LPA-ICI Applications in Image Processing

- Denoising
- Deblurring
- Derivative estimation
- Edge detection
- Inverse halftoning

Denoising

- Consider $z(x) = y(x) + \eta(x)$, where y is noise-free image and η is noise.
  - assume i.i.d. noise $\eta(\cdot) \sim \mathcal{N}\left(0, \sigma^2\right)$
- Denoising problem: obtain an estimate of y given z.
- Why use LPA-ICI for image denoising?
  - adapts to unknown smoothness of the signal
  - allows anisotropic (directional) estimation neighborhoods

- The discrete LPA estimation kernels g_h are pre-defined for a set of scales $h \in H$
  - the shapes of g_h are typically simple (e.g., square, disc, line)
  - preserve mean value: $\sum g_h(x) = 1$ (For derivative estimation, what should $\sum g_h(x)$ be equal to?)
  - example of simple isotropic averaging kernels for $h = \{1, 2, 3\}$,

$$g_2 = \begin{bmatrix} \mathbf{1} \end{bmatrix} \quad g_2 = \frac{1}{9}\begin{bmatrix} 1 & 1 & 1 \\ 1 & \mathbf{1} & 1 \\ 1 & 1 & 1 \end{bmatrix} \quad g_3 = \frac{1}{25}\begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & \mathbf{1} & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \end{bmatrix}$$

- The LPA estimate for each g_h is given by convolution

$$\hat{y}_h(x) = (z \circledast g_h)(x)$$

- **Adaptability: select adaptive scale** $h^+(x)$ for each pixel $x \in X$ **using the ICI rule**

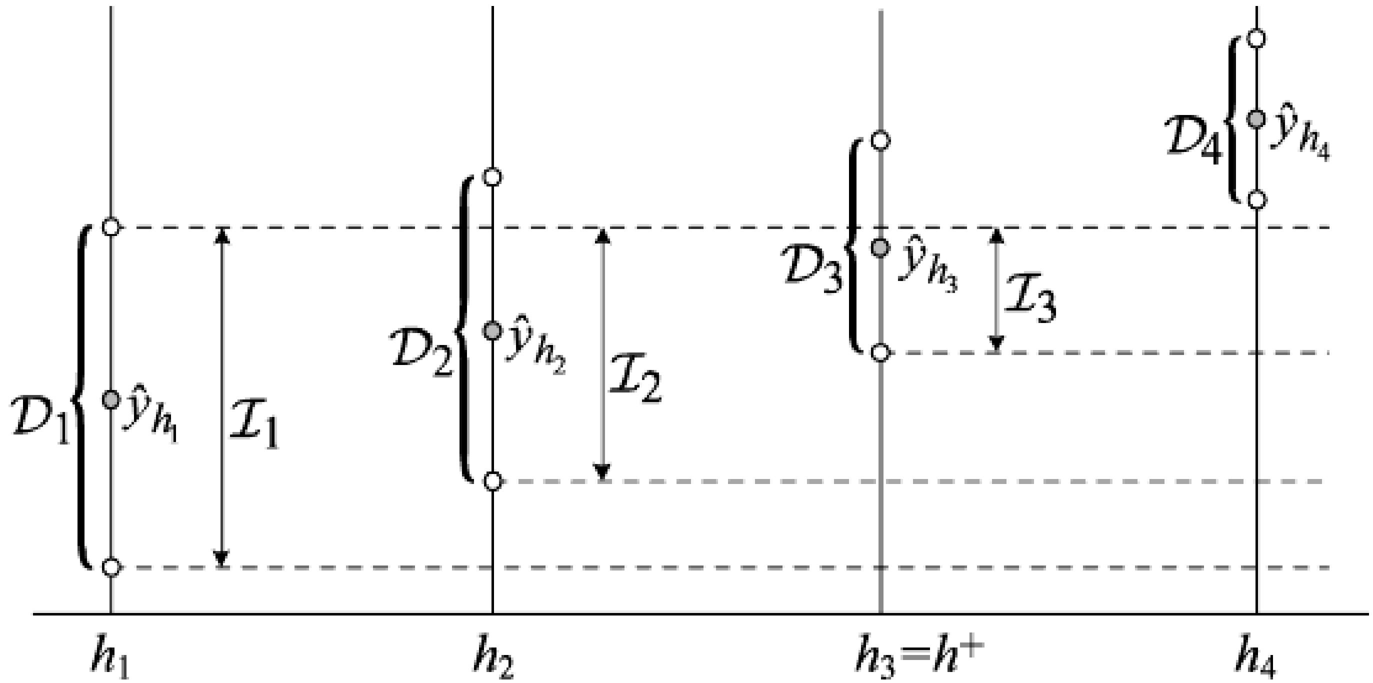

- ICI rule: Given $H = \{h_1 < \cdots < h_J\}$, choose the greatest j s.t. $\mathcal{I}_j = \bigcap_{i=1}^{j} \mathcal{D}_i$ is non-empty where $\mathcal{D}_i = \left[\hat{y}_{h_i}(x) - \Gamma\sigma_{\hat{y}_{h_i}(x)}, \hat{y}_{h_i}(x) + \Gamma\sigma_{\hat{y}_{h_i}(x)}\right]$. Result adaptive scales $h^+(x) = h_j$

- ICI yields the pointwise adaptive estimate $\hat{y}_{h^+}$
  - $h^+(x)$ is a close approximation of the optimal $h^*(x)$

