Nonlocal image processing: from sensor noise modeling to image restoration

Tampere University of Technology

Outline

- **Signal-dependent noise modeling and removal for digital imaging sensors**
- **Advanced image processing techniques:**
  shape-adaptive methods
  nonlocal transform-based methods
- **Applications:**
  denoising
  deblurring
  superresolution
  color filter array interpolation

Photon-limited image restoration for mobile camera

Goal: **To develop efficient and theoretically well-grounded approaches for restoration of still images produced by a mobile digital camera.**

- **Sensor's raw-data noise analysis methodology and modeling**
- **Algorithms for automatic estimation of noise-parameters from raw-data**
- **Spatially adaptive image denoising and deblurring**
- **Noise-robust color interpolation from Bayer pattern raw-data**
  - **- Joint interpolation-denoising and interpolation-deblurring**
- **High Dynamic Range (HDR) imaging for raw-data**

Statistical characteristics of the raw data shall be properly assessed and used in the imaging chain to improve the final image quality.

Exposure-time/noise trade-off

Digital imaging sensors can have very different performance

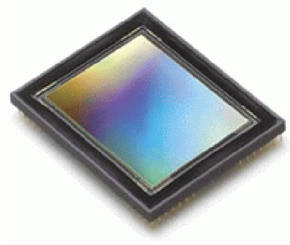

Different acquisition settings result in different noise levels in the image

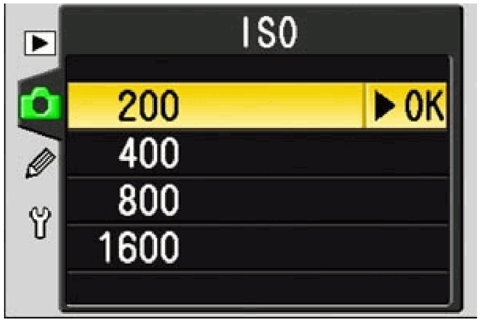

"Exposure-time/noise trade-off "

Statistical analysis of raw data: experimental measurements

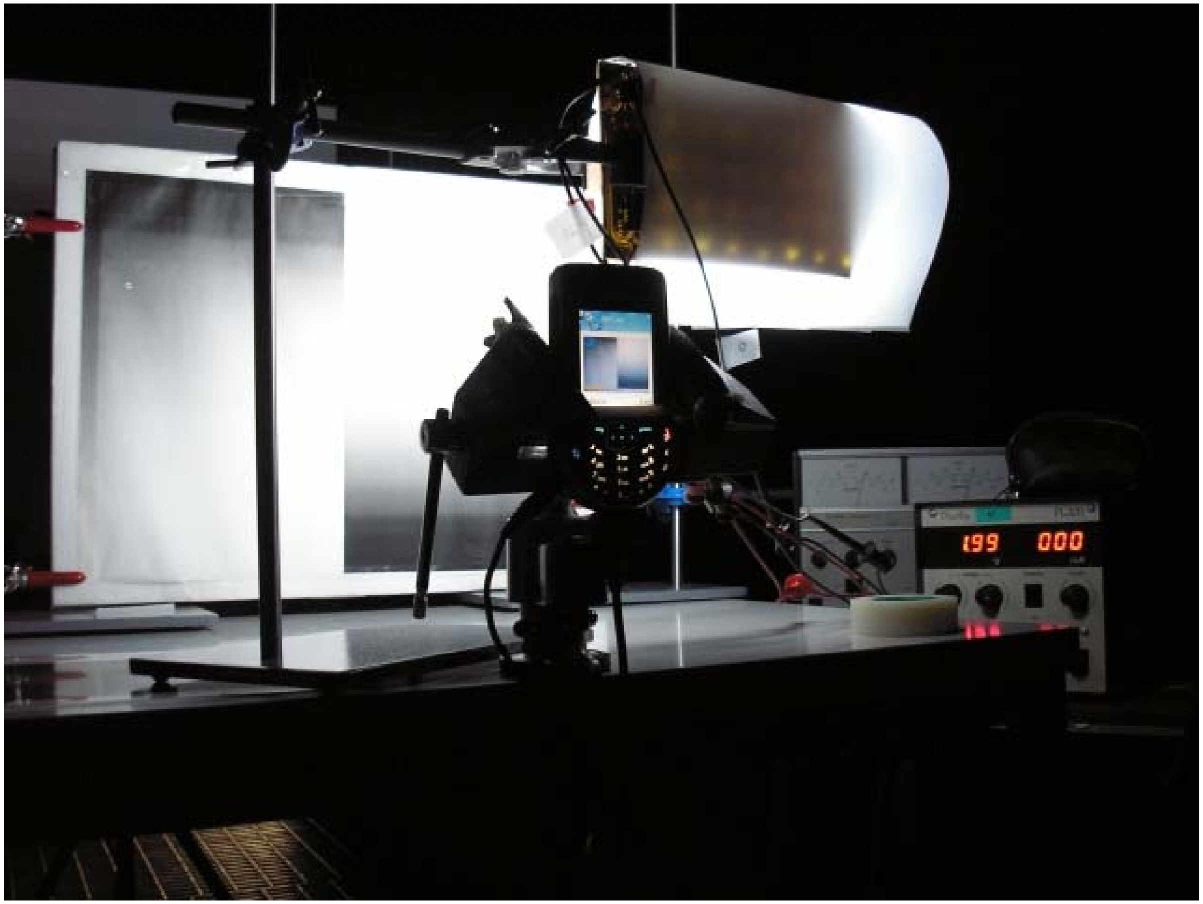

Statistical analysis of raw data: Signal-dependent noise model

The analysis of experimental data demonstrates that:

1. **The model of noise is close to the Poissonian one**
2. **Model parameters depend neither on the color channel nor on the exposure time**

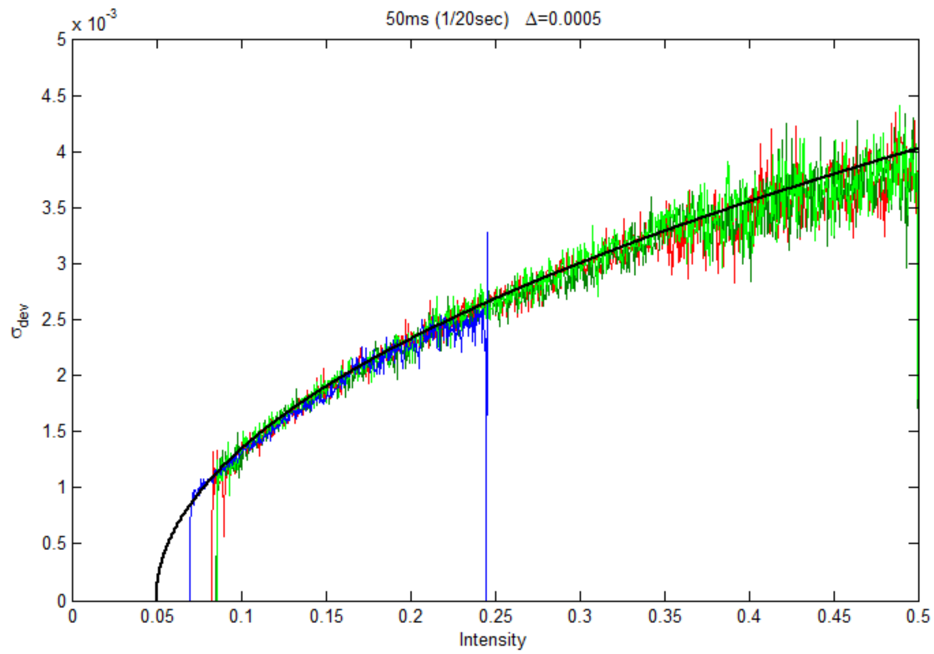

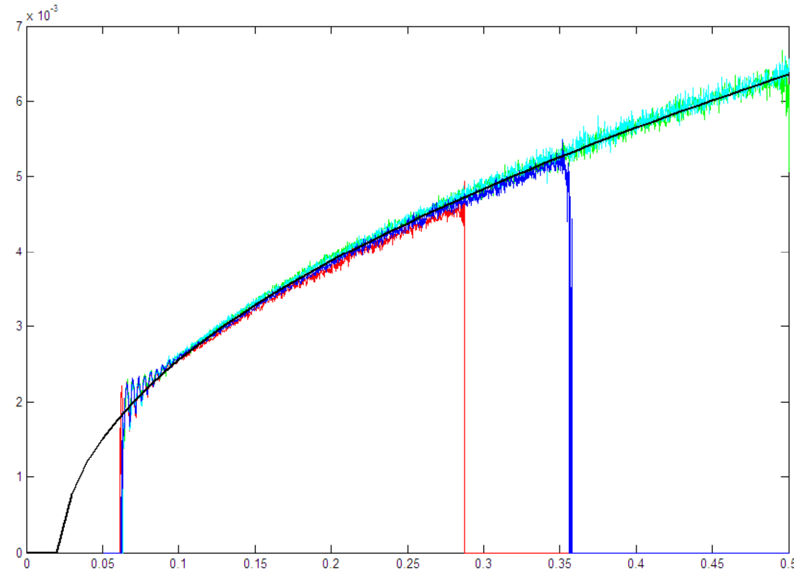

Parametric signal-dependent noise-modelling: Poissonian-Gaussian with clipping

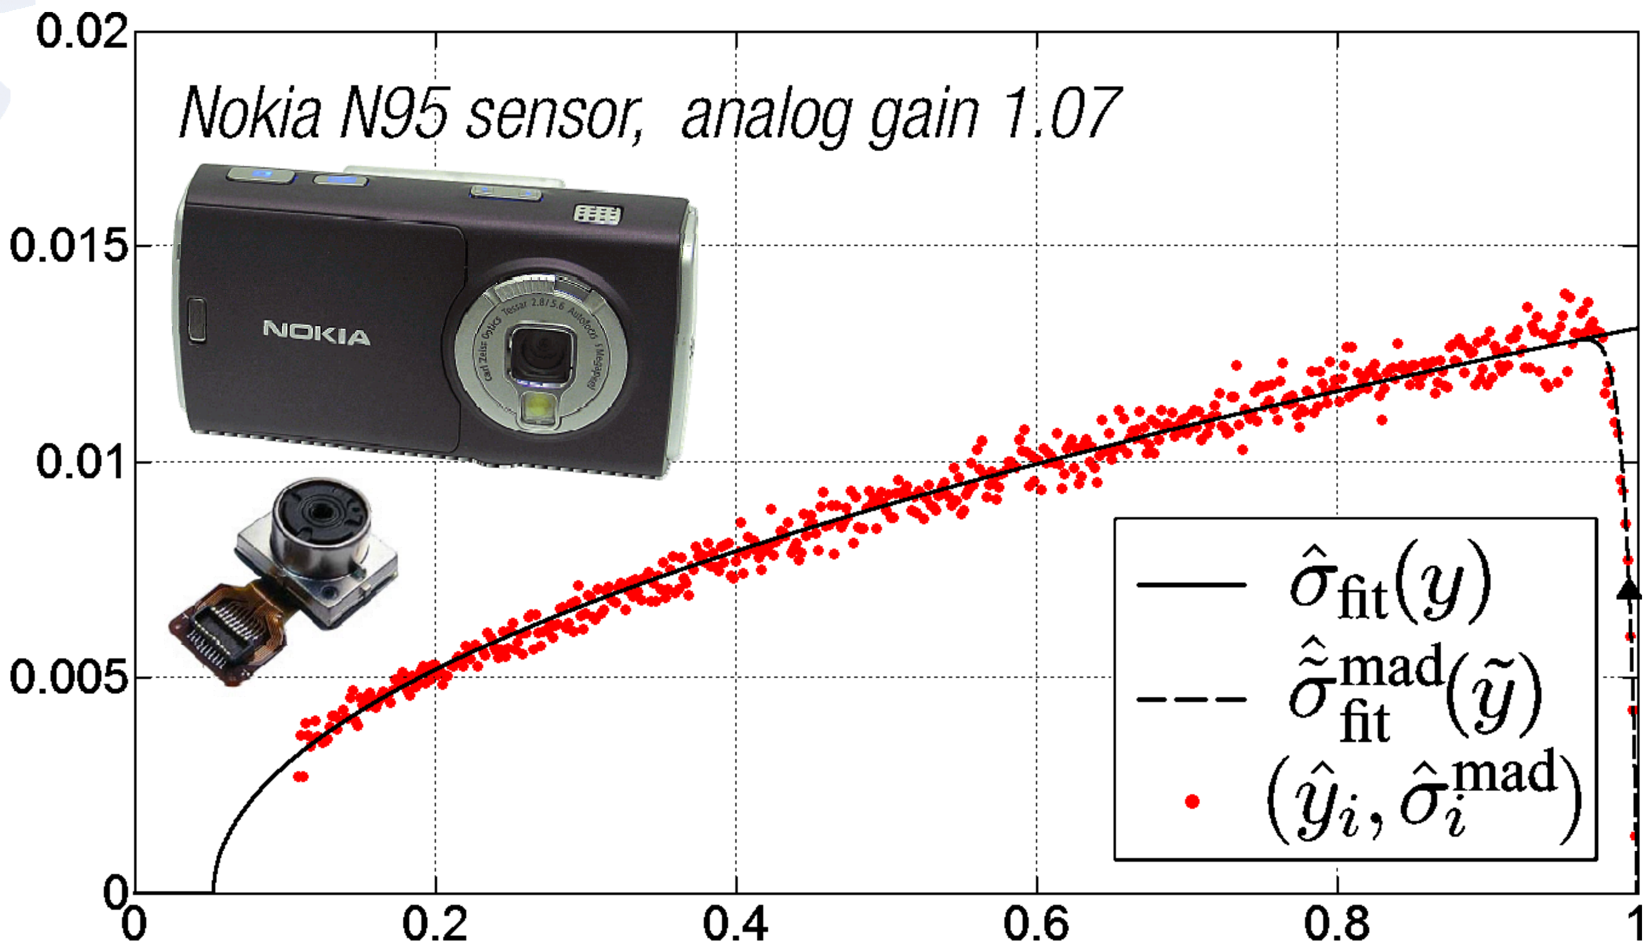

Parametric signal-dependent noise-modelling: automatic estimation from single-image raw-data (*http://www.cs.tut.fi/~foi/sensornoise.html*)

