $\frac{1}{\sigma_e^2}\|\tilde{\boldsymbol{y}} - \tilde{\boldsymbol{x}}\|_{\boldsymbol{H}^T\boldsymbol{H}}^2$ with $\frac{1}{(\sigma_e+\delta)^2}\|\tilde{\boldsymbol{y}} - \tilde{\boldsymbol{x}}\|_2^2$

  ⇨ control $\tilde{\boldsymbol{y}}$ in the null space of $\boldsymbol{H}$

  We get

$$\min_{\tilde{\boldsymbol{x}},\tilde{\boldsymbol{y}}} \frac{1}{2(\sigma_e+\delta)^2}\|\tilde{\boldsymbol{y}} - \tilde{\boldsymbol{x}}\|_2^2 + s(\tilde{\boldsymbol{x}}) \quad \text{s.t.} \quad \boldsymbol{H}\tilde{\boldsymbol{y}} = \boldsymbol{y}$$

Iterative Denoising & Backward Projections (IDBP)

- The proposed optimization problem:

$$\min_{\tilde{\boldsymbol{x}},\tilde{\boldsymbol{y}}} \frac{1}{2(\sigma_e+\delta)^2}\|\tilde{\boldsymbol{y}}-\tilde{\boldsymbol{x}}\|_2^2 + s(\tilde{\boldsymbol{x}}) \quad \text{s.t.} \quad \boldsymbol{H}\tilde{\boldsymbol{y}} = \boldsymbol{y}$$

- δ introduces a tradeoff:
  weak data-term vs. limited (effective) feasible set of $\tilde{\boldsymbol{x}}$

- We suggest: $\delta = \underset{\tilde{\delta}}{\operatorname{argmin}} \, (\sigma_e + \tilde{\delta})^2$

$$(\star) \qquad \text{s.t.} \quad \frac{1}{\sigma_e^2}\|\boldsymbol{H}^\dagger \boldsymbol{y} - \tilde{\boldsymbol{x}}\|_{\boldsymbol{H}^T\boldsymbol{H}}^2 \geq \frac{1}{(\sigma_e+\tilde{\delta})^2}\|\tilde{\boldsymbol{y}}-\tilde{\boldsymbol{x}}\|_2^2$$

$$\forall \; \tilde{\boldsymbol{x}},\tilde{\boldsymbol{y}} \quad \text{feasible}$$

Iterative Denoising & Backward Projections (IDBP)

- Since:

$$\frac{1}{\sigma_e^2}\|\boldsymbol{H}^{\dagger}\boldsymbol{y} - \tilde{\boldsymbol{x}}\|_{\boldsymbol{H}^T\boldsymbol{H}}^2 \approx \frac{1}{(\sigma_e+\delta)^2}\|\tilde{\boldsymbol{y}} - \tilde{\boldsymbol{x}}\|_2^2$$

and $\tilde{y} = \boldsymbol{H}^{\dagger} y$ solves $\boldsymbol{H}\tilde{y} = y$

- Solving new problem $\approx$ solving original problem

Iterative Denoising & Backward Projections (IDBP)

- Optimization problem:

$$\min_{\tilde{\boldsymbol{x}},\tilde{\boldsymbol{y}}} \frac{1}{2(\sigma_e+\delta)^2}\|\tilde{\boldsymbol{y}}-\tilde{\boldsymbol{x}}\|_2^2 + s(\tilde{\boldsymbol{x}}) \quad \text{s.t.} \quad \boldsymbol{H}\tilde{\boldsymbol{y}} = \boldsymbol{y}$$

- Alternating minimization:

Denoising:
$$\tilde{\boldsymbol{x}}_k = \operatorname*{argmin}_{\tilde{\boldsymbol{x}}} \frac{1}{2(\sigma_e+\delta)^2}\|\tilde{\boldsymbol{y}}_{k-1}-\tilde{\boldsymbol{x}}\|_2^2 + s(\tilde{\boldsymbol{x}})$$

$$\Rightarrow \quad \tilde{\boldsymbol{x}}_k = \mathcal{D}(\tilde{\boldsymbol{y}}_{k-1};\sigma_e+\delta)$$

Projection:
$$\tilde{\boldsymbol{y}}_k = \operatorname*{argmin}_{\tilde{\boldsymbol{y}}} \|\tilde{\boldsymbol{y}}-\tilde{\boldsymbol{x}}_k\|_2^2 \quad \text{s.t.} \quad \boldsymbol{H}\tilde{\boldsymbol{y}} = \boldsymbol{y}$$

$$\Rightarrow \quad \tilde{\boldsymbol{y}}_k = \boldsymbol{H}^\dagger\boldsymbol{y} + (\boldsymbol{I}_n - \boldsymbol{H}^\dagger\boldsymbol{H})\tilde{\boldsymbol{x}}_k$$

IDBP – algorithm

Input: $\boldsymbol{H}, \boldsymbol{y}, \sigma_e$, denoising operator $\mathcal{D}(\cdot;\sigma)$, stopping criterion. $\boldsymbol{y} = \boldsymbol{H}\boldsymbol{x} + \boldsymbol{e}$, such that $\boldsymbol{e} \sim \mathcal{N}(\boldsymbol{0}, \sigma_e^2 \boldsymbol{I}_m)$ and $\boldsymbol{x}$ is an unknown signal whose prior model is specified by $\mathcal{D}(\cdot;\sigma)$.