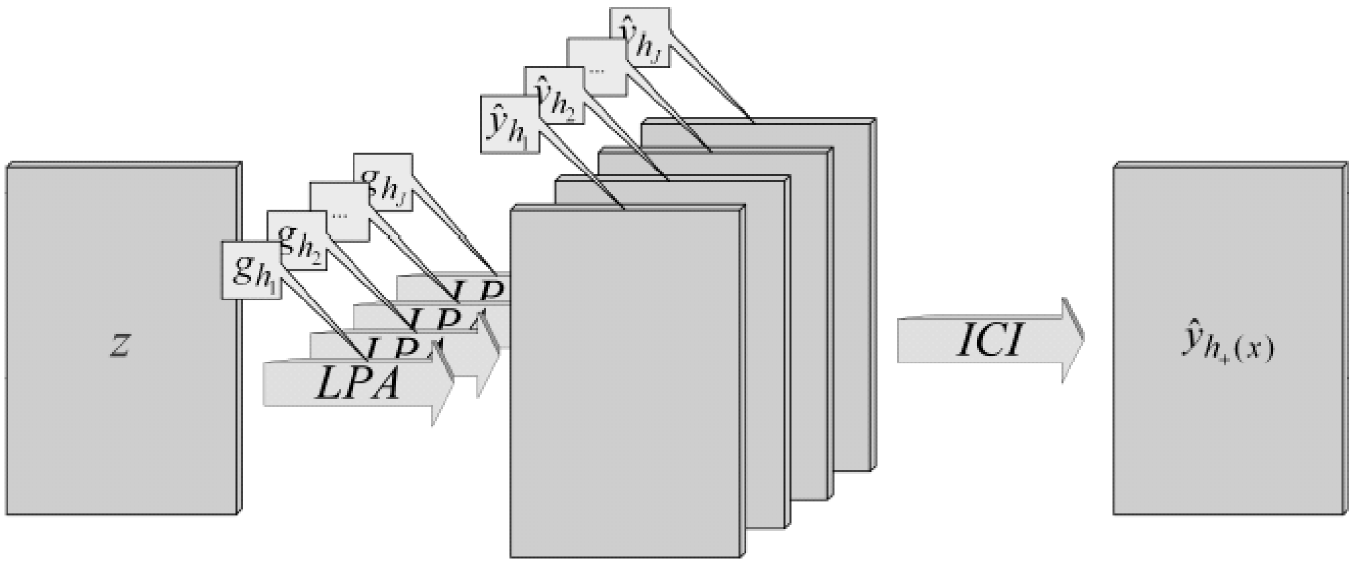

Enable anisotropic (directional) estimation

- Use directional kernels $g_{h,\theta}$ for a set of directions Θ
  - e.g. $\Theta = \{0, \frac{\pi}{2}, \pi, \frac{3}{2}\pi\}$ yields 4-directional kernels
  - $g_{h,\theta}$ are no longer symmetric with respect to the estimted pixel
  - LPA estimates are given by $\hat{y}_{h,\theta}(x) = \left(z \circledast g_{h,\theta}\right)(x)$
- Apply ICI independently to obtain $h^{+}(x, \theta)$ for each $\theta \in \Theta$, $x \in X$.

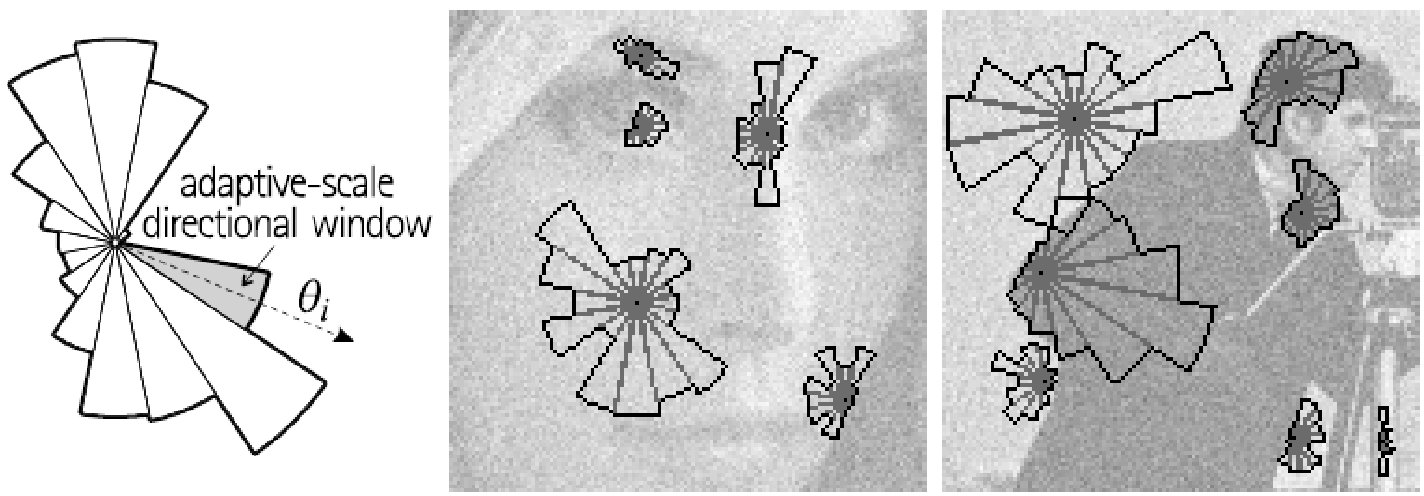

- Example of directional kernels for 3 directions. Their Fourier spectra magnitudes are shown below.