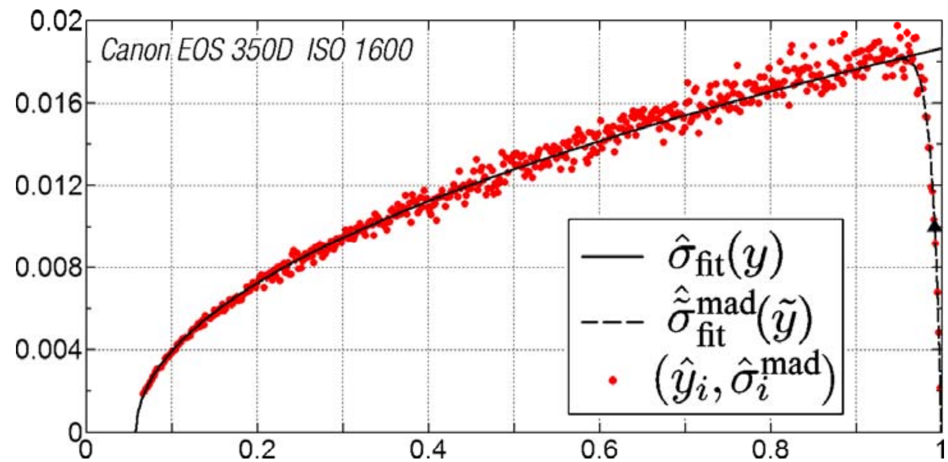

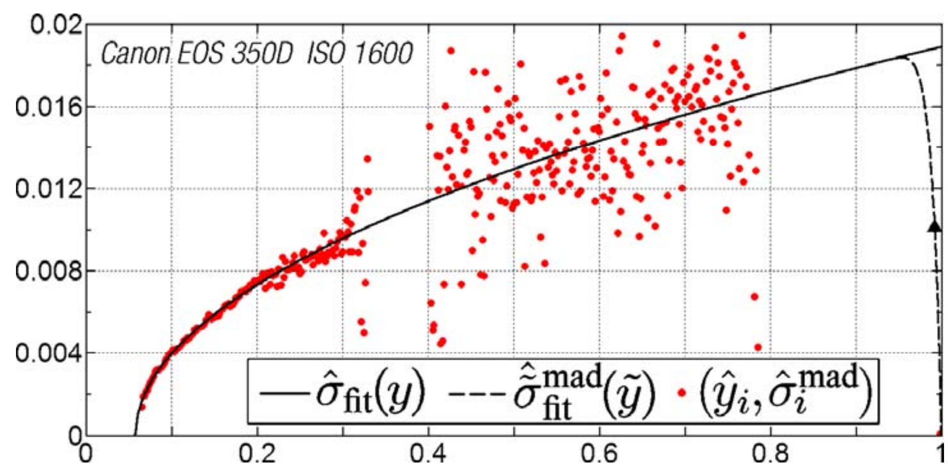

Practical modeling for raw data: idea

- Model photon-to-electron conversion using Poisson distributions (signal dependent);
- Model the other noise sources as signal-independent and Gaussian (central-limit theorem);
- Exploit normal approximation of Poisson distributions;
- The acquisition/dynamic range is limited: too dark or too bright signals are clipped;
- There can be a pedestal;
- Spatial dependencies can be ignored for normal operating conditions (go for independent noise).

Eventually, only two parameters are sufficient to describe the noise model where the raw data is described as clipped signal-dependent observations.