Output: $\hat{\boldsymbol{x}}$ an estimate for $\boldsymbol{x}$.

Initialize: $\tilde{\boldsymbol{y}}_0 =$ some initialization, $k = 0$, δ approx. satisfying (★)

while *stopping criterion not met* **do**

$k = k + 1;$

$\tilde{\boldsymbol{x}}_k = \mathcal{D}(\tilde{\boldsymbol{y}}_{k-1}; \sigma_e + \delta);$

$\tilde{\boldsymbol{y}}_k = \boldsymbol{H}^\dagger \boldsymbol{y} + (\boldsymbol{I}_n - \boldsymbol{H}^\dagger \boldsymbol{H})\tilde{\boldsymbol{x}}_k;$

end

$\hat{\boldsymbol{x}} = \tilde{\boldsymbol{x}}_k;$

IDBP – setting δ

$$(\star) \quad \frac{1}{\sigma_e^2}\|\boldsymbol{H}^\dagger \boldsymbol{y} - \tilde{\boldsymbol{x}}\|^2_{\boldsymbol{H}^T\boldsymbol{H}} \geq \frac{1}{(\sigma_e + \tilde{\delta})^2}\|\tilde{\boldsymbol{y}} - \tilde{\boldsymbol{x}}\|_2^2$$

$\forall \; \tilde{\boldsymbol{x}}, \tilde{\boldsymbol{y}}$ feasible

Proposition 1. *Set $\delta = \tilde{\delta}$. If there exist an iteration k of IDBP that violates the following condition*

$$\frac{1}{\sigma_e^2}\|\boldsymbol{y} - \boldsymbol{H}\tilde{\boldsymbol{x}}_k\|_2^2 \geq \frac{1}{(\sigma_e + \tilde{\delta})^2}\|\boldsymbol{H}^\dagger(\boldsymbol{y} - \boldsymbol{H}\tilde{\boldsymbol{x}}_k)\|_2^2,$$

Necessary condition for δ

then $\delta = \tilde{\delta}$ also violates the condition $(\star)$

IDBP for inpainting

- $\boldsymbol{H}$ is a selection of m rows of $\boldsymbol{I}_n$
- $\tilde{\boldsymbol{y}}_k$ simply takes observed pixels from $\boldsymbol{y}$ and missing pixels from $\tilde{\boldsymbol{x}}_k$
- We have $\|\boldsymbol{y} - \boldsymbol{H}\tilde{\boldsymbol{x}}_k\|_2 = \|\boldsymbol{H}^\dagger(\boldsymbol{y} - \boldsymbol{H}\tilde{\boldsymbol{x}}_k)\|_2$
  Proposition 1 suggests using $\delta = 0$
  (If $\sigma_e = 0$ use a small positive δ)