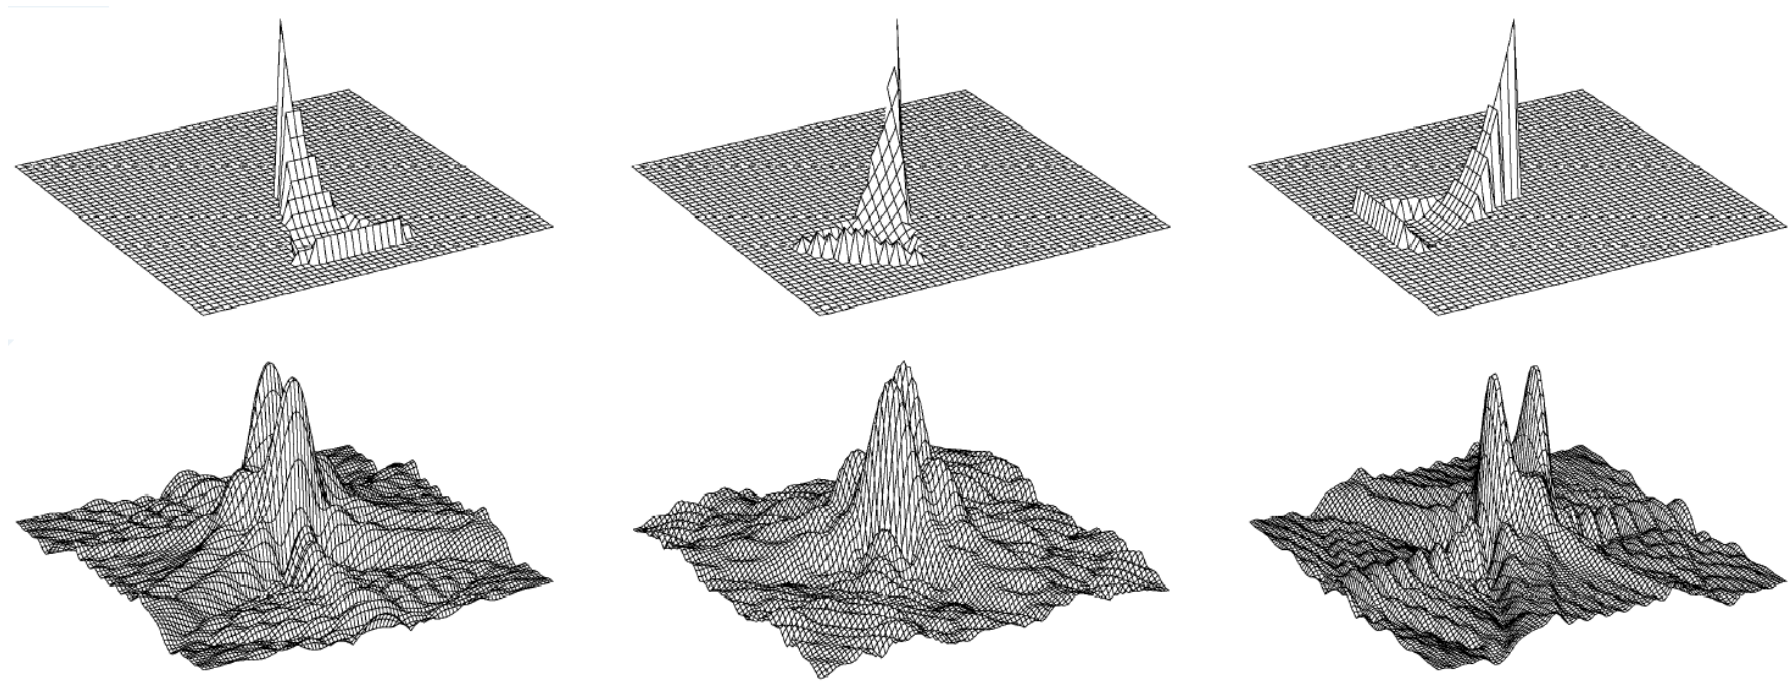

- How to fuse the directional adaptive-scale estimates $\hat{y}_{h^+,\theta}(x)$? Use weighted averaging

$$\begin{aligned}\hat{y}(x) &= \sum_{\theta\in\Theta} \lambda(x,\theta)\,\hat{y}_{h^+,\theta}(x) \\ \lambda(x,\theta) &= \frac{\sigma^{-2}_{\hat{y}_{h^+,\theta}(x)}}{\sum_{\theta'\in\Theta} \sigma^{-2}_{\hat{y}_{h^+,\theta'}(x)}},\end{aligned}$$

  where for readability $h^+(x,\theta)$ is denoted as h^+.

- Properties of the weights:
  - inversely proportional to the variance – noisier estimates have smaller weights
  - for uniform kernels, the fused estimate is the average over the anisotropic neighborhood
  - assuming that $\hat{y}_{h^+,\theta}$ are independent and unbiased, the fusing is the ML estimate of $y(x)$ given $\left\{\hat{y}_{h^+,\theta}\right\}_{\theta\in\Theta}$

- Example of adaptive scales in the case of 8 directions
  - $\Theta = \{0, \frac{\pi}{4}, \frac{\pi}{2}, \frac{3\pi}{4}, \pi, \frac{5}{4}\pi, \frac{3}{2}\pi, \frac{7}{4}\pi\}$, $H = \{1, 2, 3, 5, 6, 8\}$
  - $g_{h,\theta}$ are uniform lines (with length $h \in H$ and direction $\theta \in \Theta$)

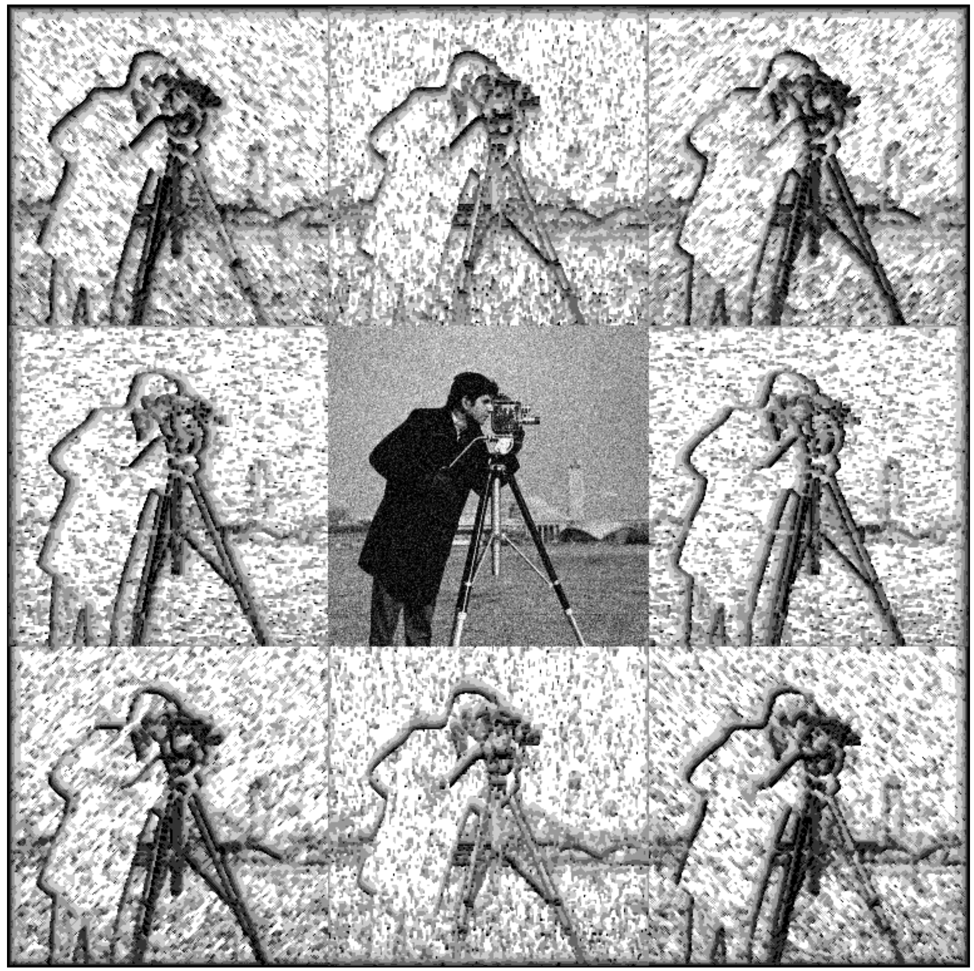

- Corresponding $\hat{y}_{h^+,\theta}(x)$ in the case of 8 directions (adaptive scales $h^+(x,\theta)$ used)and the fused estimate in the center