Experiment: clipped noisy data

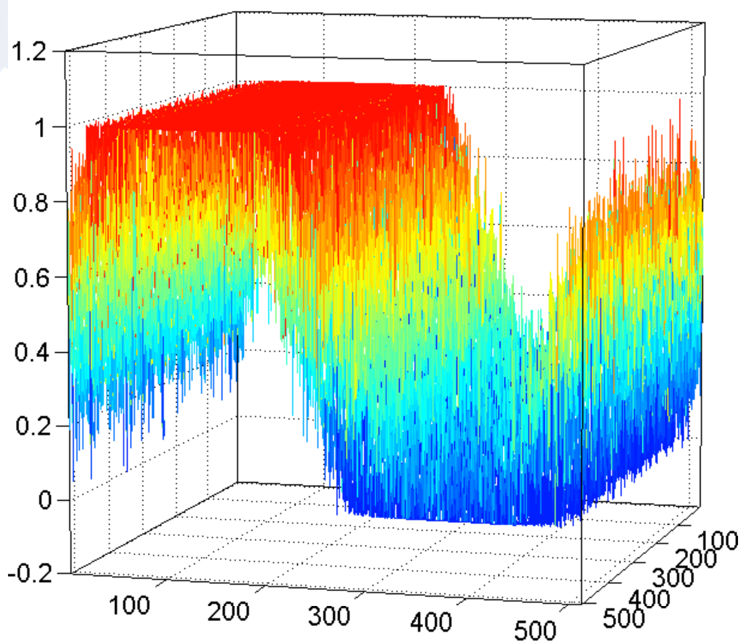

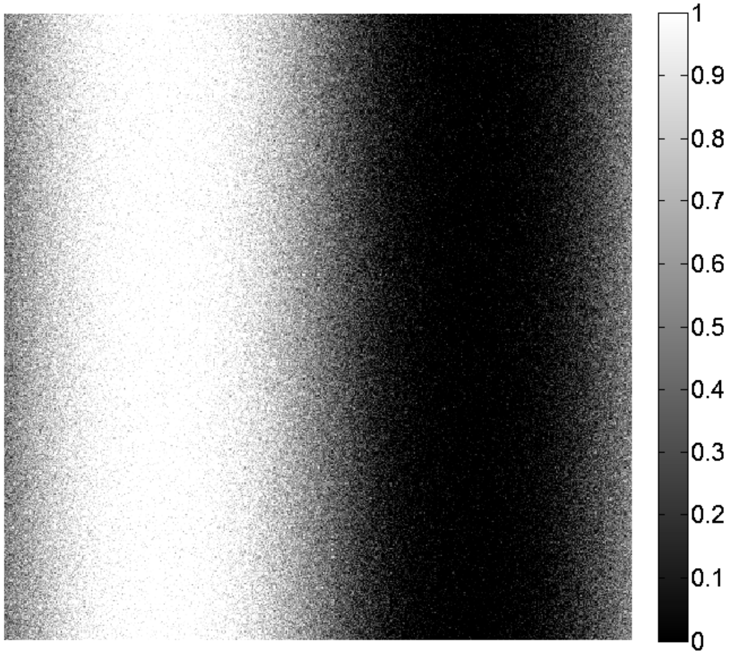

Original image : $y(x1, x2) = 0.7 \sin(2\pi\, x1/512) + 0.5$

Experiment: Noise Estimation

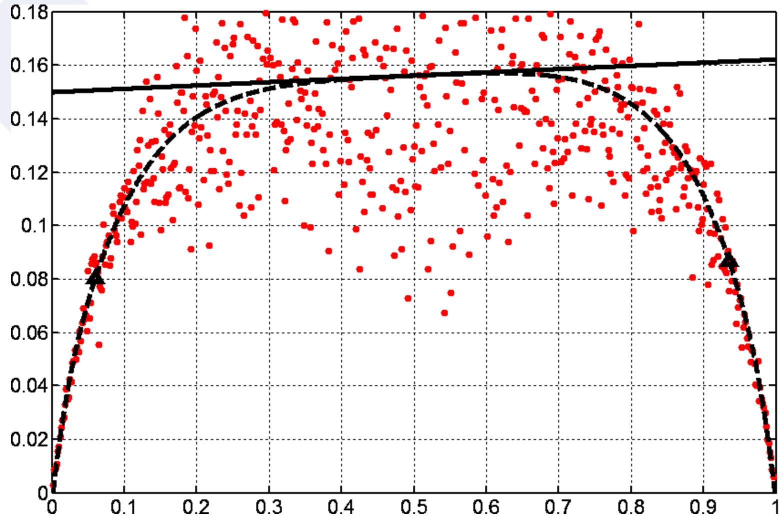

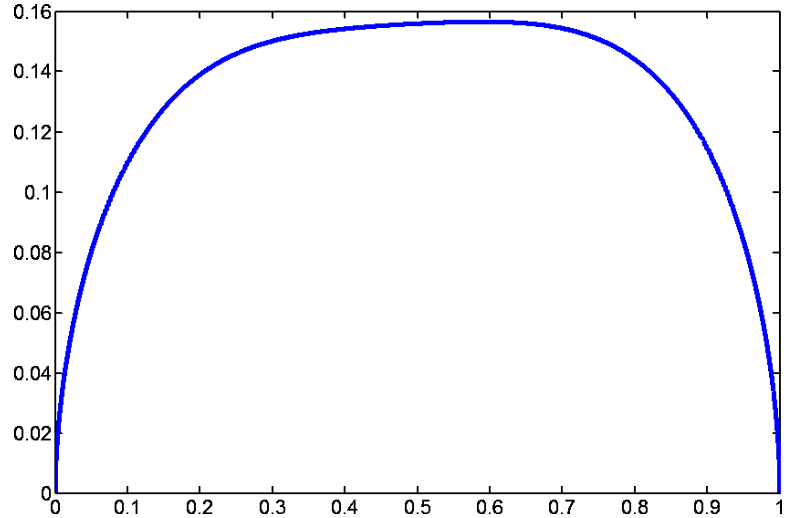

estimation and fitting a = 0.0038, b = 0.022

st.dev.-function .σ

Experiment: denoised estimate after variance stabilization before declipping

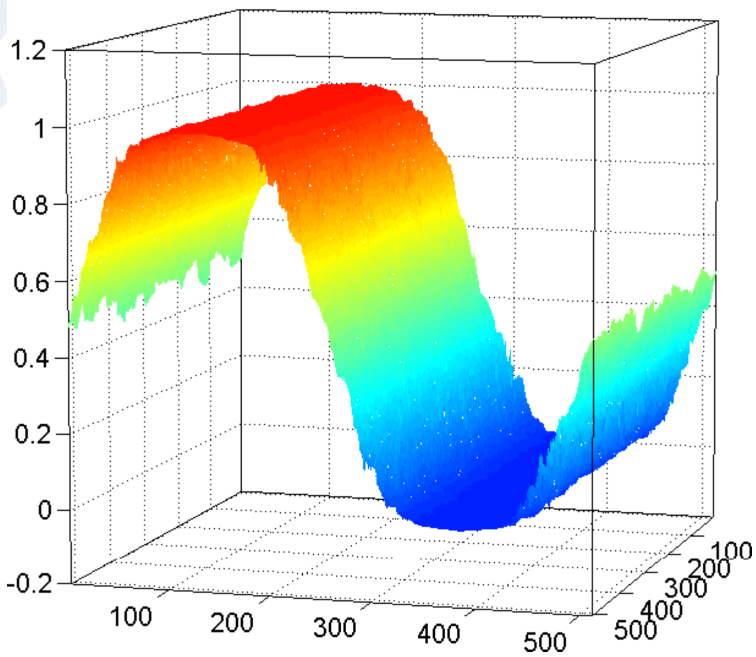

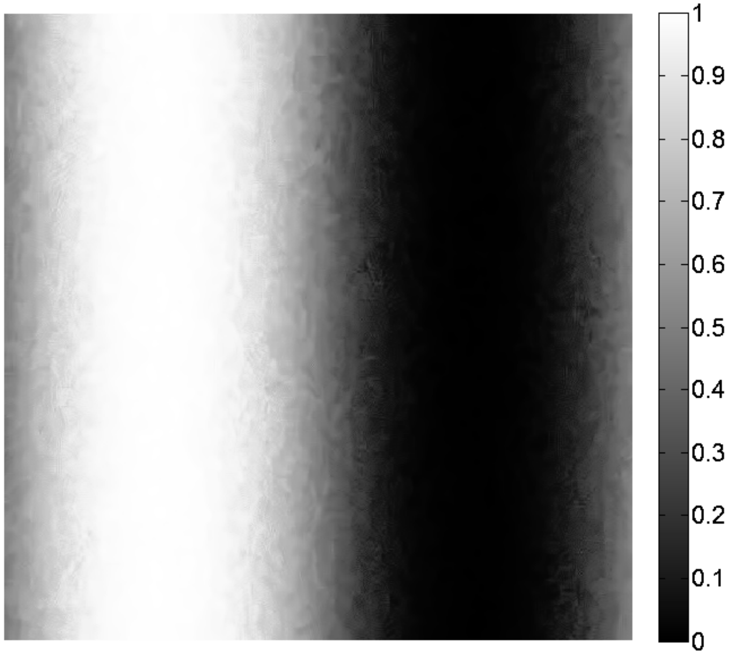

Experiment: declipped estimate

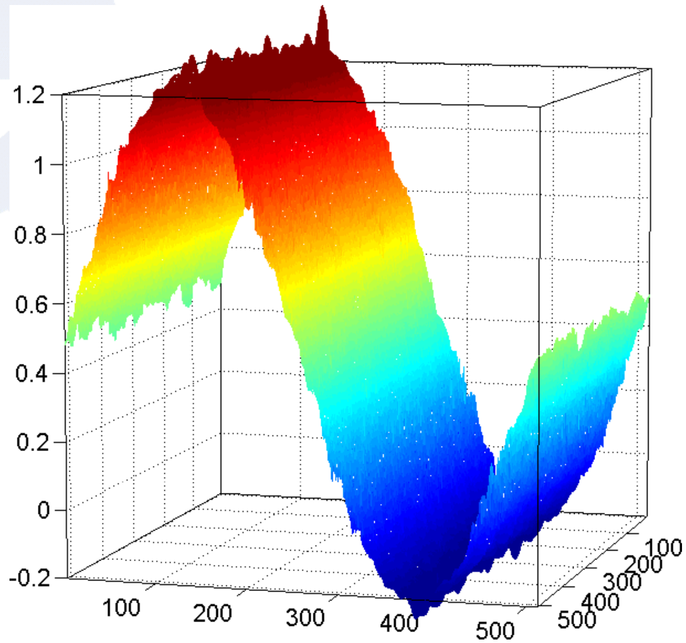

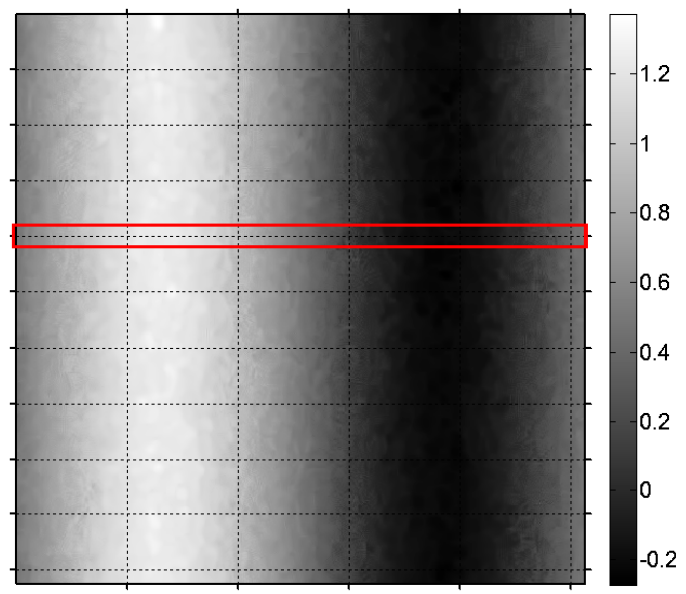

Experiment: declipped estimate (crosssection)

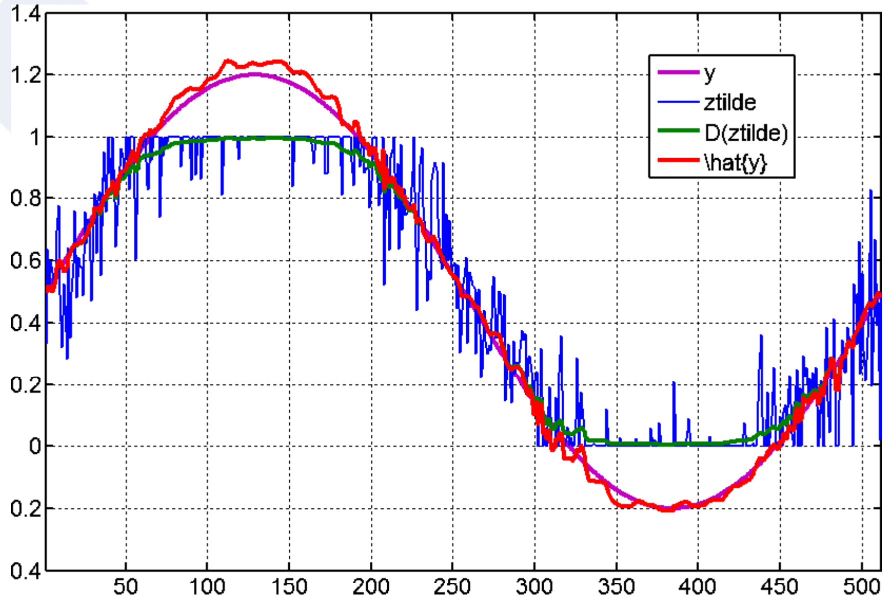

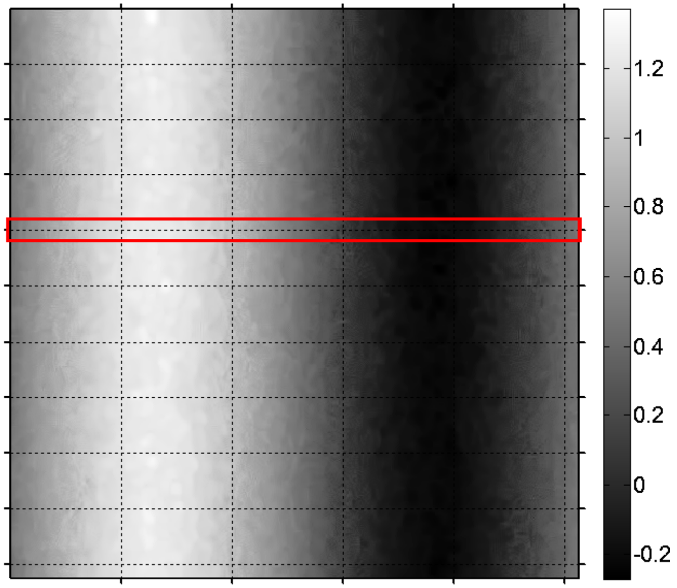

Real experiment: (Raw-data from Fujiflm FinePix S9600, ISO 1600)

Real experiment: Denoising before declipping

Real experiment: Denoising after declipping

Real experiment: Denoising after declipping (crossection)

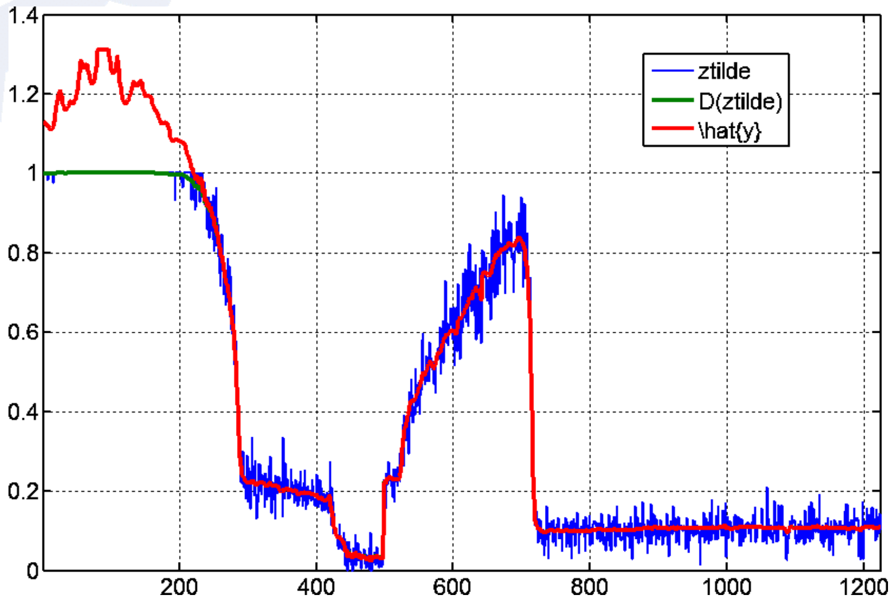

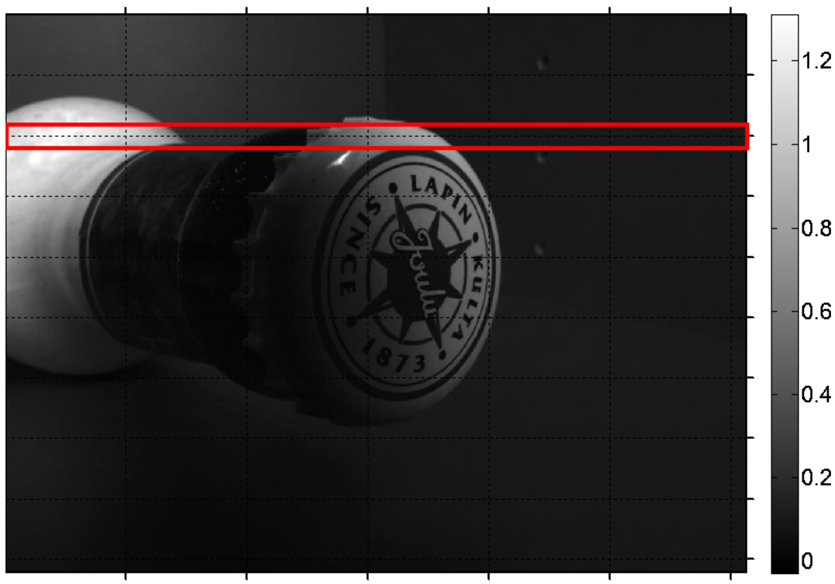

Transform methods in Image and Video Processing

- ***Local Approximation Signal and Image Processing (LASIP) Project***

LASIP project is dedicated to investigations in a wide class of novel efficient adaptive signal processing techniques.

Statistical methods for restoration from noisy and blurred observations of one-dimensional signals, images, 3D microscopy and video data were recently proposed.

More info: www.cs.tut.fi/~lasip/

Transform methods in Image and Video Processing
LASIP Project

LPA estimates, bias and variance, and asymptotic *MSE*

The observation model is $z = y + \eta$, where y is the true signal and η is noise. Let $\hat{y}_h$ denote the *LPA* estimate and the *LPA* kernel corresponding to different values of a scale parameter h:

$$\hat{y}_h = z \circledast g_h \qquad \text{where} \quad g_h = g\left(\cdot/h\right)$$

$$\text{Bias:} \quad b_{\hat{y}_h(x)} = y\left(x\right) - \left(y \circledast g_h\right)\left(x\right) \quad (\eta \text{ zero-mean and independent})$$

$$\text{Variance:} \quad \sigma^2_{\hat{y}_h(x)} = \left(\sigma^2_z \circledast g^2_h\right)\left(x\right) \qquad (\text{if } \sigma^2_z \equiv \sigma^2 \text{ then } \sigma^2_{\hat{y}_h(x)} \equiv \sigma^2_z \left\|g_h\right\|^2_2)$$

The following asymptotic expressions for the bias, variance and *MSE* of hold:

$$b_{\hat{y}_h} = ch^a, \qquad \sigma^2_{\hat{y}_h} = dh^{-b}, \qquad l_{\hat{y}_h(x)} = c^2 h^{2a} + dh^{-b}.$$

Transform methods in Image and Video Processing
LASIP: Intersection of Confidence Intervals (*ICI*) rule