Inpainting experiment

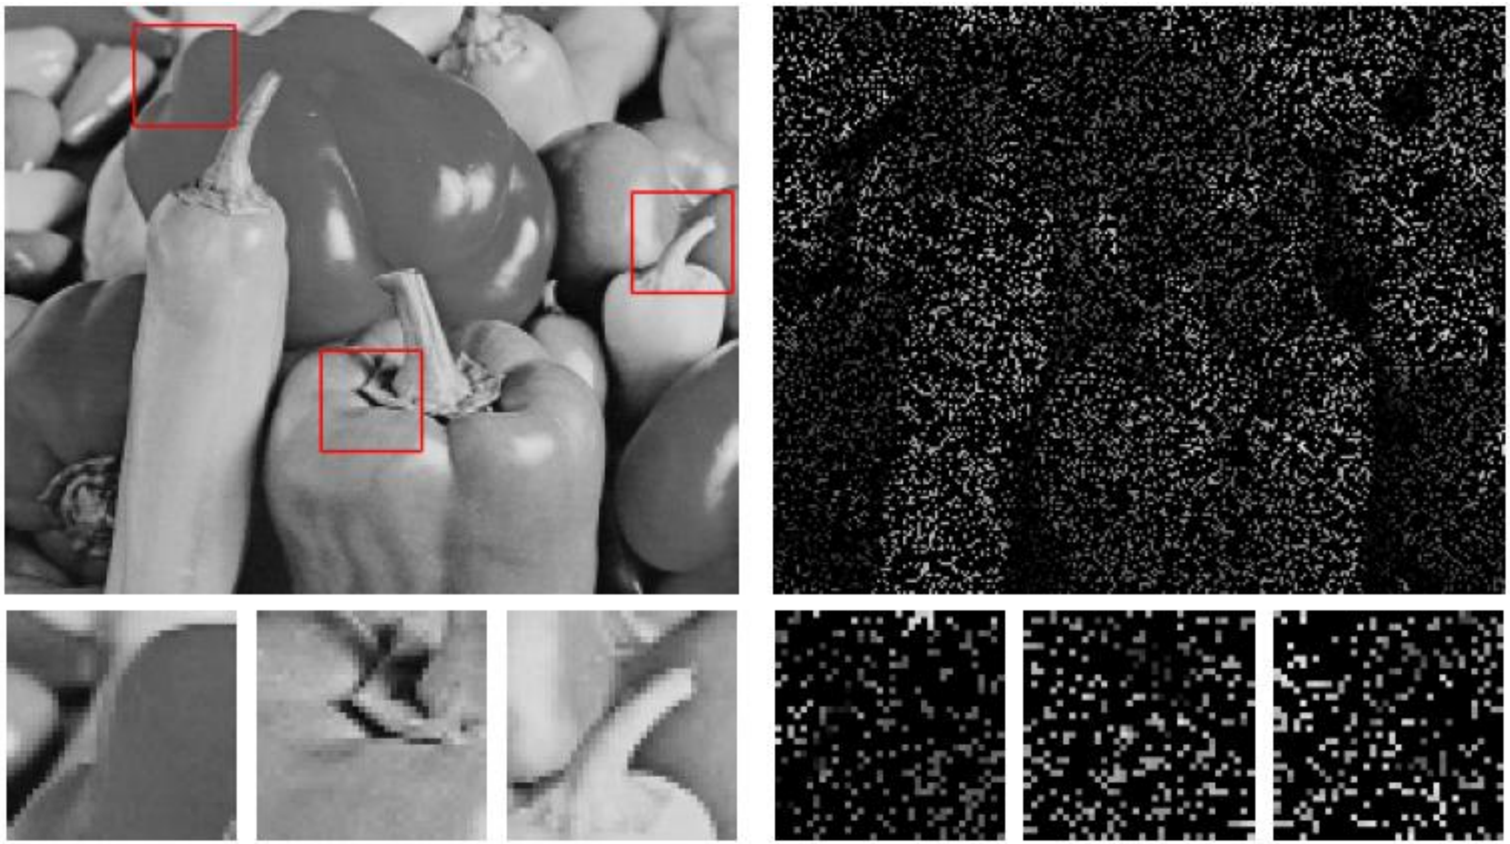

(a) Original image (b) Subsampled and noisy image

80% missing pixels

$\sigma_e = 12$

Inpainting experiment

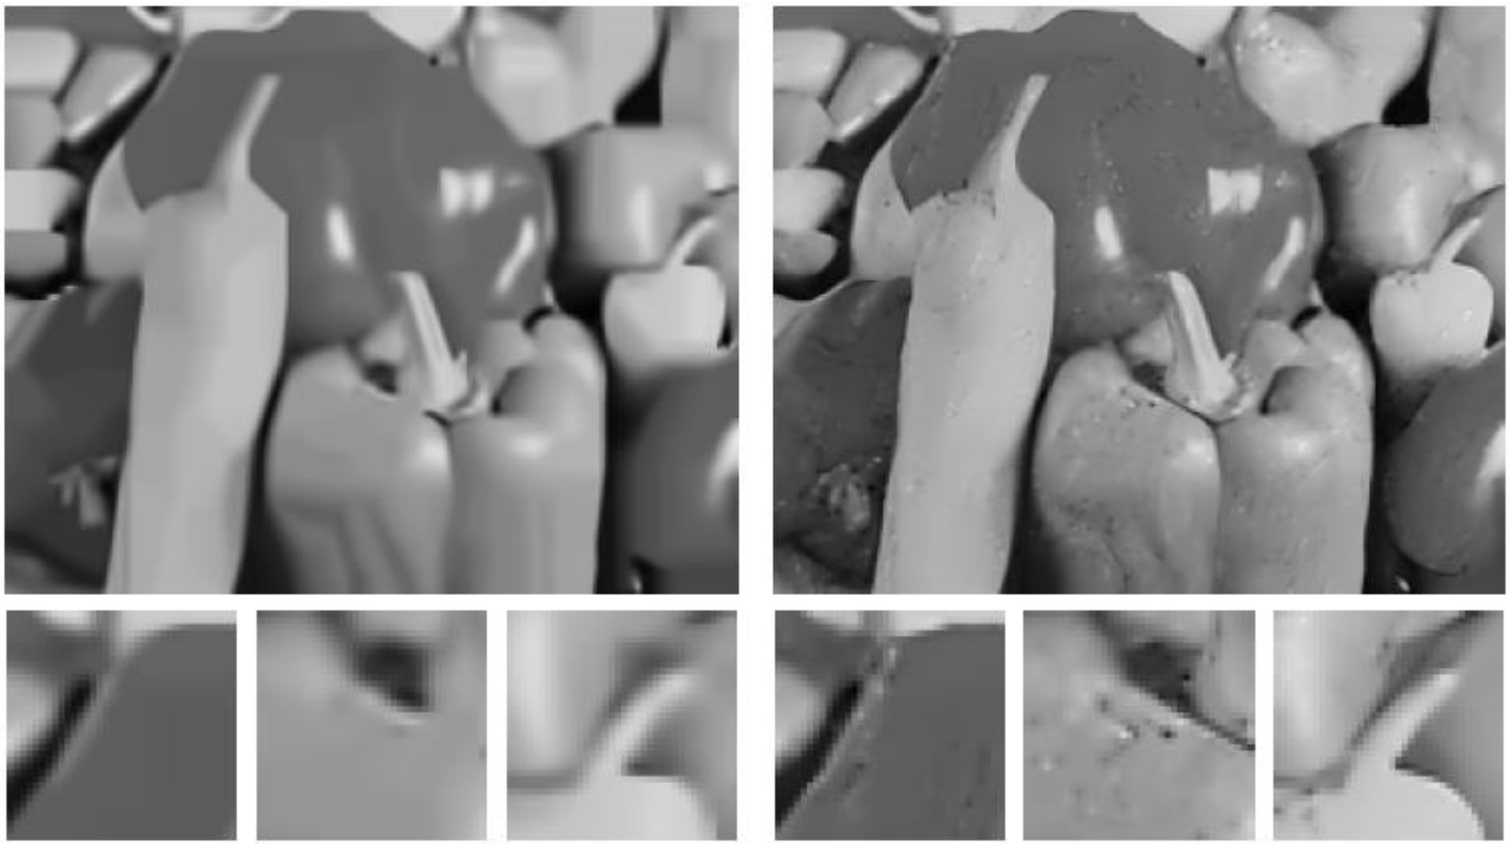

(c) P&P-BM3D (d) IRCNN

Tuning 2 params., PSNR=26.56 ~25 CNNs, PSNR=26.94

Inpainting experiment

(e) IDBP-BM3D

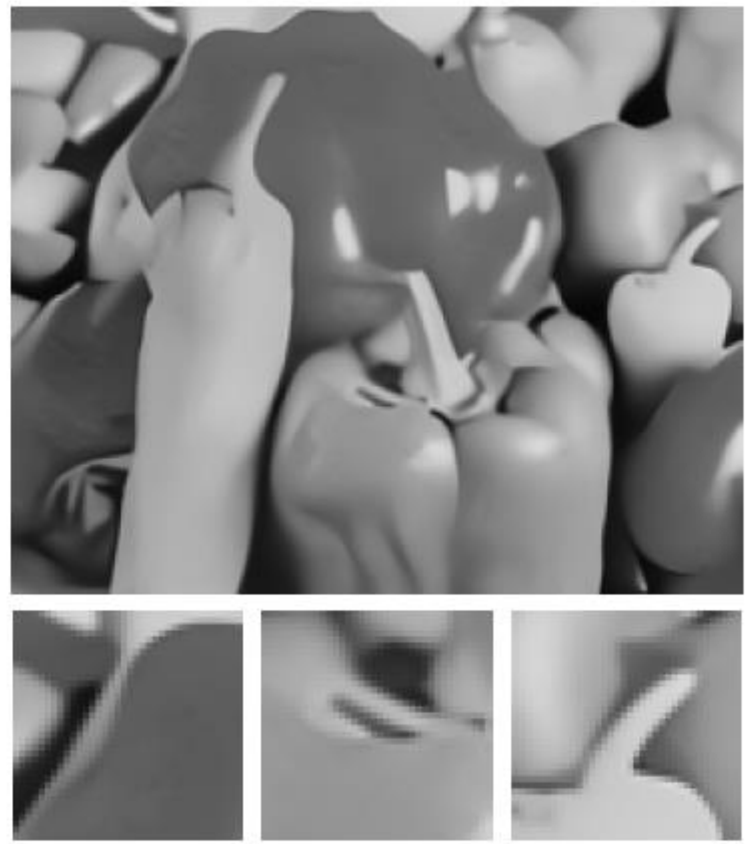

(f) IDBP-CNN

Fixed $\delta = 0$, PSNR=26.79

1 CNN, PSNR=27.17

IDBP for deblurring

- $\boldsymbol{H}^\dagger$ must be approximated, e.g. by standard Tikhonov regularization in freq. domain:

$$\tilde{\boldsymbol{g}} \triangleq \frac{\mathcal{F}^*\{\boldsymbol{h}\}}{|\mathcal{F}\{\boldsymbol{h}\}|^2 + \epsilon \cdot \sigma_e^2}$$

$$\tilde{\boldsymbol{y}}_k = \mathcal{F}^{-1}\left\{\tilde{\boldsymbol{g}}\left(\mathcal{F}\{\boldsymbol{y}\} - \mathcal{F}\{\boldsymbol{h}\}\mathcal{F}\{\tilde{\boldsymbol{x}}_k\}\right)\right\} + \tilde{\boldsymbol{x}}_k$$

- So – do we have to tune 2 parameters (δ, ϵ)?
- Note that Proposition 1 can still be computed

$$\frac{1}{\sigma_e^2}\left\|\boldsymbol{y} - \mathcal{F}^{-1}\left\{\mathcal{F}\{\boldsymbol{h}\}\mathcal{F}\{\tilde{\boldsymbol{x}}_k\}\right\}\right\|_2^2 \geq \frac{1}{(\sigma_e + \delta)^2}\left\|\mathcal{F}^{-1}\left\{\tilde{\boldsymbol{g}}\left(\mathcal{F}\{\boldsymbol{y}\} - \mathcal{F}\{\boldsymbol{h}\}\mathcal{F}\{\tilde{\boldsymbol{x}}_k\}\right)\right\}\right\|_2^2$$

RHS can be reduced by increasing δ or ϵ

IDBP – setting δ for deblurring

- Recall: LHS/RHS<1 ➡ violates Prop. 1 ➡ violates (★)
- We observed: Pairs of (δ, ϵ) that give good results indeed satisfy the condition in Prop.1 at all iterations.