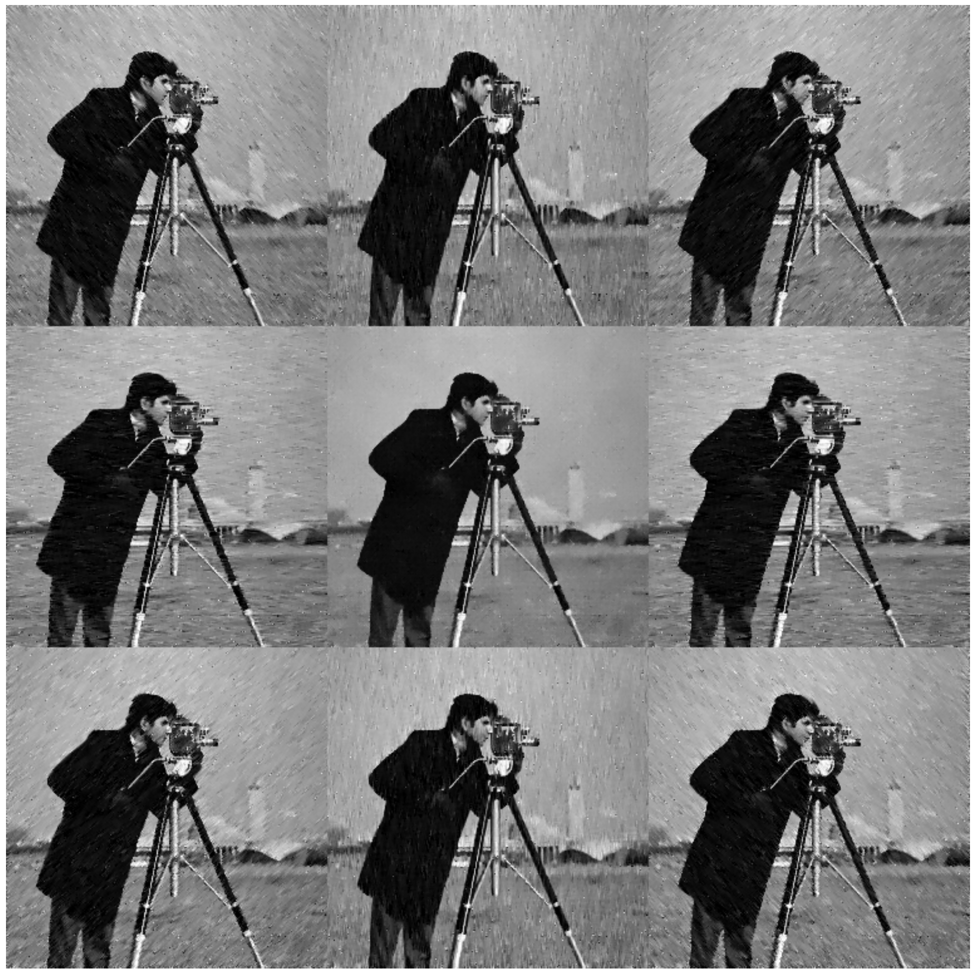

Denoising results

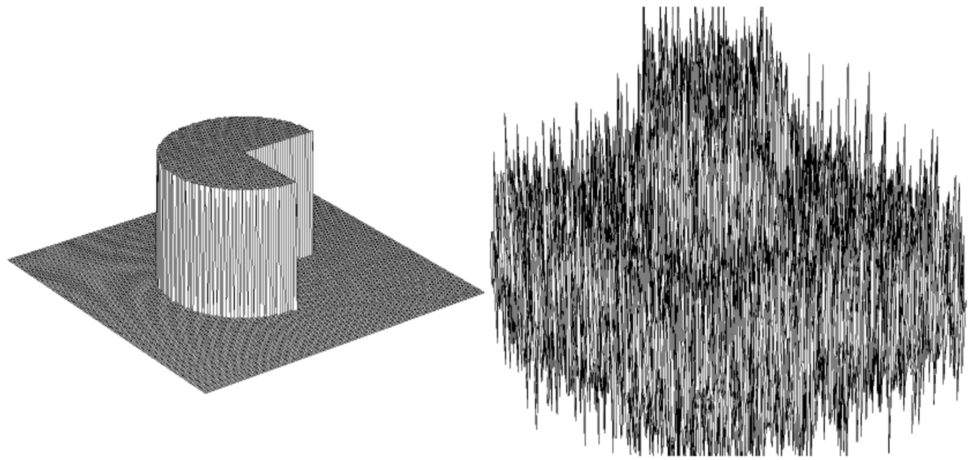

Original signal and observation (σ=0.5)

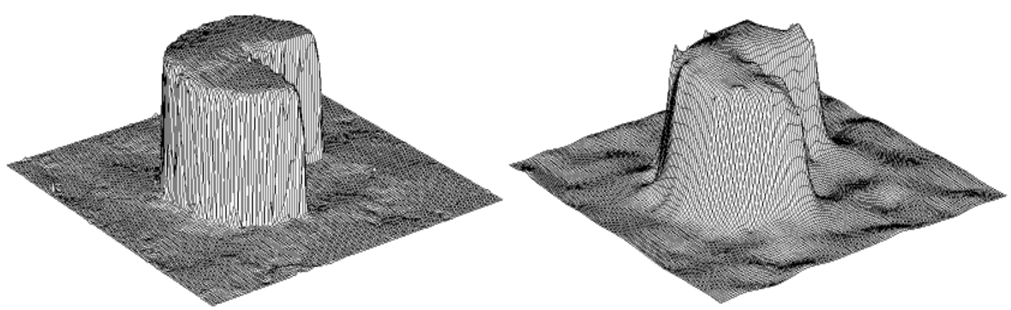

Anisotropic vs. isotropic estimation

Original *Cameraman*

Noisy observation ($\sigma = 0.1$)