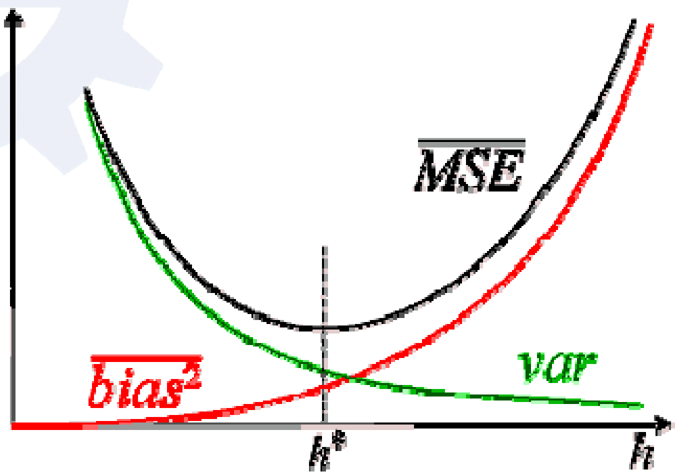

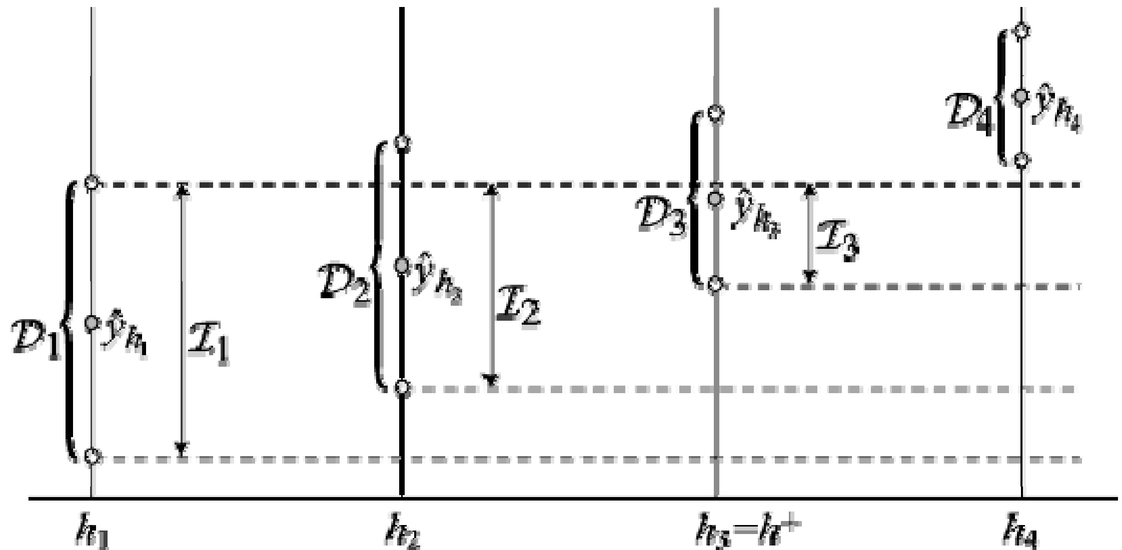

The estimates $\hat{y}_h(x)$ are calculated for a set $H = \{h_j\}_{j=1}^J$ of increasing scales. The *ICI* rule yields a pointwise adaptive estimate $\hat{y}_{h^+}(x)$, where for every x an adaptive scale $h^+(x) \in H$ is used; $h^+(x) \approx h^*(x)$. The *ICI* rule is as follows. *Consider the intersection of confidence intervals* $\mathcal{I}_j = \bigcap_{i=1}^j \mathcal{D}_i$, *where* $\mathcal{D}_i = [\hat{y}_{h_i}(x) - \Gamma\sigma_{\hat{y}_{h_i}}, \hat{y}_{h_i}(x) + \Gamma\sigma_{\hat{y}_{h_i}}]$ *and* $\Gamma > 0$ *is a threshold parameter, and let* j^+ *be the largest of the indexes* j *for which* $\mathcal{I}_j$ *is non-empty,* $\mathcal{I}_{j^+} \neq \varnothing$ *and* $\mathcal{I}_{j^+ + 1} = \varnothing$. *Then,* h^+ *is defined as* $h^+ = h_{j^+}$ *and the adaptive estimate is* $\hat{y}_{h^+}(x)$.

Transform methods in Image and Video Processing

LASIP: HOW LPA-ICI WORKS

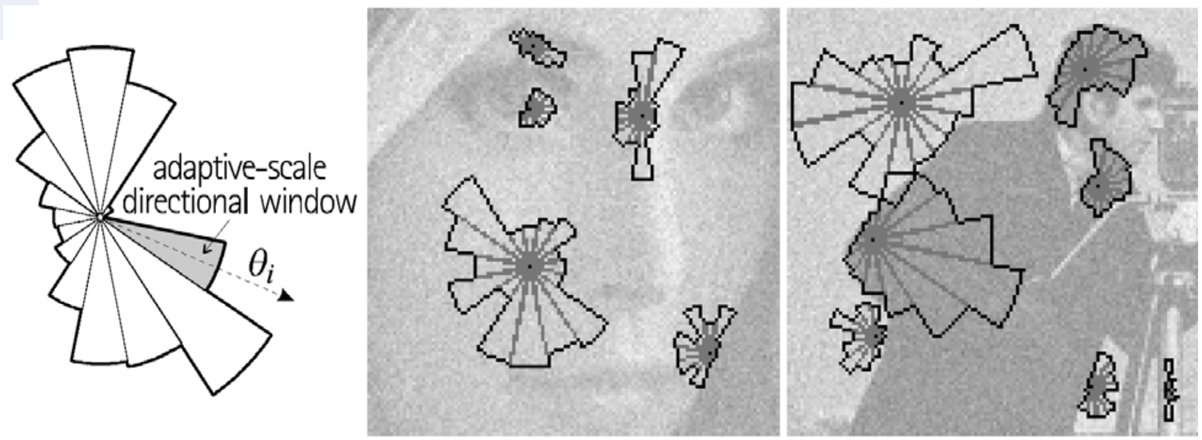

Figure 1: Anisotropic local approximations achieved by combining a number of adaptive-scale directional windows. The examples show some of these windows selected by the directional *LPA-ICI* for the noisy *Lena* and *Cameraman* images.

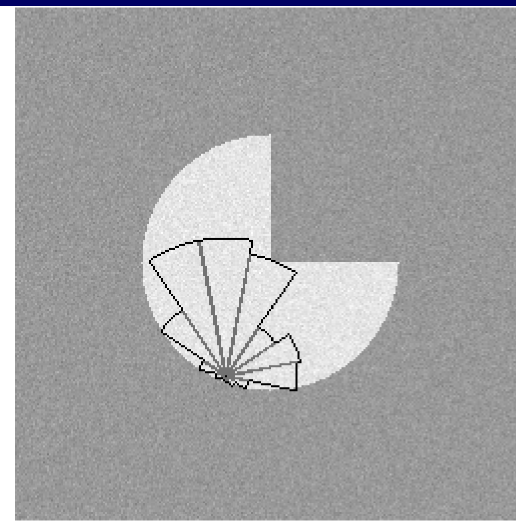

Anisotropic filtering based on pointwise adaptive nonparametric regression: Denoising and Deblurring

Removing of signal-dependent noise from images obtained by camera-phone

Extended Depth of Field Through Digital Inverse Filtering

Anisotropic filtering based on pointwise adaptive nonparametric regression: Super-resolution imaging

Available low resolution data

Combination of many low-resolution images (high-resolution blurred image)

Sharp high resolution data

Shape-adaptive DCT image filtering

By demanding the local fit of a polynomial model, we are able to avoid the presence of singularities or discontinuities within the transform support. In this way, we ensure that data are represented sparsely in the transform domain, significantly improving the effectiveness of shrinkage (e.g., thresholding).

noisy image and adaptive-shape neighborhood

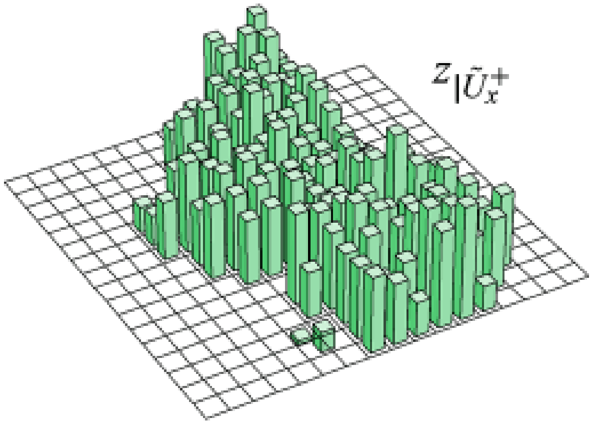

noisy data extracted from the neighborhood

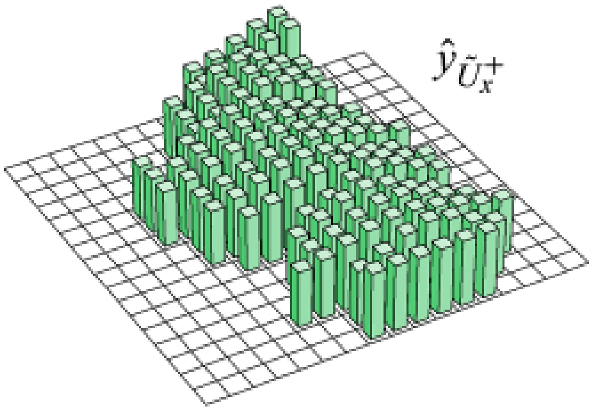

after hard-thresholding in SA-DCT domain

Shape-adaptation: use directional LPA-ICI

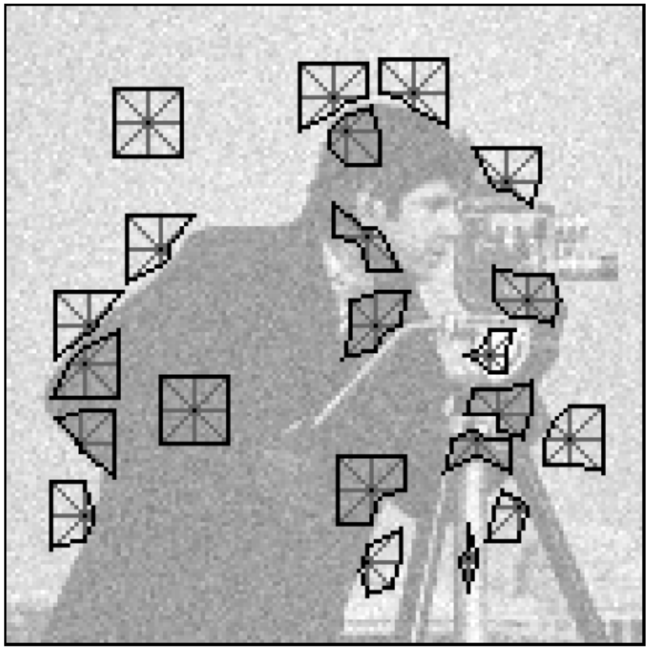

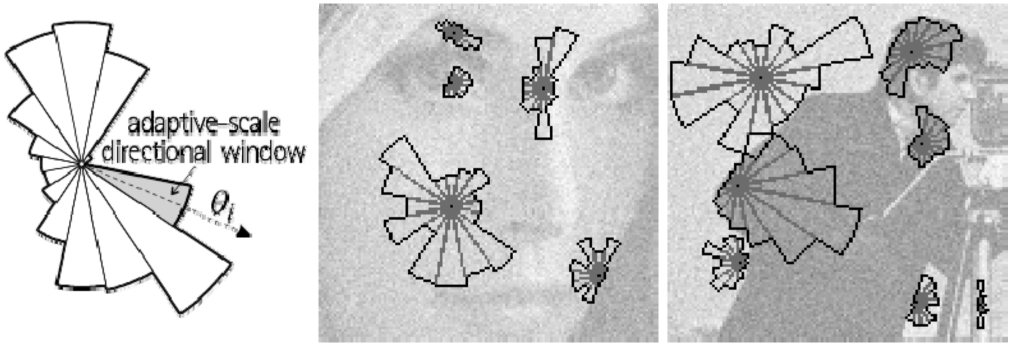

Shape-adaptive DCT image filtering

Pointwise SA-DCT: anisotropic neighborhoods

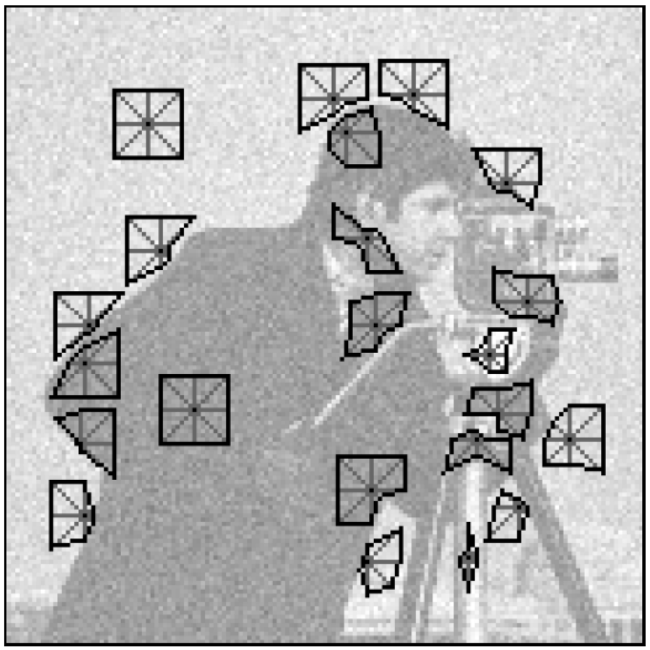

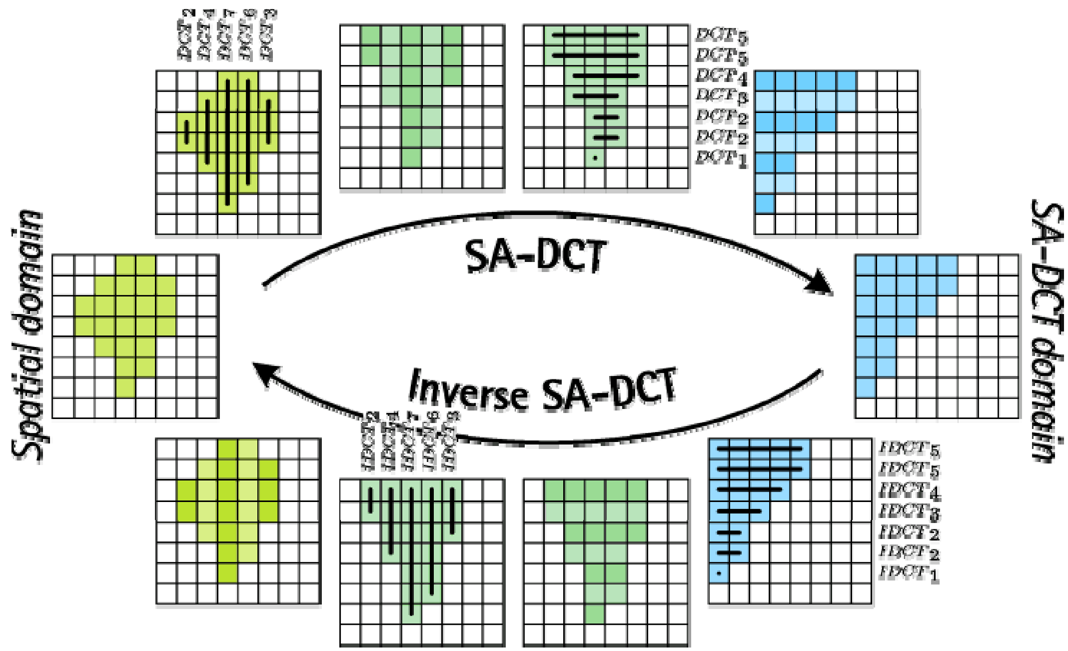

Shape-adaptive DCT image filtering

- Direct generalization of the classical block-DCT (B-DCT);
- On rectangular domains (e.g., squares) the SA-DCT and B-DCT coincide;
- Comparable computational complexity as the separable B-DCT (fast algorithms);
- SA-DCT is part of the MPEG-4 standard;
- Efficient (low-power) hardware implementations available.

