Two-stage image denoising by principal component analysis with local pixel grouping

2014.2.13

Hsin-Hui Chen

Outline

- Principal component analysis (PCA)
- Local pixel grouping (LPG)
- LPG-PCA

LPG-PCA

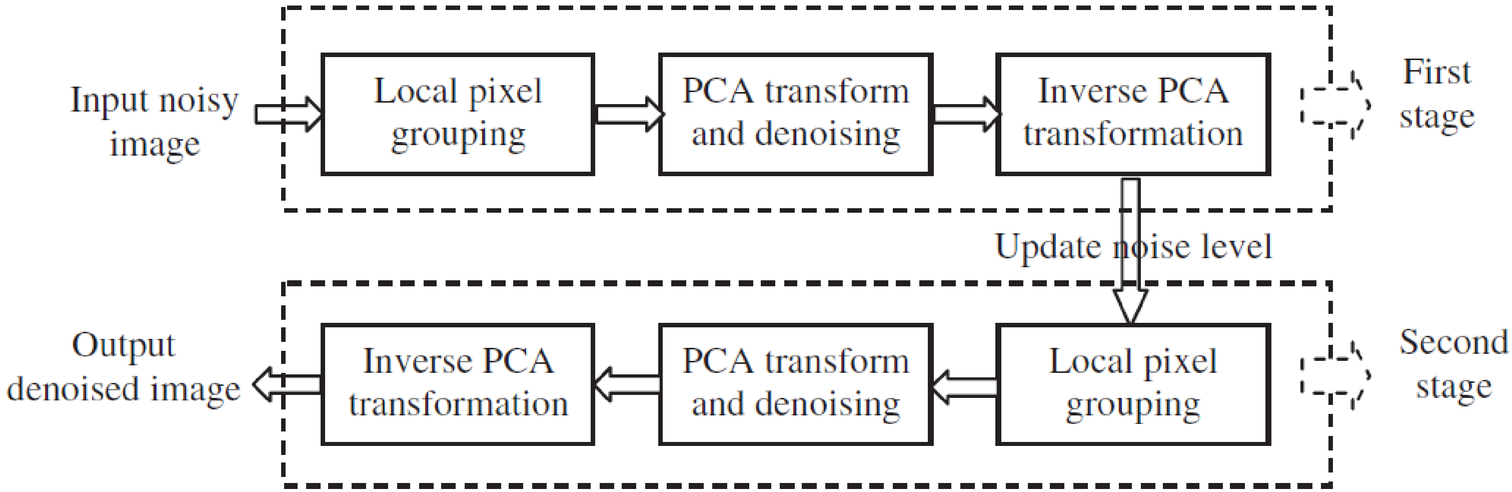

Input noisy image
Local pixel grouping
PCA transform and denoising
Inverse PCA transformation
First stage
Update noise level
Output denoised image
Inverse PCA transformation
PCA transform and denoising
Local pixel grouping
Second stage

LPG-PCA

Denote by $\mathbf{x} = [x_1 \quad x_2 \quad \ldots \quad x_m]^T$ an m-component vector variable

$$\mathbf{X} = \begin{bmatrix} x_1^1 & x_1^2 & \cdots & x_1^n \\ x_2^1 & x_2^2 & \cdots & x_2^n \\ \vdots & \vdots & \vdots & \vdots \\ x_m^1 & x_m^2 & \cdots & x_m^n \end{bmatrix}$$

LPG-PCA

The ith row of sample matrix $\mathbf{X}$,

$$X_i = [x_i^1 \;\; x_i^2 \;\; \ldots \;\; x_i^n]$$

The mean value of X_i is calculated as

$$\mu_i = \frac{1}{n}\sum_{j=1}^{n} X_i(j)$$

the sample vector X_i is centralized as

$$\overline{X}_i = X_i - \mu_i = [\overline{x}_i^1 \;\; \overline{x}_i^2 \;\; \ldots \;\; \overline{x}_i^n]$$

LPG-PCA

The goal of PCA is to find an orthonormal transformation matrix $\mathbf{P}$ to de-correlate $\overline{\mathbf{X}}$, i.e. $\overline{\mathbf{Y}} = \mathbf{P}\overline{\mathbf{X}}$ so that the co-variance matrix of $\overline{\mathbf{Y}}$ is diagonal.