Anisotropic (8-directional) LPA-ICI
ISNR 8.2 dB

Wavelet shrinkage
ISNR 7.8 dB

Non-blind deblurring

- Consider the observation model

$$z(x) = (y \circledast v)(x) + \eta(x)$$

  – v is a priori known Point-Spread Function (PSF),

  – η is additive i.i.d. Gaussian noise with variance σ^2

- The model can be represented in discrete Fourier domain as

$$Z = YV + \tilde{\eta},$$

  where capital letters denote the DFT of the corresponding signals and $\tilde{\eta}$ is i.i.d. Gaussian with variance σ^2 (just as η)

- The inverse problem is badly conditioned
  - the naive inverse $Z/V = (YV + \tilde{\eta})/V$ results in very strong noise amplification (or is not possible at all if V has exact zeros).
  - **regularization required!**
- Solve the deblurring (inverse) problem by:
  - performing regularized inversion (RI)
  - use the LPA-ICI denoising to improve the regularization by smoothing the solution.
- Why LPA-ICI denoising?
  - relatively mild regularization parameters used (leaving both more noise and more details)
  - improved detail preservation.

- Algorithm (RI part)

  - Step 1. Tikhonov regularized inversion

$$
\begin{aligned}
\hat{Y}^{RI} &= \arg\min_{\Psi} \left(\|\Psi V - Z\|^2 + \varepsilon_1^2 \|\Psi\|^2 \right) \\
&= \frac{\bar{V}}{|V|^2 + \varepsilon_1^2} Z \\
&= \frac{\bar{V}}{|V|^2 + \varepsilon_1^2} (YV + \tilde{\eta}) \\
&= \left(\frac{|V|^2}{|V|^2 + \varepsilon_1^2} Y + \frac{\bar{V}}{|V|^2 + \varepsilon_1^2} \tilde{\eta} \right)
\end{aligned}
$$

  - Step 2. Further LPA-ICI regularization: $\hat{y}_{h,\theta}^{RI} = \mathcal{F}^{-1}\left\{ \hat{Y}^{RI} G_{h,\theta} \right\}$

  - Step 3. ICI selects the adaptive scales to obtain $\hat{y}_{h^+,\theta}^{RI}$ for all $\theta \in \Theta$.

  - Step 4. Fuse $\hat{y}_{h^+,\theta\in\Theta}^{RI}$ to obtain the estimate is $\hat{y}^{RI}$.

- Algorithm (RWI part)

  – Step 1. Regularized Wiener inversion (using the $\hat{y}^{RI}$ as reference estimate)

$$\hat{Y}^{RWI} = \left(\frac{\bar{V} \left| \hat{Y}^{RI} \right|^2}{\left| V \hat{Y}^{RI} \right|^2 + \alpha_{RWI}^2 \sigma^2} \right) Z$$

  – Step 2. Further LPA-ICI regularization: $\hat{y}_{h,\theta}^{RWI} = \mathcal{F}^{-1} \left\{ \hat{Y}^{RWI} G_{h,\theta} \right\}$

  – Step 3. ICI selects the adaptive scales to obtain $\hat{y}_{h^+,\theta}^{RWI}$ for all $\theta \in \Theta$.

  – Step 4. Fuse the directional estimates $\hat{y}_{h^+,\theta \in \Theta}^{RWI}$ by weighted averaging where the weights are inversely proportional to their corresponding variances. The resultant and final estimate is $\hat{y}^{RWI}$