Before our work on SA-DCT filtering, the SA-DCT had been used only for image and video compression.

Pointwise SA-DCT: denoising results

A fragment of Cameraman: noisy observation (σ=25, PSNR=20.14dB), BLS-GSM estimate (Portilla et al.) (PSNR=28.35dB), and the proposed Pointwise SA-DCT estimate (PSNR=29.11dB).

Pointwise SA-DCT: deblocking results

JPEG coded Cameraman with 2 different quality levels and the results of post-filtering using SA-DCT

Pointwise SA-DCT: deblurring results

Images blurred & noisy are deblurred & denoised by SA-DCT filter.

Pointwise SA-DCT: extension to color, motivation

Luminance-chrominance decompositions: structural correlation

R
G
B

color transformation

$$\mathbf{A}_{opp} = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{\sqrt{6}} & 0 & \frac{-1}{\sqrt{6}} \\ \frac{1}{3\sqrt{2}} & \frac{-\sqrt{2}}{3} & \frac{1}{3\sqrt{2}} \end{bmatrix}$$

Y

U

V

Pointwise SA-DCT: structural contraint in luminance-chrominance space

Use for all three channels the adaptive neighborhoods defined by the anisotropic LPA-ICI for the luminance channel.

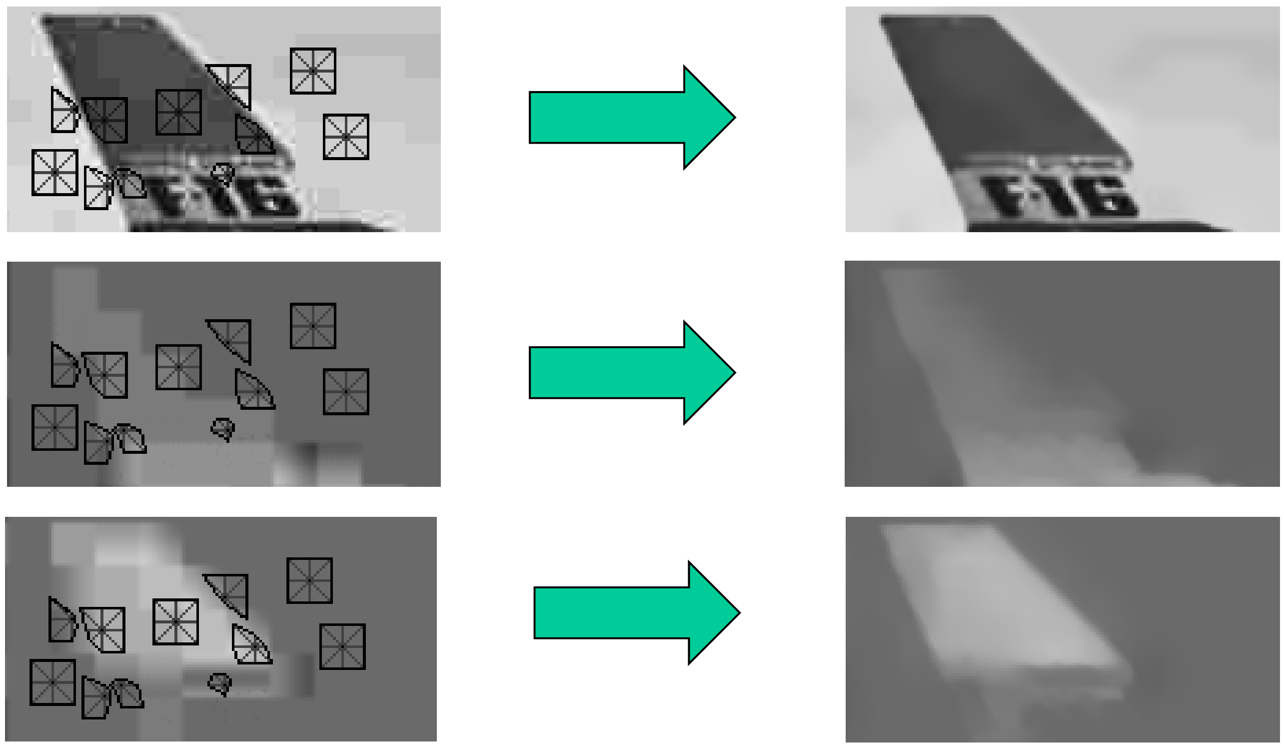

Pointwise SA-DCT: deblocking results

JPEG-compressed
(Q=10, 0.25bpp, PSNR=26.87dB)

Pointwise SA-DCT deblocking
(PSNR=28.30dB)

Pointwise SA-DCT: deblocking results

Pointwise SA-DCT: denoising results



Fragments of the noisy F-16 (σ=30, PSNR=18.59dB), of ProbShrink-MB (Pizurica et al.) estimate (PSNR=30.50dB), and of Pointwise SA-DCT estimate (PSNR=31.59dB).

Block-Matching and 3D filtering (BM3D) denoising algorithm

- Generalizes NL-means and overcomplete transform methods
- Current state-of-the-art denoising method

K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, “Image denoising with block-matching and 3D filtering”, Proc. SPIE Electronic Imaging 2006, Image Process.: Algorithms and Systems V, no. 6064A-30, San Jose (CA), USA, Jan. 2006.

--- , “Image denoising by sparse 3D transform-domain collaborative filtering”, IEEE Trans. Image Process., vol. 16, no. 8, pp. 2080-2095, Aug. 2007.

Block-matching and grouping

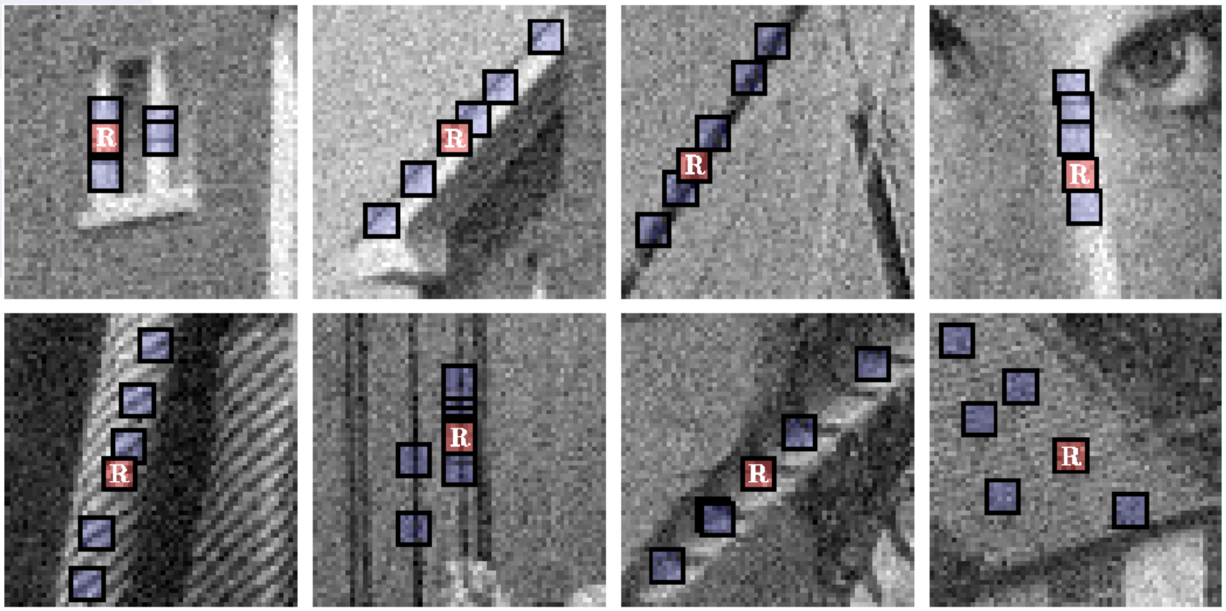

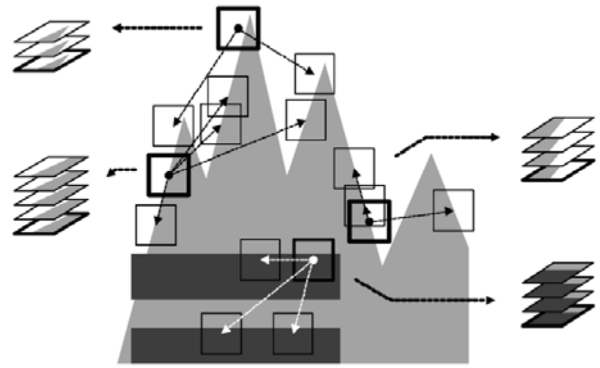

Groups are characterized by both:

- intra-block correlation between the pixels of each grouped block (natural images);
- inter -block correlation between the corresponding pixels of different blocks (grouped block are similar);

BM3D: Collaborative filtering

- Each grouped block collaborates for the filtering of all others, and vice versa.
- Provides individual estimates for all grouped blocks (not necessarily equal).
- Realized as shrinkage in a 3-D transform domain.

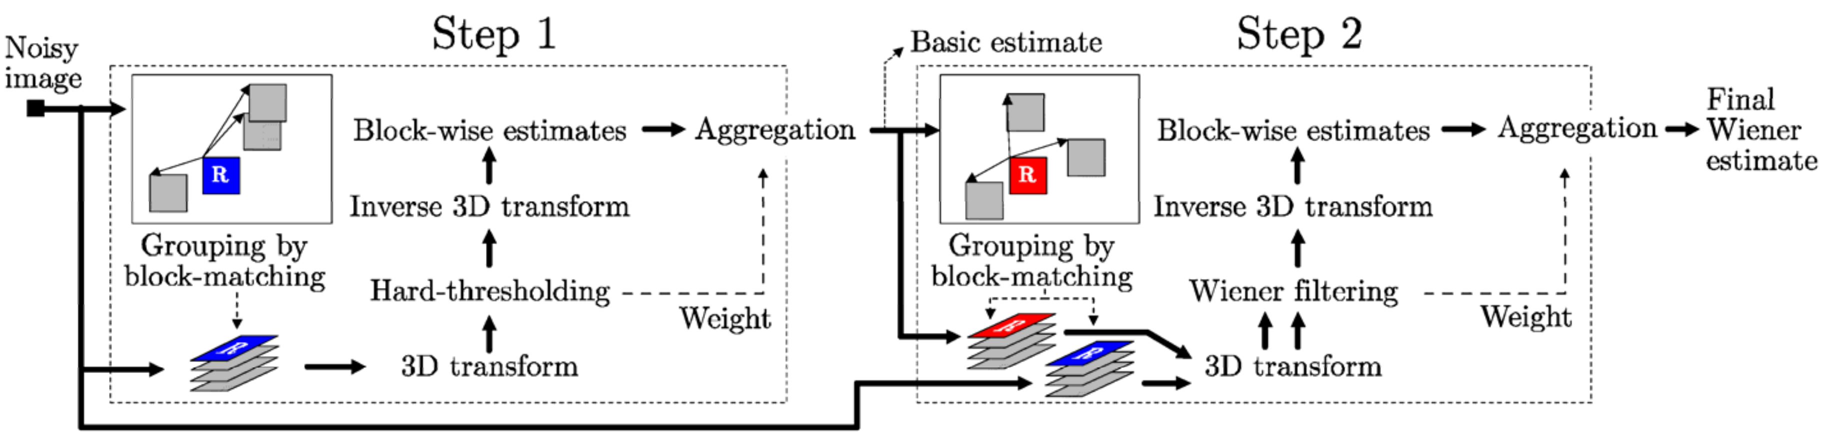

BM3D with Shape-Adaptive PCA (BM3D-SAPCA)

Main ingredients:

- Local Polynomial Approximation - Intersection of Confidence Intervals (LPA-ICI) to adaptively select support for 2-D transform;
- Block-Matching to enable non-locality;
- Shape-Adaptive PCA (SA-PCA);
- Shape-Adaptive DCT low-complexity 2-D transform on arbitrarily-shaped domains (when SA-PCA is not feasible).