LPG-PCA

Finally, the co-variance matrix of the centralized dataset is calculated as

$$\boldsymbol{\Omega} = \frac{1}{n}\overline{\mathbf{X}}\overline{\mathbf{X}}^T$$

Since the co-variance matrix $\boldsymbol{\Omega}$ is symmetrical, it can be written as:

$$\boldsymbol{\Omega} = \boldsymbol{\Phi}\boldsymbol{\Lambda}\boldsymbol{\Phi}^{\mathrm{T}}$$

where $\boldsymbol{\Phi} = [\phi_1 \;\; \phi_2 \;\; \ldots \;\; \phi_m]$ is the $m \times m$ orthonormal eigenvector matrix and $\boldsymbol{\Lambda} = diag\{\lambda_1, \lambda_2, \ldots, \lambda_m\}$ is the diagonal eigenvalue matrix with $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_m$.

LPG-PCA

The terms $\phi_1, \phi_2, \ldots, \phi_m$ and $\lambda_1, \lambda_2, \ldots, \lambda_m$ are the eigenvectors and eigenvalues of $\mathbf{\Omega}$. By setting

$$\mathbf{P} = \mathbf{\Phi}^{\mathbf{T}}$$

$\overline{\mathbf{X}}$ can be decorrelated, i.e. $\overline{\mathbf{Y}} = \mathbf{P}\overline{\mathbf{X}}$ and $\mathbf{\Lambda} = (1/n)\overline{\mathbf{Y}}\overline{\mathbf{Y}}^T$.

LPG-PCA

An important property of PCA is that it fully de-correlates the original dataset $\overline{\mathbf{X}}$. Generally speaking, the energy of a signal will concentrate on a small subset of the PCA transformed dataset, while the energy of noise will evenly spread over the whole dataset. Therefore, the signal and noise can be better distinguished in the PCA domain.

LPG-PCA

for an underlying pixel to be denoised, we set a $K \times K$ window centered on it and denote by $\mathbf{x} = [x_1 \quad \dots \quad x_m]^T$, $m = K^2$, the vector containing all the components within the window.

Since the observed image is noise corrupted, we denote by

$$\mathbf{x}_\upsilon = \mathbf{x} + \boldsymbol{\upsilon}$$

the noisy vector of $\mathbf{x}$, where $\mathbf{x}_\upsilon = [x_1^\upsilon \quad \dots \quad x_m^\upsilon]^T$, $\boldsymbol{\upsilon} = [\upsilon_1 \quad \dots \quad \upsilon_m]^T$

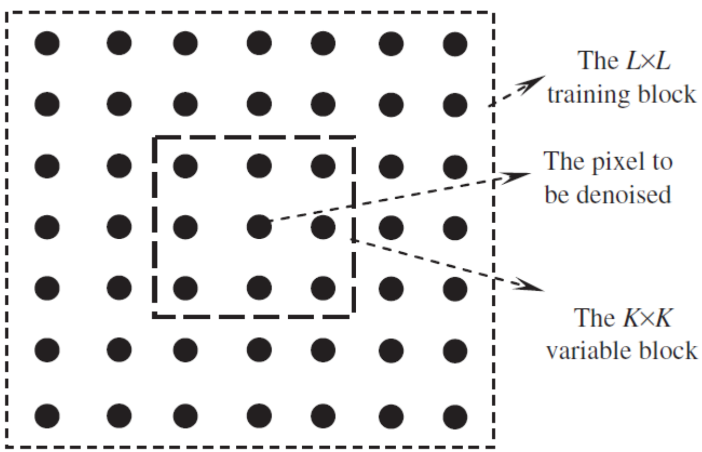

LPG-PCA

Next we centralize dataset $\mathbf{X}_\upsilon$.

$$\overline{\mathbf{X}}_\upsilon = \overline{\mathbf{X}} + \mathbf{V}$$

by computing the covariance matrix of $\overline{\mathbf{X}}$, denoted by $\mathbf{\Omega}_{\overline{\mathbf{x}}}$, the PCA transformation matrix $\mathbf{P}_{\overline{\mathbf{x}}}$ can be obtained.