- Anisotropic LPA-ICI RI-RWI deconvolution algorithm flowchart

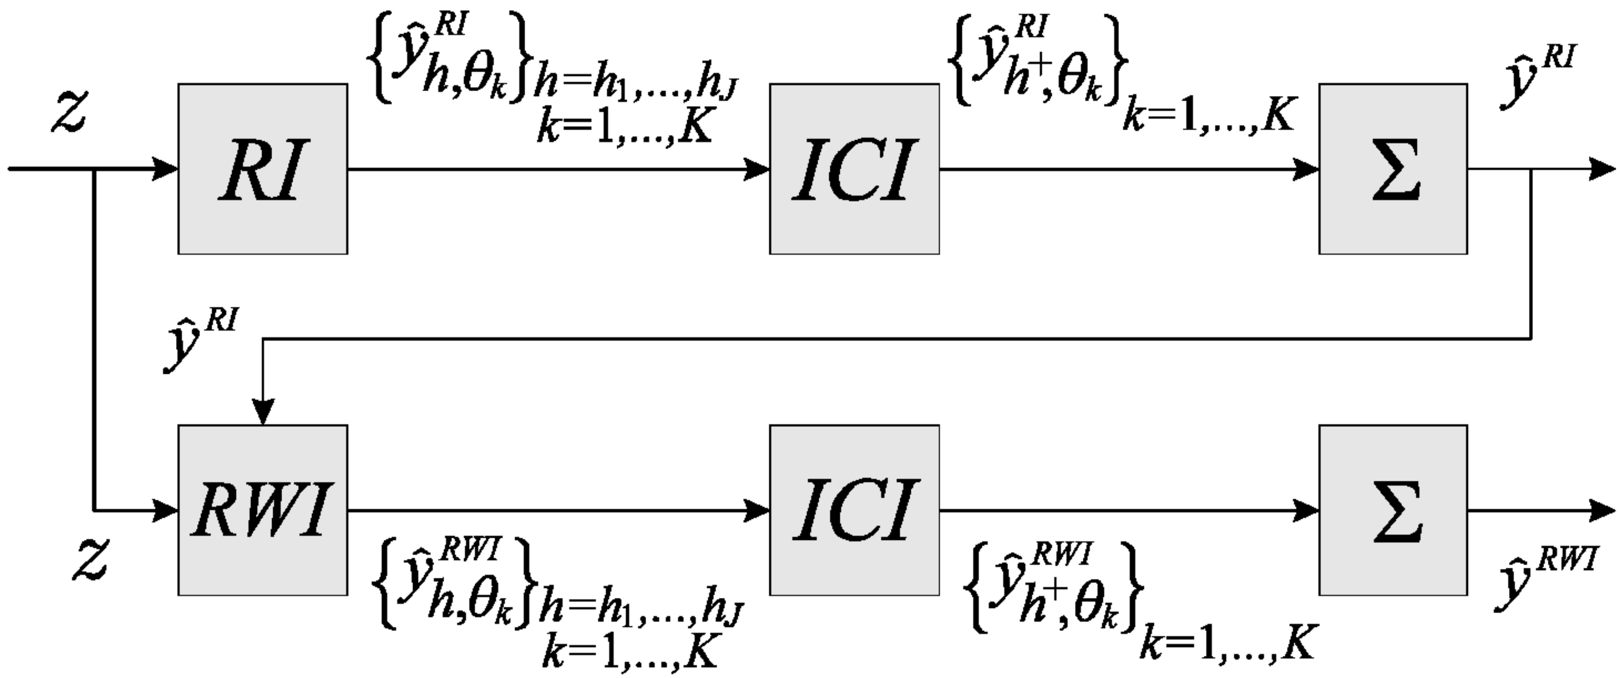

Blur: 9x9 uniform "box-car"
BSNR 40dB

Anisotropic LPA-ICI RI-RWI estimate
ISNR 8.29dB

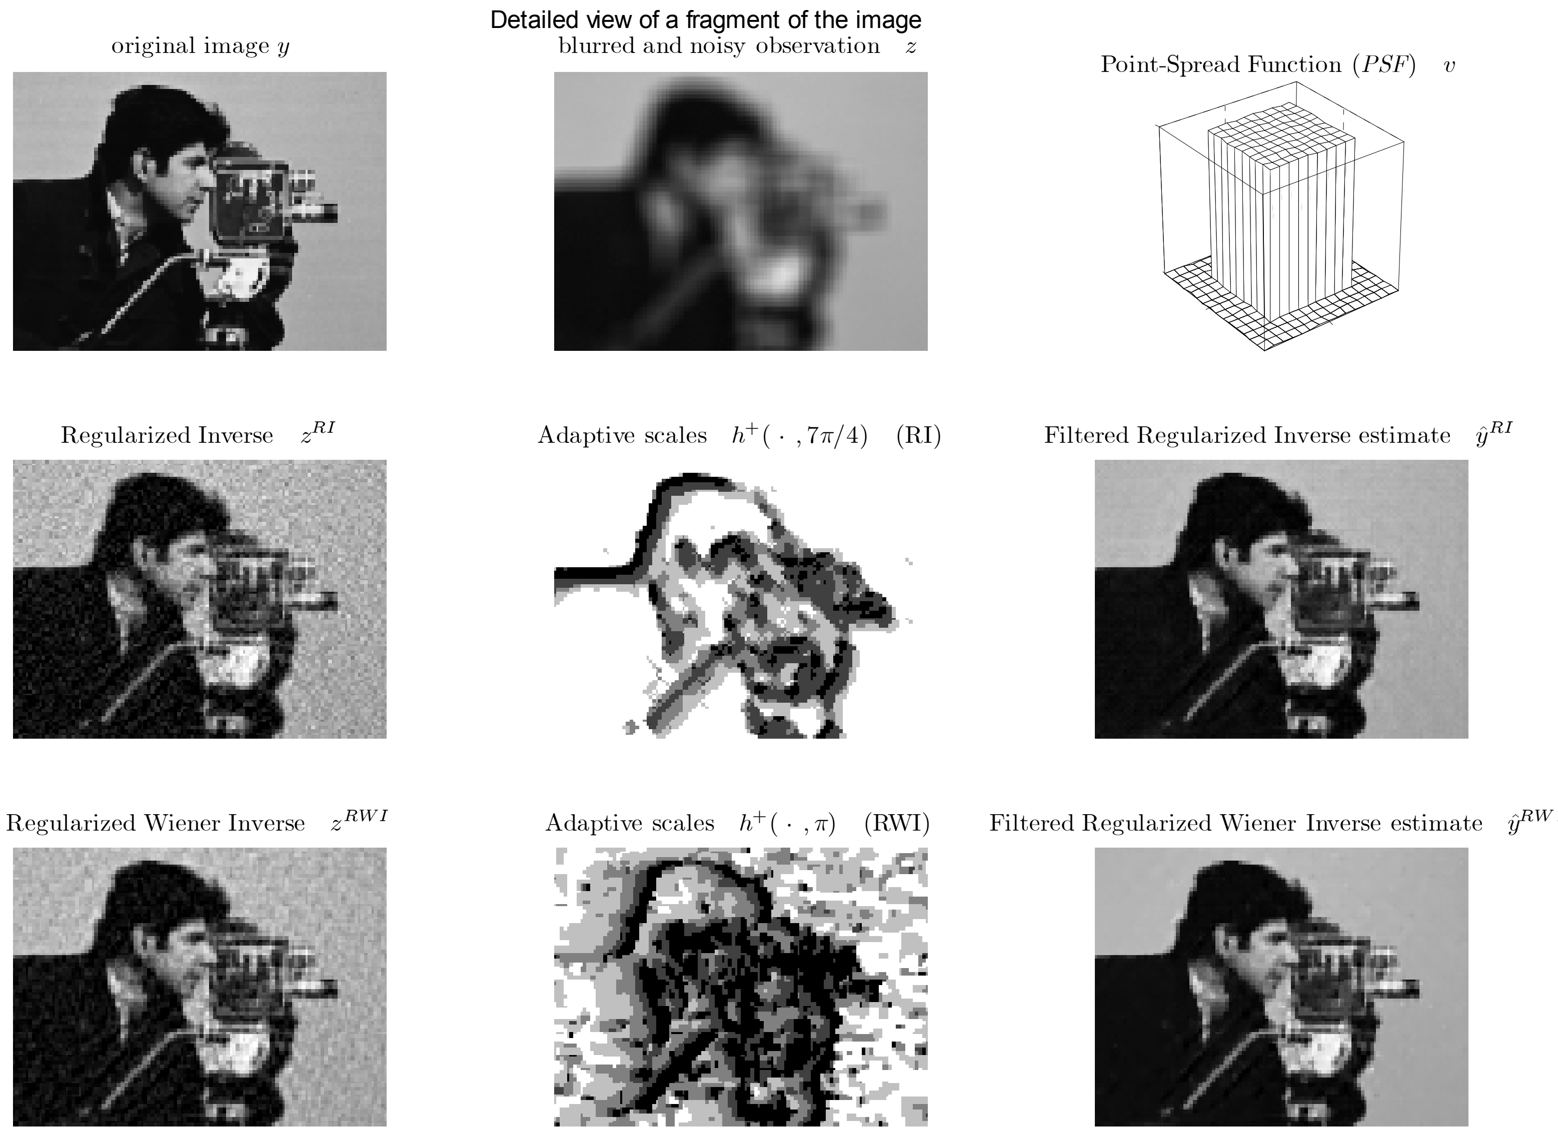

Detailed view of a fragment of the image
original image y
blurred and noisy observation z
Point-Spread Function (PSF) v
Regularized Inverse z^{RI}
Adaptive scales $h^{+}(\cdot, 7\pi/4)$ (RI)
Filtered Regularized Inverse estimate $\hat{y}^{RI}$
Regularized Wiener Inverse z^{RWI}
Adaptive scales $h^{+}(\cdot, \pi)$ (RWI)
Filtered Regularized Wiener Inverse estimate $\hat{y}^{RWI}$

Derivative estimation from noisy blurred images

- Use the RI-RWI deblurring algorithm but replace the estimation (smoothing kernels) $g_{h,\theta}$ with derivative estimation kernels.

- Then, $\hat{y}_{h^+,\theta}^{RWI}$ is an estimate of $\partial_\theta y$ (derivative of y in direction θ).

- For example, for $\theta = \pi/4$, when using the 8-directional LPA-ICI, we compute the derivative as $\hat{\partial}_{\pi/4} = \left(\hat{y}_{h^+,\pi/4}^{RWI} - \hat{y}_{h^+,5\pi/4}^{RWI}\right)$