K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, .BM3D Image Denoising with Shape-Adaptive Principal Component Analysis., Proc. Workshop on Signal Processing with Adaptive Sparse Structured Representations (SPARS.09), Saint-Malo, France, April 2009.

BM3D-SAPCA

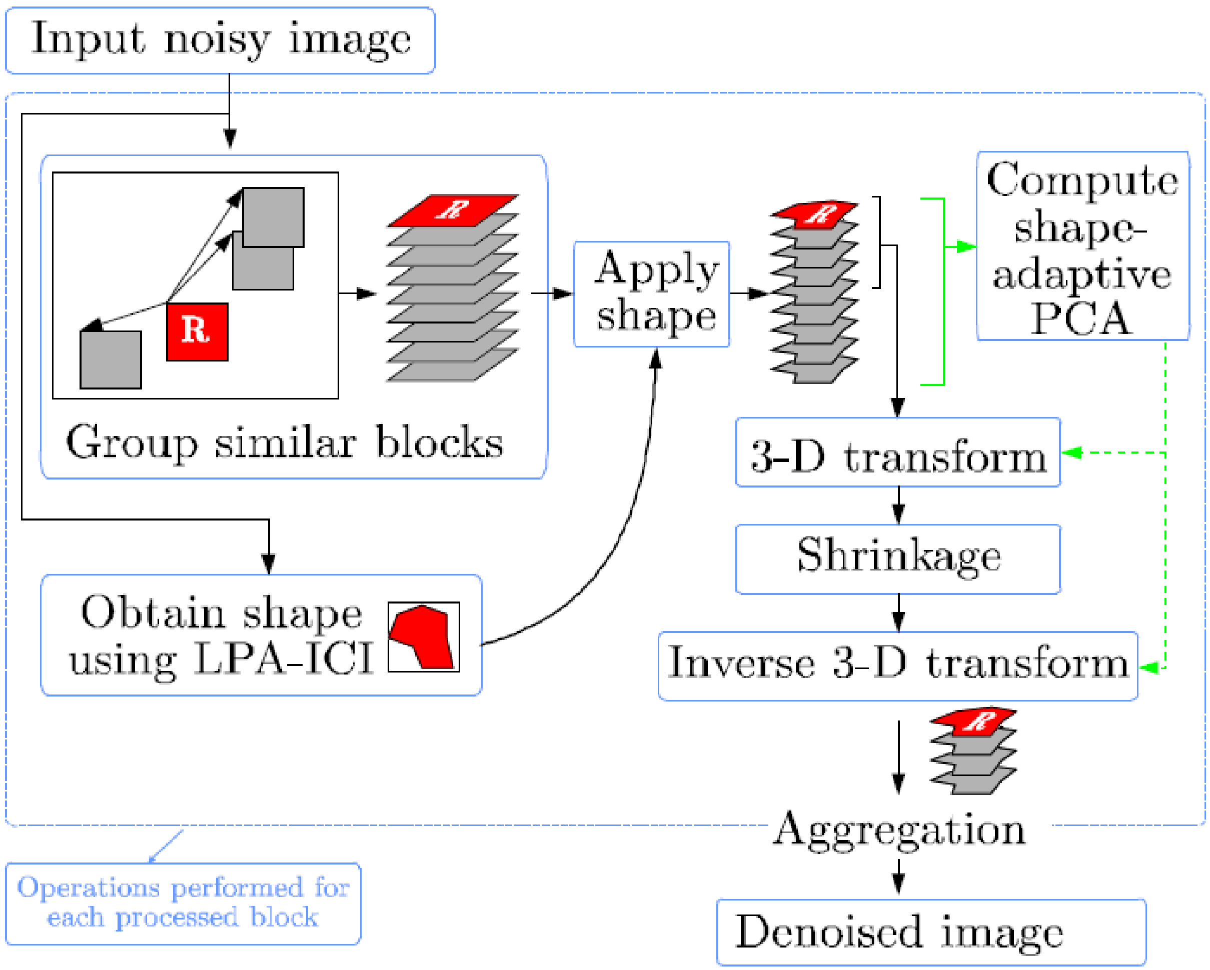
Input noisy image
R
Group similar blocks
Apply shape
Compute shape-adaptive PCA
3-D transform
Shrinkage
Obtain shape using LPA-ICI
Inverse 3-D transform
Aggregation
Operations performed for each processed block
Denoised image