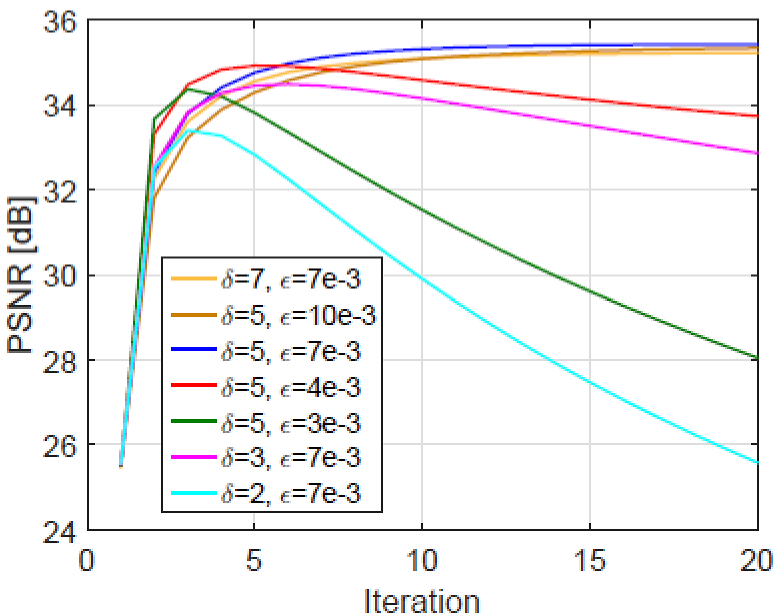

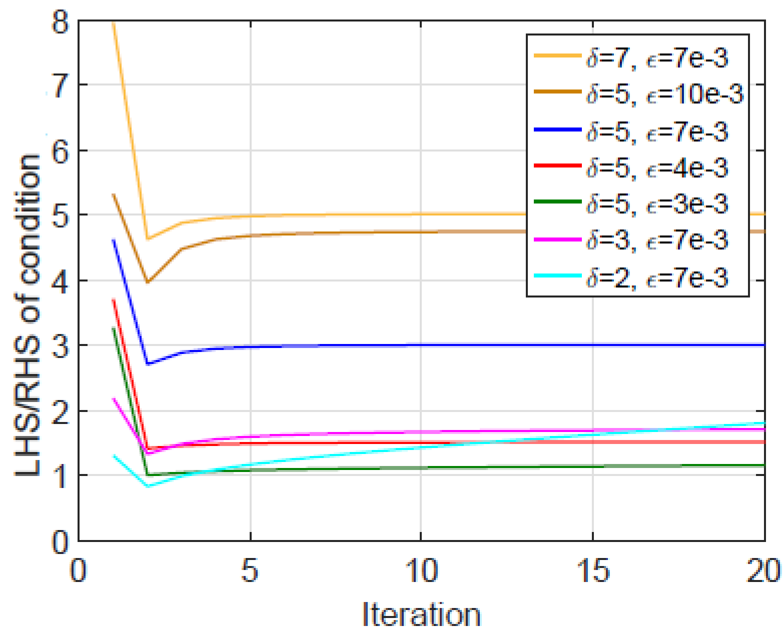

IDBP – setting δ for deblurring

- Prop. 1 can be used for automatic parameter tuning: Fix δ (e.g. $\delta = 5$) and increase ϵ until reaching some confidence margin (e.g. $\tau = 3$) for the inequality in Prop.1.