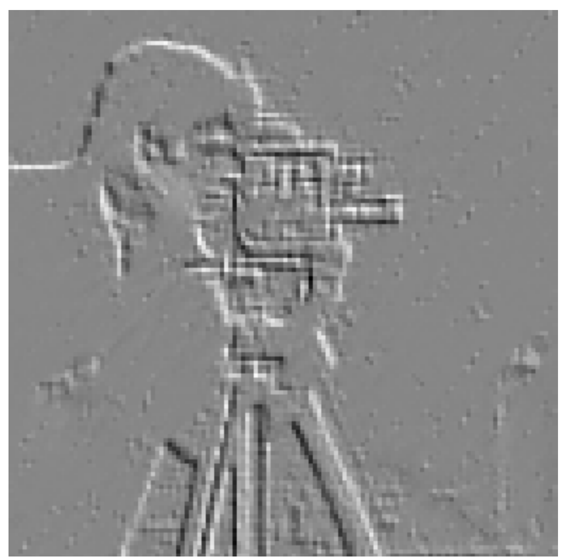

- For edge detection, one can use the sum of the absolute values of the directional derivatives

- Example for the 8-directional LPA-ICI, $\sum_{k=1}^{4} \hat{\partial}_{\theta_k}$

Inverse halftoning by LPA-ICI RI-RWI deconvolution

- The Inverse halftoning problem is to recover a grayscale image from a binary halftone

- Error-diffusion using Floyd-Steinberg or Jarvis masks are common halftoning methods

	•	7	
3	5	1	/16

Floyd-Steinberg

		•	7	5	
3	5	7	5	3	/48
1	3	5	5	3	

Jarvis et al.

- We can do a coarse linear modeling of the halftoning as convolution with additive colored noise:

$$Z = PY + Q\tilde{\eta}$$

  - P and Q can be computed from the error-diffusion masks.
  - The difference with the LPA-ICI RI-RWI's modeling is the colored noise component $Q\tilde{\eta}$

- The Fourier spectra P and Q in the case of Floyd-Steinberg halftoning. Most of their energy in the high frequencies – unlike the case of blurring.

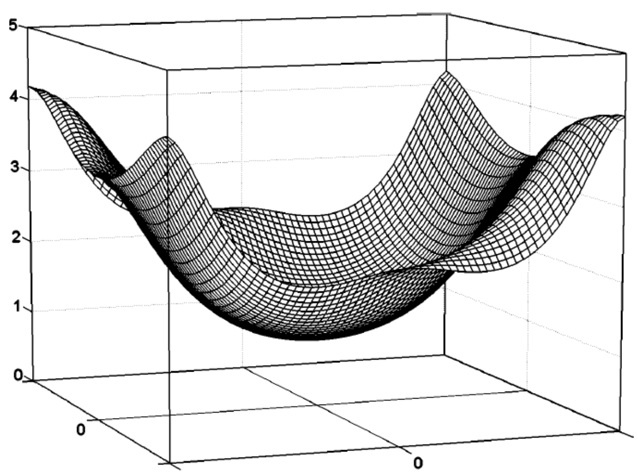

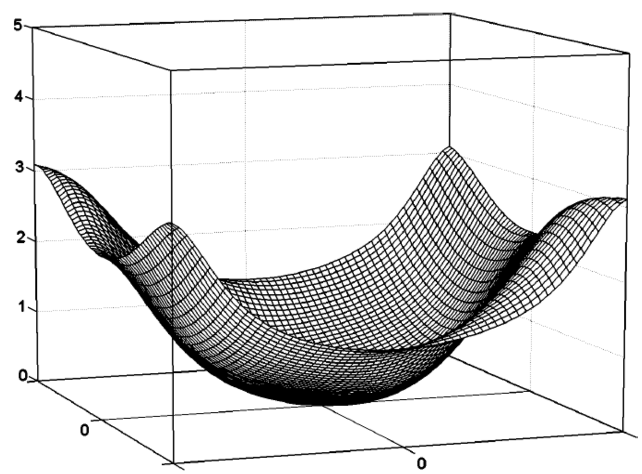

- We extend the RI-RWI deconvolution as follows.