Comparison of BM3D-SAPCA with other filters

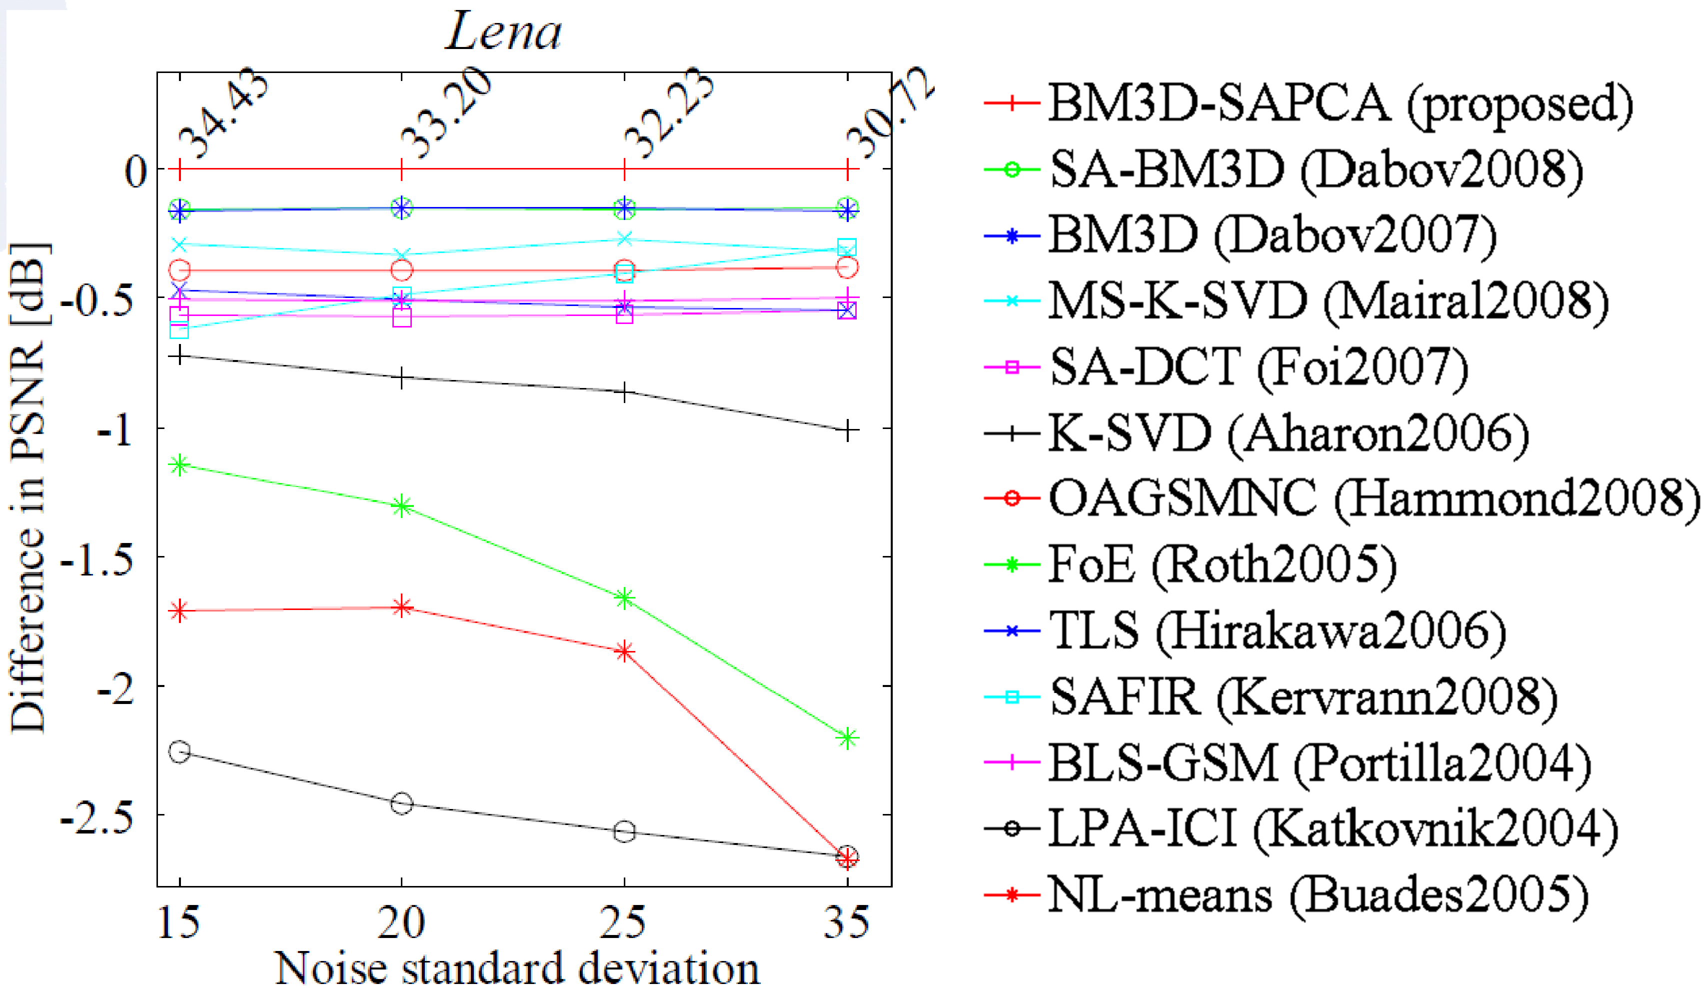

Comparison of BM3D-SAPCA with other filters

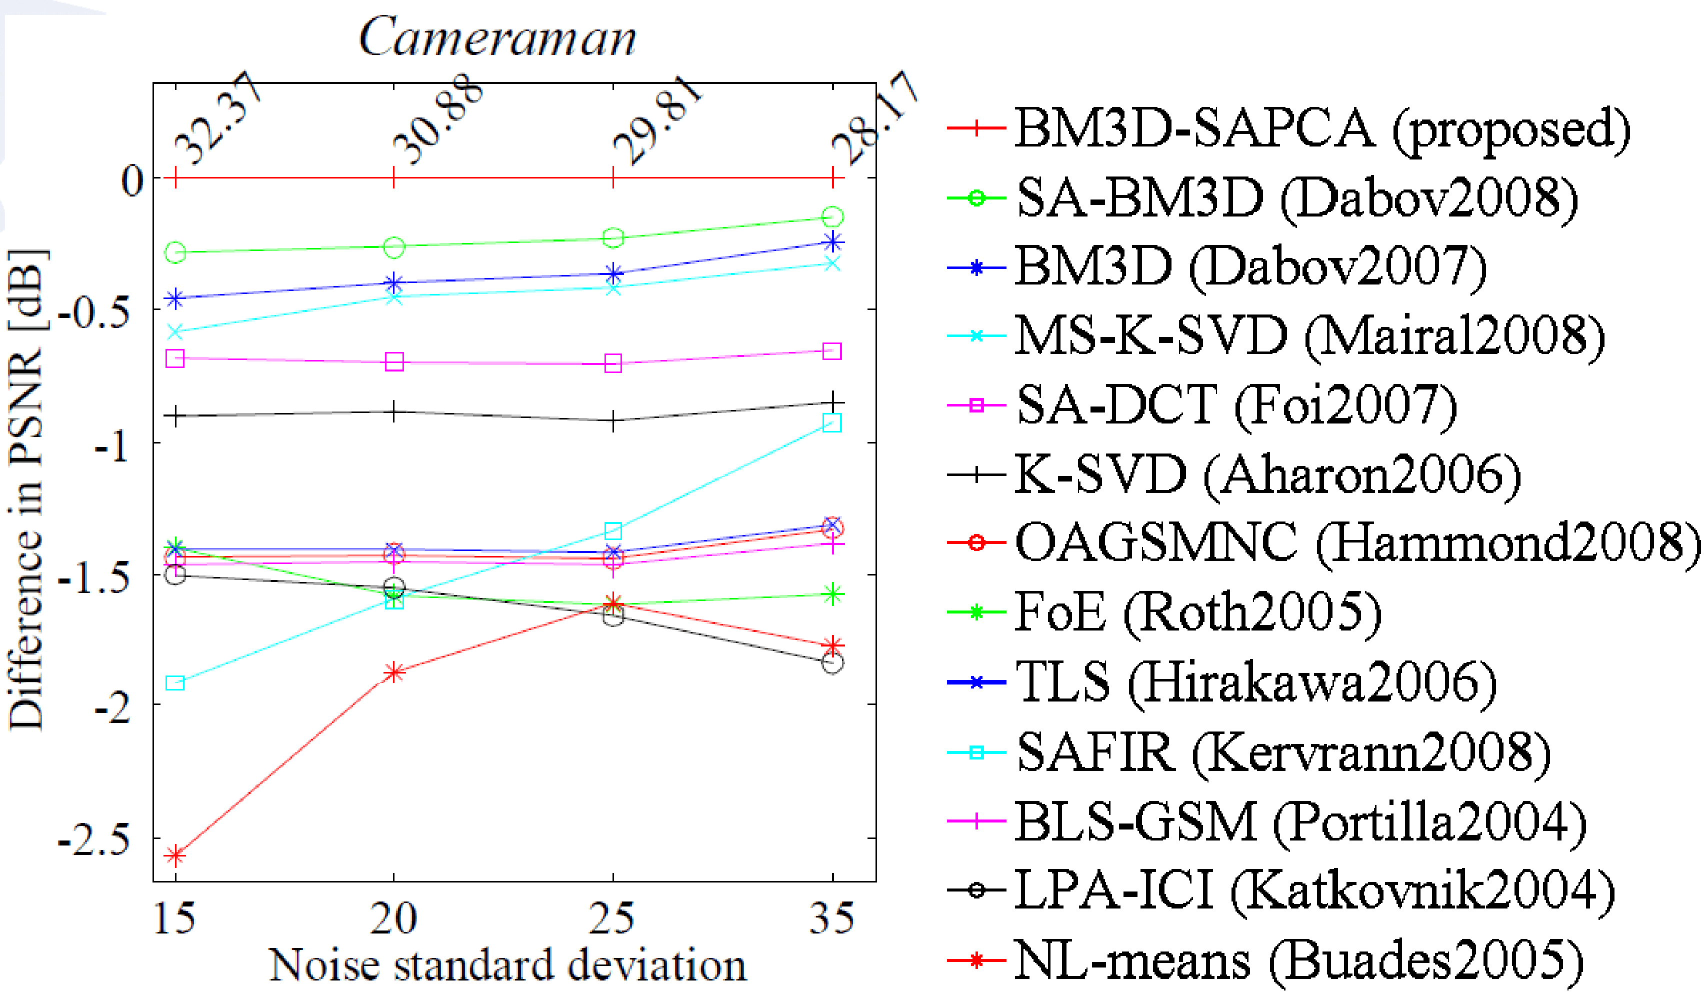

Comparison of BM3D-SAPCA with other filters

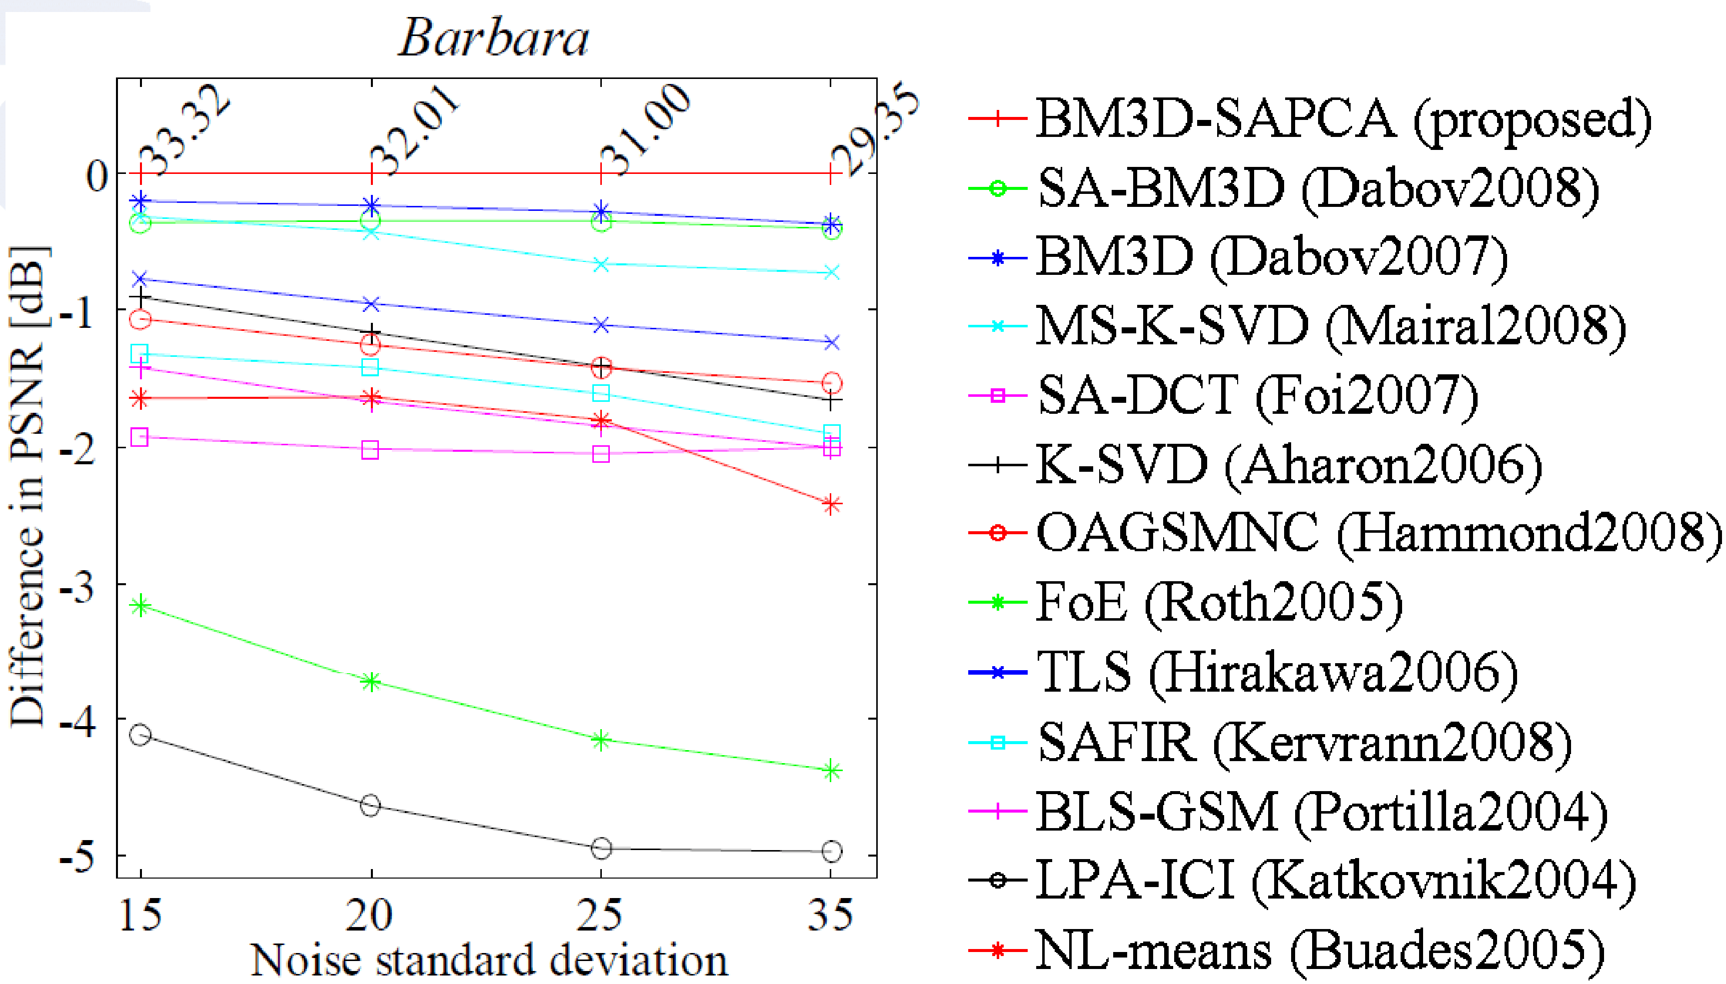

Comparison of BM3D-SAPCA with other filters (PSNR, SSIM)

Original

Noisy, σ = 35

BM3D (27.82, 0.8207)

P.SADCT (27.51, 0.8143) SA-BM3D (28.02, 0.8228) BM3D-SAPCA (28.16, 0.8269)

Interpolation for Bayer Pattern

Original scene

Color Filter Array

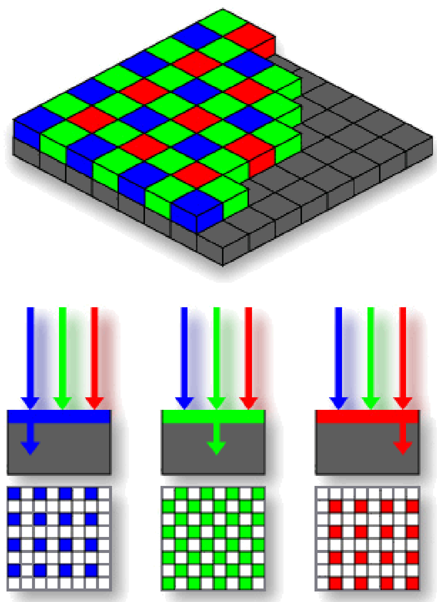

Observation

Color Interpolation

Competitiveness with state-of-the-art techniques

The proposed CFAI technique adapts to spatial properties of an image

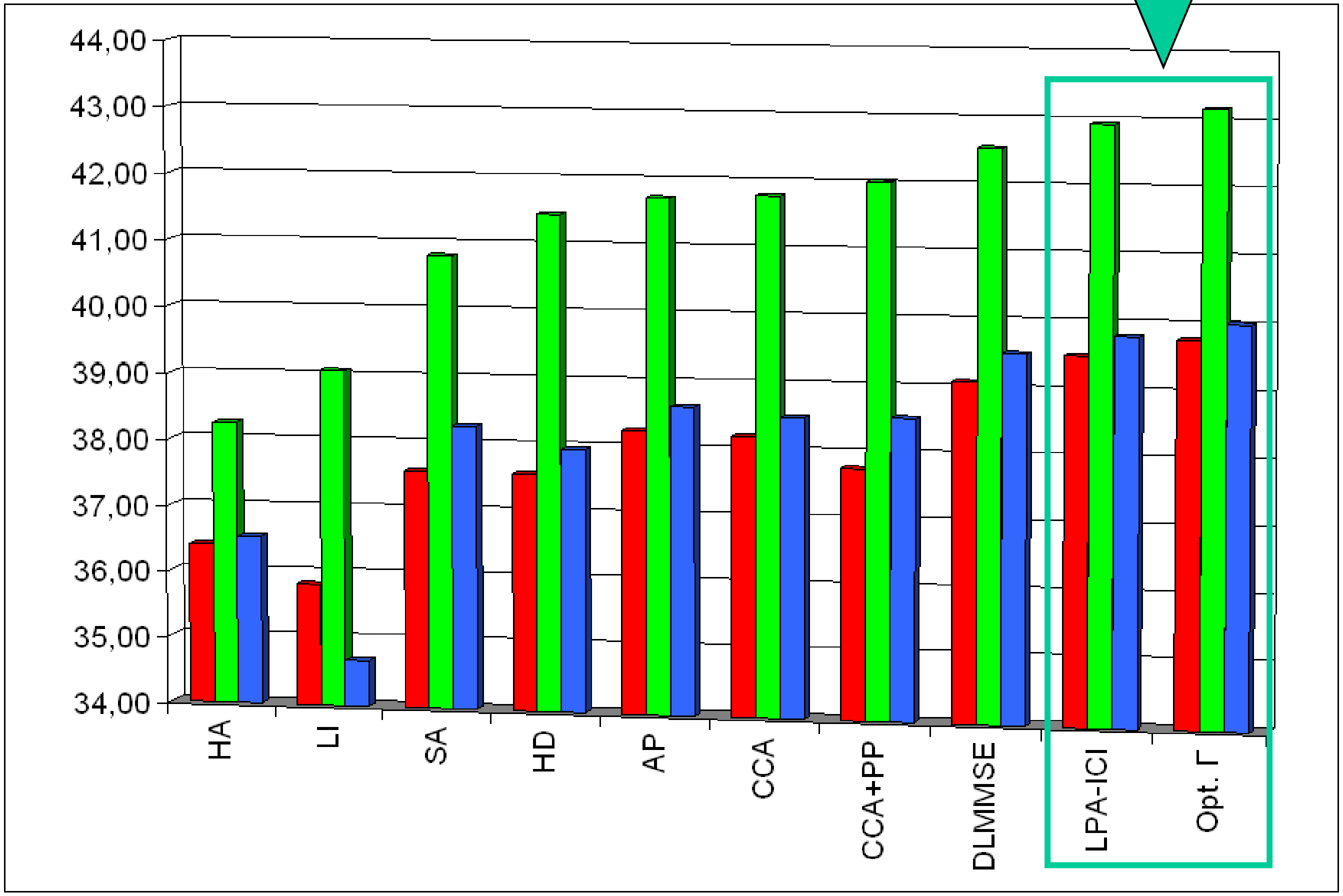

Conventional Approach for Noiseless Data (Hamilton-Adams)