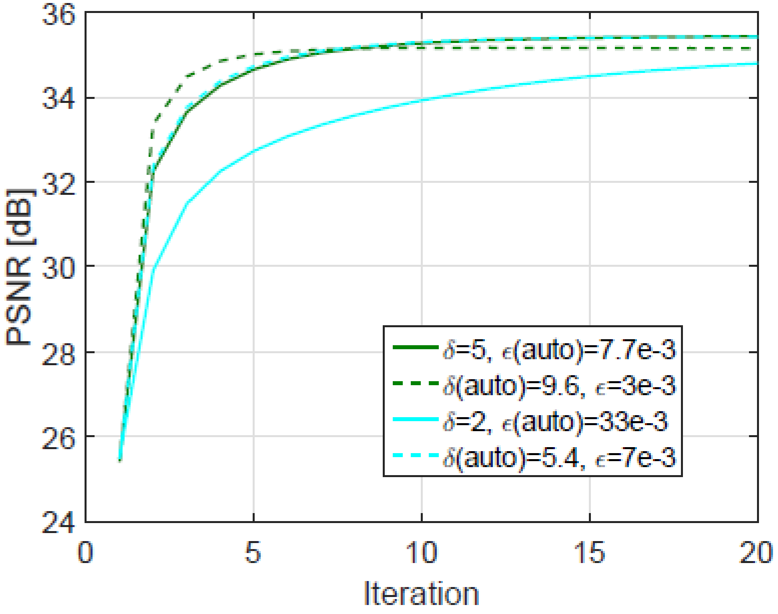

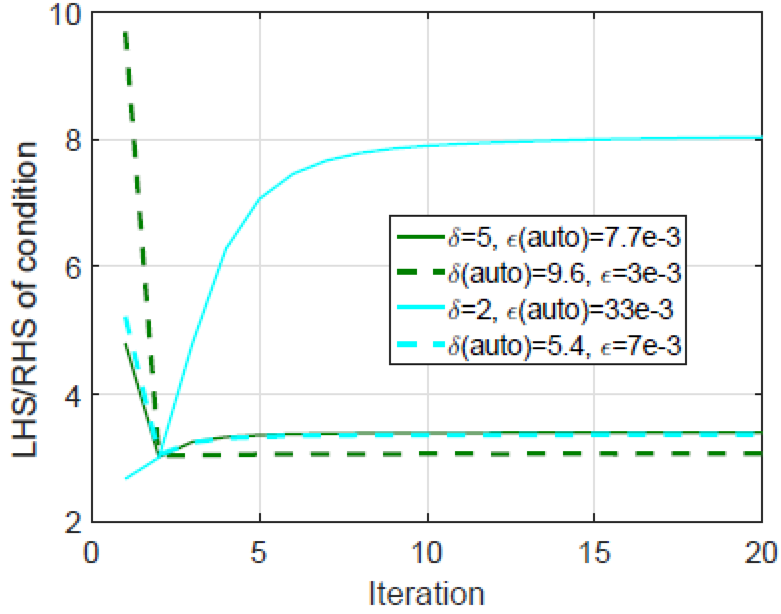

Deblurring experiment

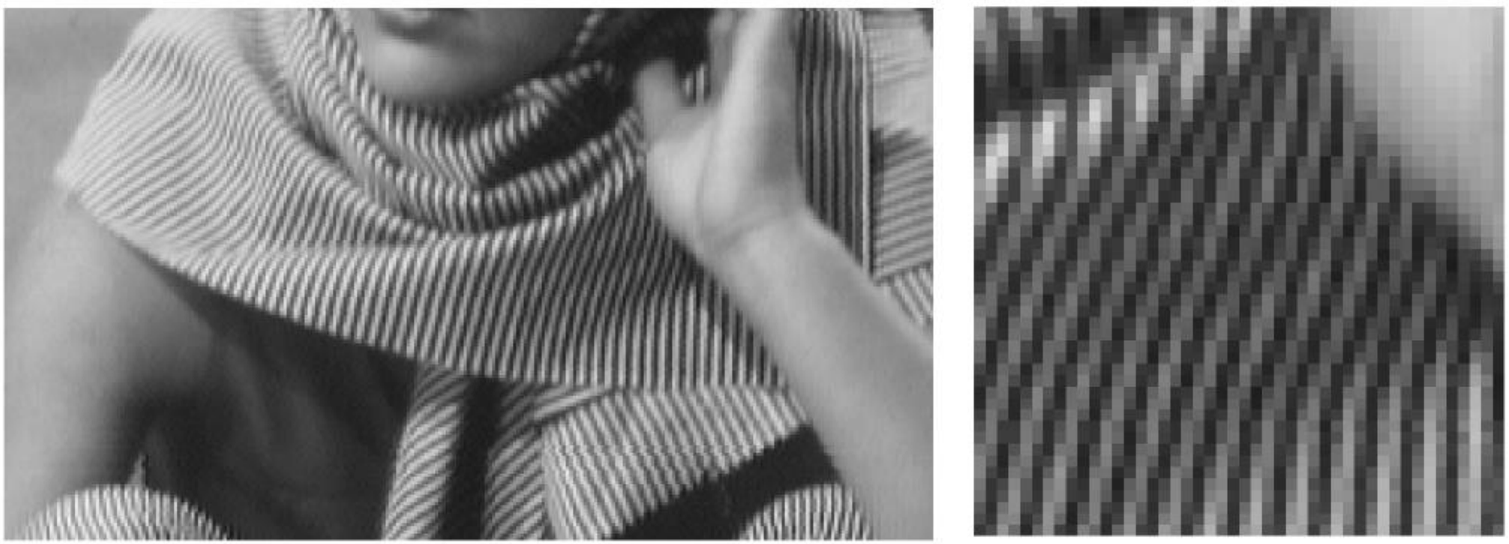

(a) Original image

Deblurring of *Barbara* image, Scenario 4

Deblurring experiment

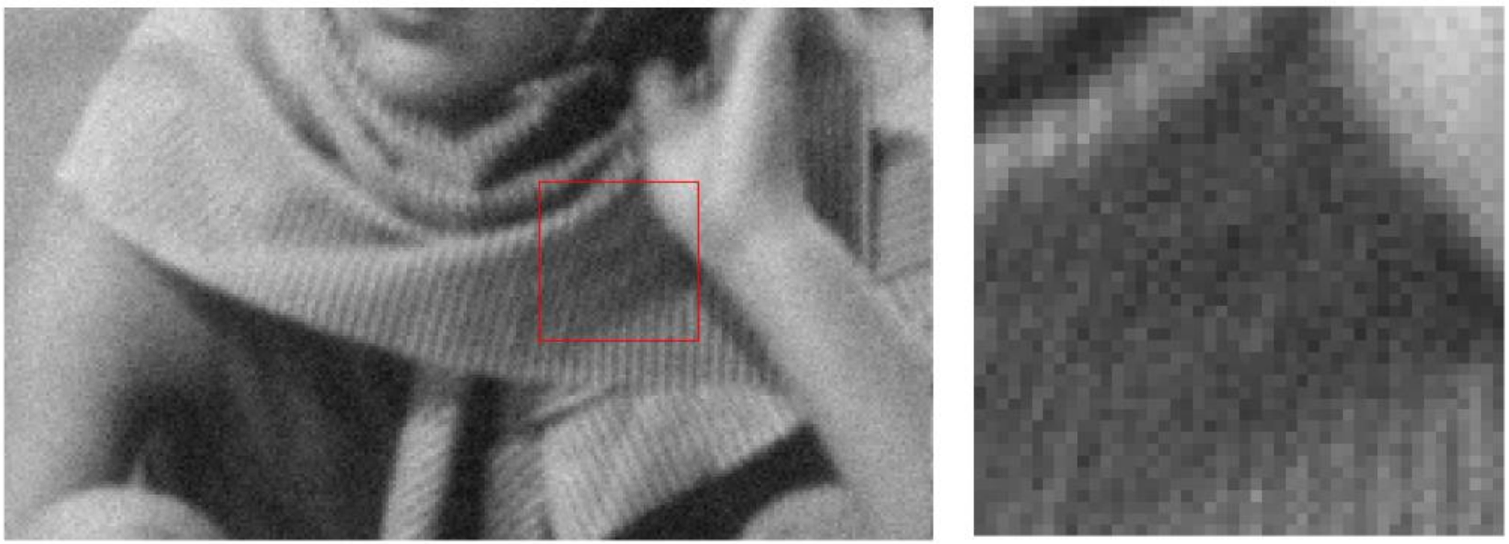

(b) Blurred and noisy image

Deblurring of *Barbara* image, Scenario 4

Deblurring experiment

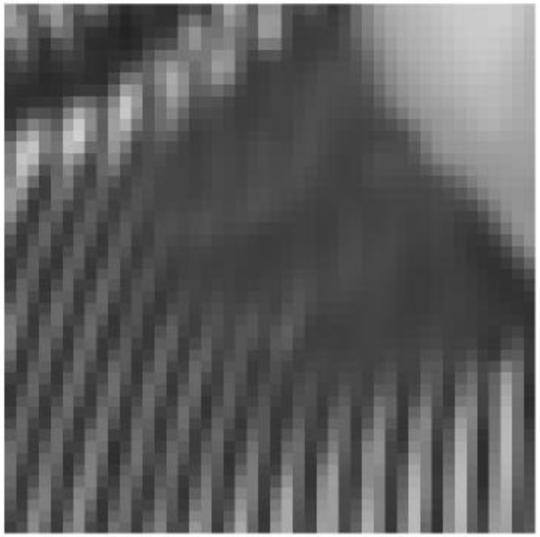

(c) IDD-BM3D

Many params., many iterations, PSNR=26.10

Deblurring of *Barbara* image, Scenario 4

Deblurring experiment

(d) P&P-BM3D

Tuning 2 params., PSNR=25.72

Deblurring of *Barbara* image, Scenario 4

Deblurring experiment

(f) Auto-tuned IDBP-BM3D

Fixed δ & tuned ϵ, PSNR=26.94

Deblurring of *Barbara* image, Scenario 4

Deblurring experiment

(a) Original image

Deblurring of *cameraman* image, Scenario 3.

Deblurring experiment

(b) Blurred and noisy image

Deblurring of *cameraman* image, Scenario 3.

Deblurring experiment

(c) IRCNN

~25 CNNs, PSNR=31.07

Deblurring of *cameraman* image, Scenario 3.

Deblurring experiment

(d) IDBP-CNN

1 CNN, PSNR=31.32

Deblurring of *cameraman* image, Scenario 3.

IDBP – advantages for blind-deblurring

- Most blind-deblurring methods:
  1. estimate only the blur kernel
  2. use non-blind deblurring
- Many non-blind deblurring algorithms require tuning per kernel (using several clean & blurry pairs)
- IDBP has automatic parameter tuning!
  Use larger confidence margin due to inexact kernel