- In the RI step,

$$\hat{Y}_{h,\theta}^{RI} = \frac{\bar{P} G_{h,\theta}}{|P|^2 + |Q|^2 \varepsilon_1^2} Z$$

- In the RWI step,

$$\hat{Y}_{h,\theta}^{RWI} = \frac{\bar{P} \left| \hat{Y}^{RI} \right|^2 G_{h,\theta}}{\left| P \hat{Y}^{RI} \right|^2 + \alpha_{RWI}^2 |Q|^2 \sigma^2} Z$$

- Anisotropic RI-RWI inverse halftoning results

Lena, Floyd-Steinberg halftone

LPA-ICI RI-RWI inverse halftoning

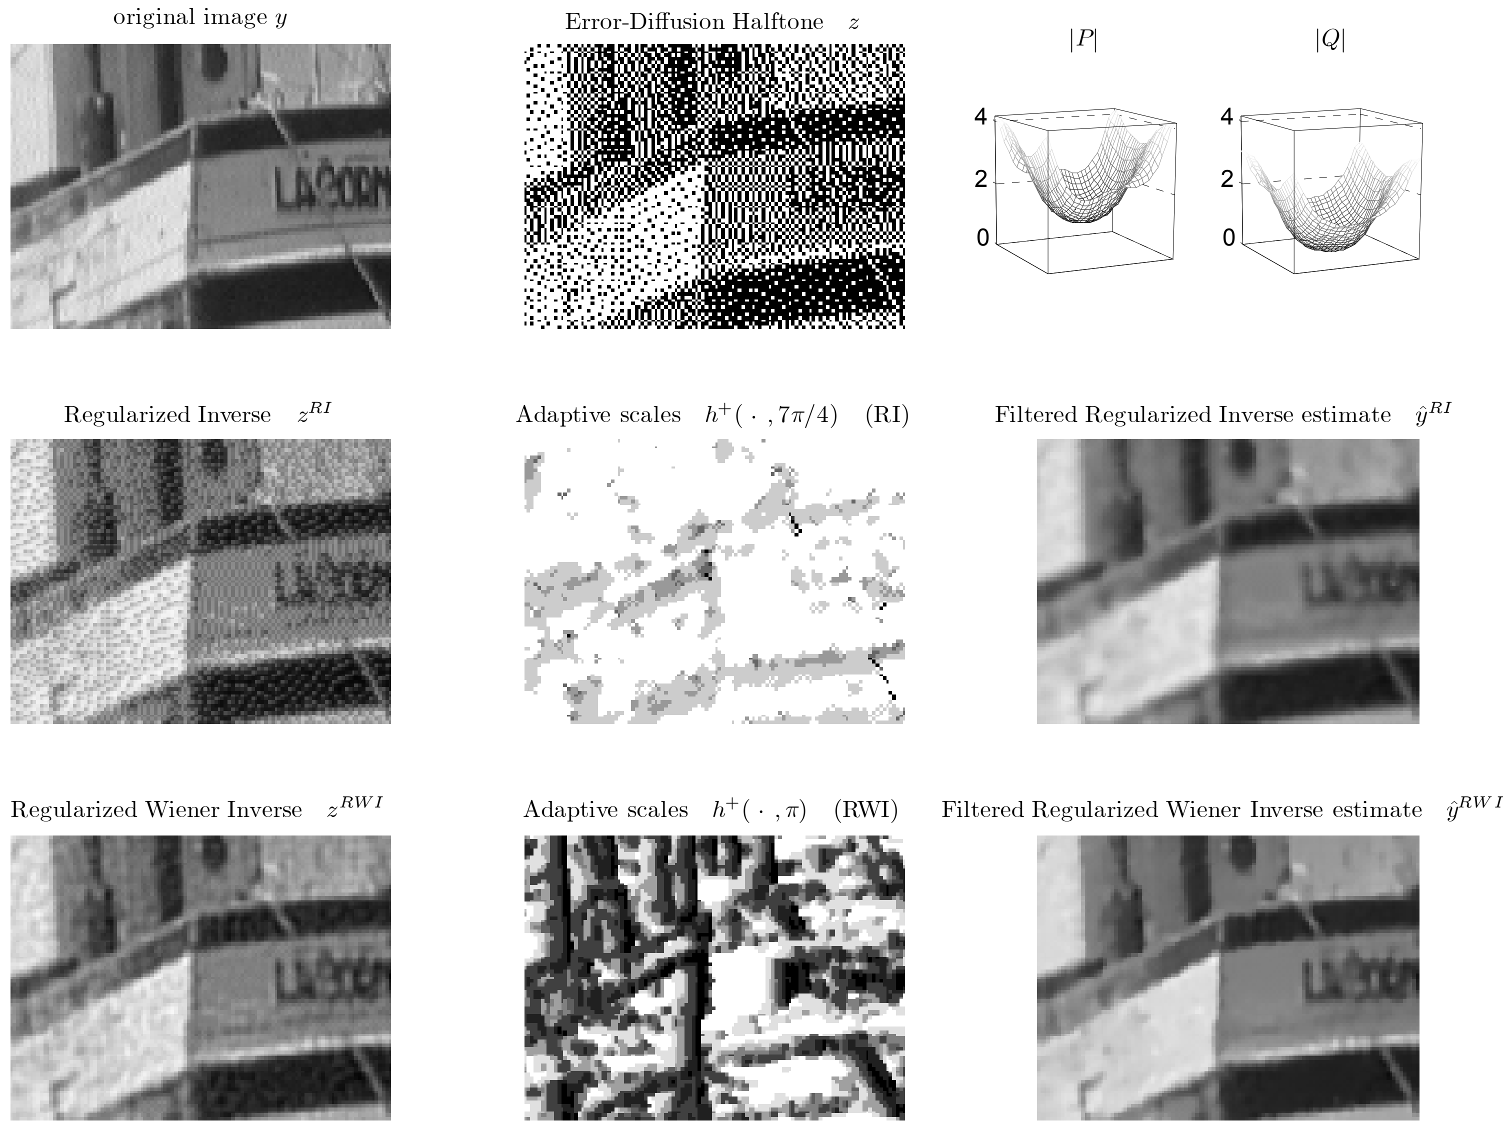
original image y
Error-Diffusion Halftone z
$|P|$
$|Q|$
4
2
0
4
2
0
Regularized Inverse z^{RI}
Adaptive scales $h^{+}(\,\cdot\,, 7\pi/4)$ (RI)
Filtered Regularized Inverse estimate $\hat{y}^{RI}$
Regularized Wiener Inverse z^{RWI}
Adaptive scales $h^{+}(\,\cdot\,, \pi)$ (RWI)
Filtered Regularized Wiener Inverse estimate $\hat{y}^{RWI}$

- Matlab code and publications can be found at:

`http://www.cs.tut.fi/~lasip`