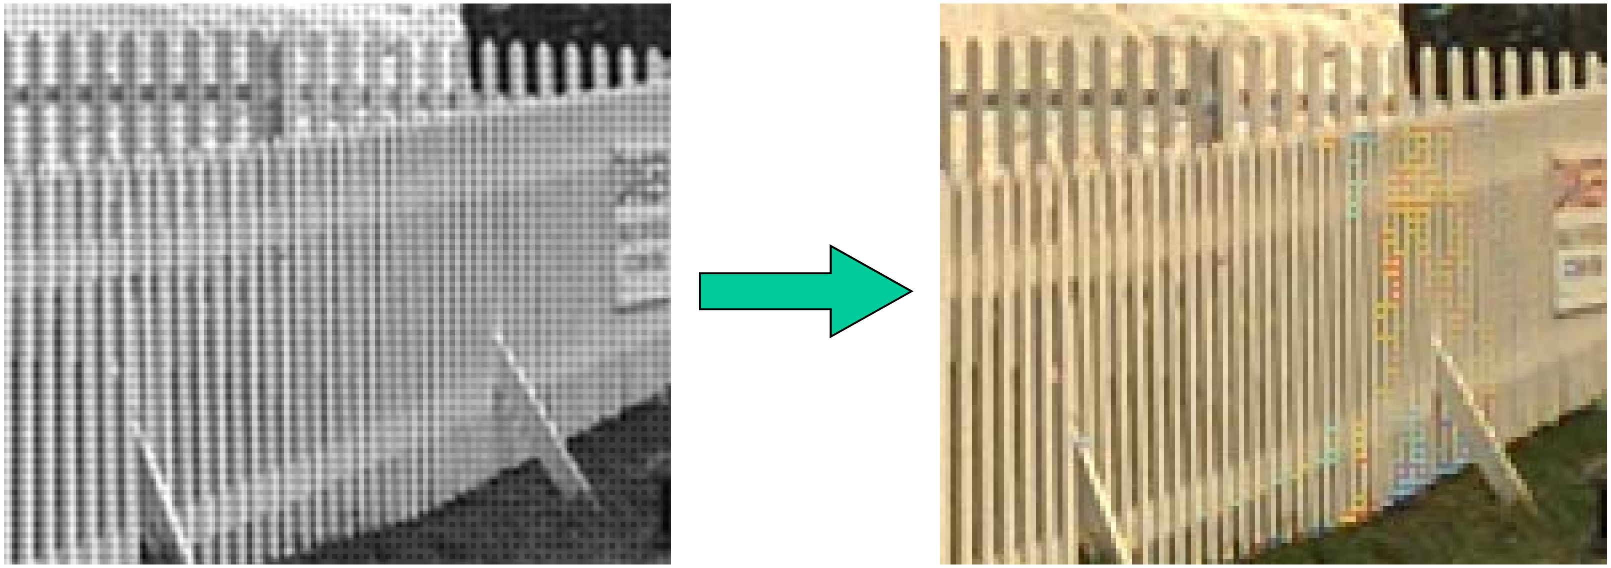

Proposed Approach for Noiseless Data (Spatially-Adaptive LPA-ICI)

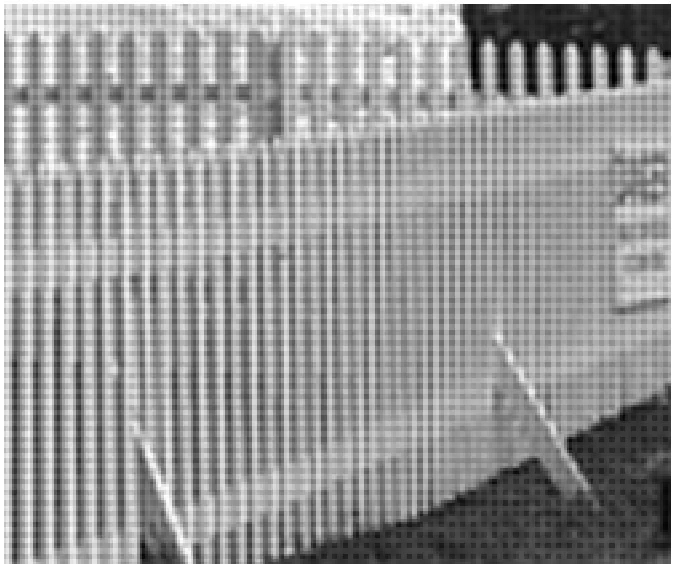

CROSS-COLOR BM3D FILTERING OF NOISY RAWDATA: Examples

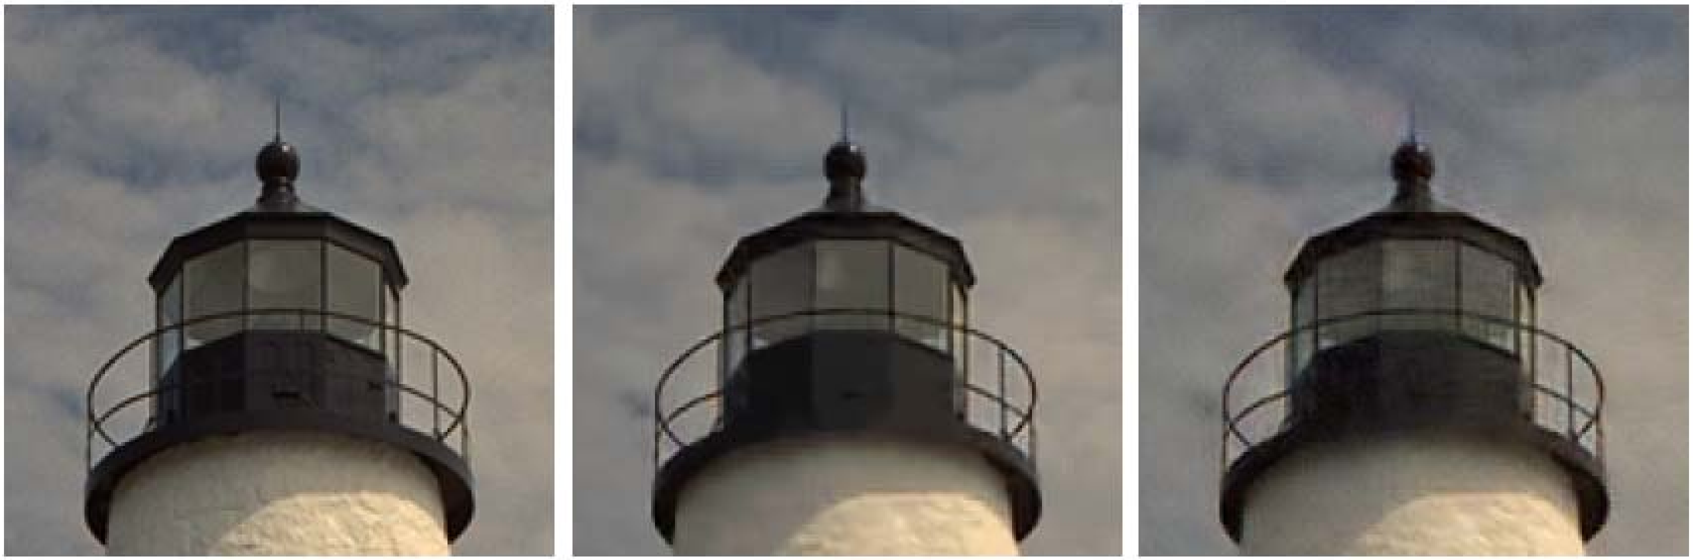

From left to right: ground truth, proposed denoising + interpolation , denoising (Zhang, L. et al (2009) + interpolation . Gaussian noise (σ = 12/255)

CROSS-COLOR BM3D FILTERING OF NOISY RAWDATA: Signal dependent noise

	$\sigma(y)=\sqrt{ay+b}$		$a=0.004,\ b=0.02^2$	
(07)		R	**34.1**	32.7
		G	**34.9**	33.5
		B	**34.3**	33.1
(08)	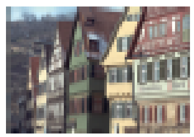	R	**29.5**	28.3
		G	**30.4**	29.2
		B	**29.7**	28.5
(19)		R	**32.0**	31.0
		G	**32.6**	31.7
		B	**32.7**	31.7
(23)	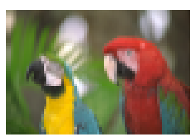	R	**34.2**	33.7
		G	**35.4**	34.7
		B	**34.8**	34.2

PSNR (dB) of denoised and demosaicked images corrupted by signal dependent noise (Computed excluding a 20-pixel border)

First column – using BM3D, second column – using method by Zhang et al(2009)

BM3D for upsampling and super-resolution

Image **upsampling** or **zooming**, can be de.ned as the process of resampling a single low-resolution (LR) image on a high-resolution grid.

The process of combining a sequence of undersampled and degraded low-resolution images in order to produce a single high-resolution image is commonly referred to as a **Super-resolution** (SR) reconstruction.

Modern SR methods (e.g., Protter et al. 2008, Ebrahimi and Vrscay 2008) are based on the nonlocal means (NLM) filtering paradigm (Buades-Coll-Morel, 2005).

• No explicit registration: one-to-one pixel mapping between frames is replaced by a one-to-many mapping.

The BM3D and V-BM3D algorithms share with the NLM the idea of exploiting nonlocal similarity between blocks. However, in (V-)BM3D a more powerful transform-domain modeling is used.

BM3D based superresolution

Multistage iterative reconstruction

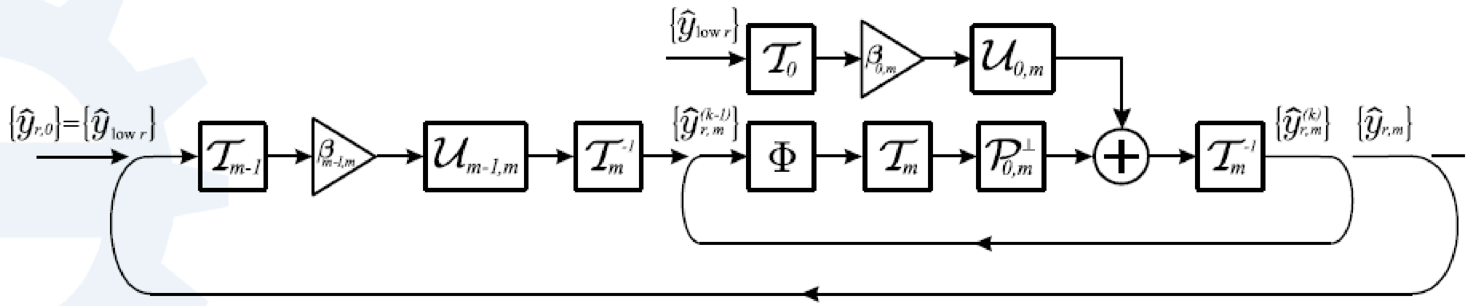

$$
\begin{cases}
\hat{y}_{r,0} = y_{\text{low } r} & \texttt{(algorithm input)} \\
\hat{y}_{r,m} = \hat{y}_{r,m}^{(k_{\text{final } m})} & \texttt{(stage input)} \\
\hat{y}_{r,m}^{(0)} = \mathcal{T}_m^{-1}\left(\mathcal{U}_{m-1,m}\left(\beta_{m-1,m}\mathcal{T}_{m-1}\left(\hat{y}_{r,m-1}\right)\right)\right) & \\
\hat{y}_{r,m}^{(k)} = \mathcal{T}_m^{-1}\left(\mathcal{U}_{0,m}\left(\beta_{0,m}\mathcal{T}_0\left(y_{\text{low } r}\right)\right) + \mathcal{P}_{0,m}^{\perp}\left(\mathcal{T}_m\left(\Phi\left(r, \left\{\hat{y}_{r,m}^{(k-1)}\right\}_{r=1}^{R}, \sigma_{k,m}\right)\right)\right)\right) &
\end{cases}
$$

m	stage number	$m = 1, \ldots, M$
k	iteration number	$k = 0, \ldots, k_{\text{final } m}$
$\hat{y}_{r,m}^{(k)}$	estimate for $\hat{y}_r$ on iter. k of stage m	
$\mathcal{T}_m$	transform	
Φ	spatially adaptive filter (V-BM3D)	
$\sigma_{k,m}$	parameter controlling the strength of the filter	$\sigma_{k,m} = \sigma_{k,m-1} - \Delta_m$

Image upsampling x 4 (pixel replication)

Image upsampling x 4 in wavelet domain (Danielyan et al. EUSIPCO 2008)

Video superresolution comparison with (Protter et. al.)

Nearest neighbor Ground truth Protter et. al. Proposed

1. M. Protter, M. Elad, H. Takeda, and P. Milanfar, .Generalizing the Non-Local-Means to Super-Resolution Reconstruction., IEEE Trans. Image Process., 2008.

2. A. Danielyan, A. Foi, V. Katkovnik, and K. Egiazarian, .Image upsampling via spatially adaptive block-matching filtering, EUSIPCO2008, Lausanne, Switzerland, Aug. 2008.

Examples: Video denoising using V-BM3D

Examples: Video denoising using V-BM3D

Examples: Video denoising using V-BM3D