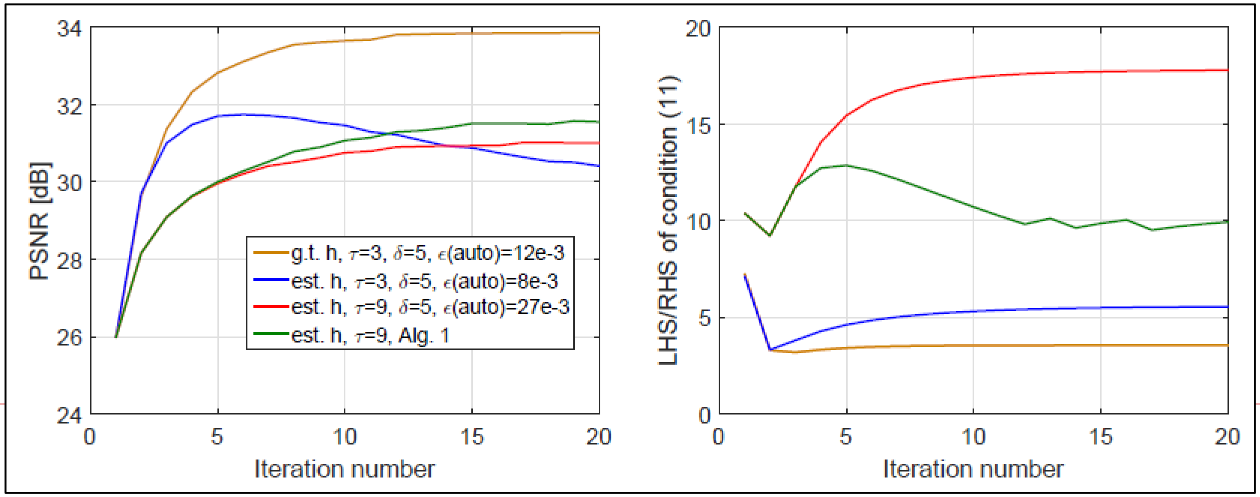

Deblurring (estimated kernel) experiment

EPLL IDBP-BM3D

Deblurring (estimated kernel) experiment

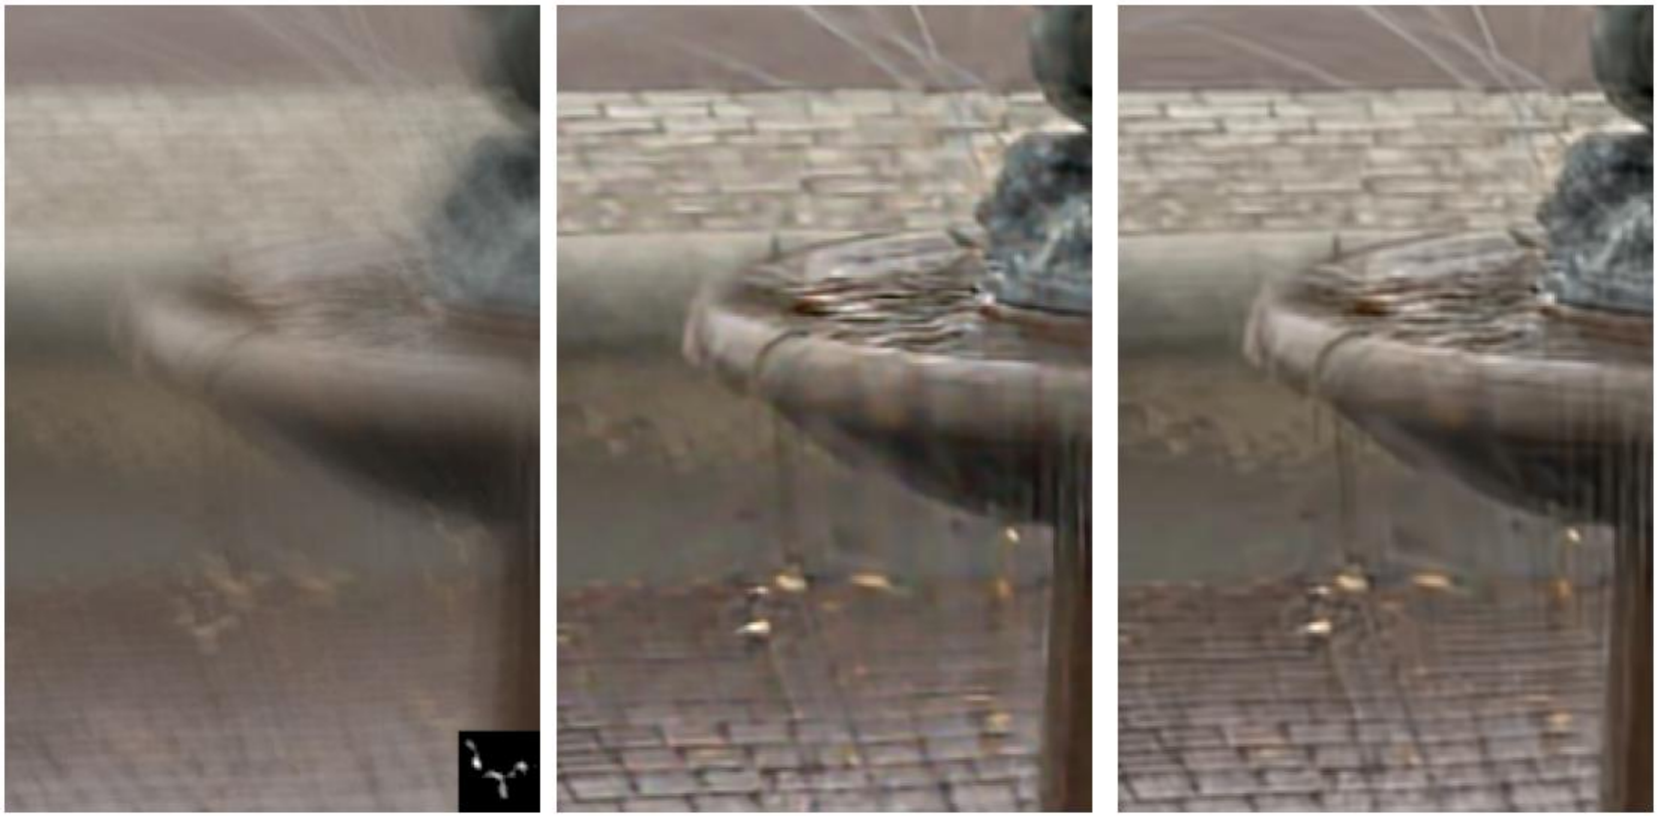

EPLL IDBP-BM3D

Thank you

Many experiments and mathematical analysis can be found in:

T. Tirer and R. Giryes, "Image Restoration by Iterative Denoising and Backward Projections," Accepted to IEEE Transactions on Image Processing, 2018.

T. Tirer and R. Giryes, "An Iterative Denoising and Backwards Projections Method and its Advantages for Blind Deblurring,‟ IEEE International Conference on Image Processing (ICIP), 2018.

Code: https://github.com/tomtirer/IDBP
https://tirertom.wixsite.com/homepage