LPG-PCA

However, the available dataset $\mathbf{X}_\upsilon$ is noise corrupted so that $\mathbf{\Omega}_{\overline{\mathbf{X}}}$ cannot be directly computed. With the linear model (3.5), we have

$$\mathbf{\Omega}_{\overline{\mathbf{X}}_\upsilon} = \frac{1}{n}\overline{\mathbf{X}}_\upsilon\overline{\mathbf{X}}_\upsilon^T = \frac{1}{n}\left(\overline{\mathbf{X}}\overline{\mathbf{X}}^T + \overline{\mathbf{X}}\mathbf{V}^T + \mathbf{V}\overline{\mathbf{X}}^T + \mathbf{V}\mathbf{V}^T\right)$$

Since $\overline{\mathbf{X}}$ and $\mathbf{V}$ are uncorrelated, items $\overline{\mathbf{X}}\mathbf{V}^T$ and $\mathbf{V}\overline{\mathbf{X}}^T$ will be nearly zero matrices and thus:

$$\mathbf{\Omega}_{\overline{\mathbf{X}}_\upsilon} \approx \frac{1}{n}\left(\overline{\mathbf{X}}\overline{\mathbf{X}}^T + \mathbf{V}\mathbf{V}^T\right) = \mathbf{\Omega}_{\overline{\mathbf{X}}} + \mathbf{\Omega}_\upsilon$$

where $\mathbf{\Omega}_{\overline{\mathbf{X}}} = (1/n)\overline{\mathbf{X}}\overline{\mathbf{X}}^T$ and $\mathbf{\Omega}_\upsilon = (1/n)\mathbf{V}\mathbf{V}^T$.

LPG-PCA

we can decompose $\boldsymbol{\Omega}_{\overline{\mathbf{x}}}$ as

$$\boldsymbol{\Omega}_{\overline{\mathbf{x}}} = \boldsymbol{\Phi}_{\overline{\mathbf{x}}} \boldsymbol{\Lambda}_{\overline{\mathbf{x}}} \boldsymbol{\Phi}_{\overline{\mathbf{x}}}^{\mathbf{T}}$$

where $\boldsymbol{\Phi}_{\overline{\mathbf{x}}}$ is the $m \times m$ orthonormal eigenvector matrix and $\boldsymbol{\Lambda}_{\overline{\mathbf{x}}}$ is the diagonal eigenvalue matrix.

LPG-PCA

Since $\mathbf{\Phi}_{\overline{\mathbf{x}}}$ is an orthonormal matrix, we can write $\mathbf{\Omega}_{\upsilon}$ as

$$\mathbf{\Omega}_{\upsilon} = (\sigma^2\mathbf{I})\mathbf{\Phi}_{\overline{\mathbf{x}}}\mathbf{\Phi}_{\overline{\mathbf{x}}}^{\mathbf{T}} = \mathbf{\Phi}_{\overline{\mathbf{x}}}(\sigma^2\mathbf{I})\mathbf{\Phi}_{\overline{\mathbf{x}}}^{\mathbf{T}} = \mathbf{\Phi}_{\overline{\mathbf{x}}}\mathbf{\Omega}_{\upsilon}\mathbf{\Phi}_{\overline{\mathbf{x}}}^{\mathbf{T}}$$

Thus we have

$$\begin{aligned}\mathbf{\Omega}_{\overline{\mathbf{x}}_{\upsilon}} &= \mathbf{\Omega}_{\overline{\mathbf{x}}} + \mathbf{\Omega}_{\upsilon} = \mathbf{\Phi}_{\overline{\mathbf{x}}}\mathbf{\Lambda}_{\overline{\mathbf{x}}}\mathbf{\Phi}_{\overline{\mathbf{x}}}^{\mathbf{T}} + \mathbf{\Phi}_{\overline{\mathbf{x}}}(\sigma^2\mathbf{I})\mathbf{\Phi}_{\overline{\mathbf{x}}}^{\mathbf{T}} \\ &= \mathbf{\Phi}_{\overline{\mathbf{x}}}(\mathbf{\Lambda}_{\overline{\mathbf{x}}} + \sigma^2\mathbf{I})\mathbf{\Phi}_{\overline{\mathbf{x}}}^{\mathbf{T}} = \mathbf{\Phi}_{\overline{\mathbf{x}}}\mathbf{\Lambda}_{\overline{\mathbf{x}}_{\upsilon}}\mathbf{\Phi}_{\overline{\mathbf{x}}}^{\mathbf{T}}\end{aligned}$$

LPG-PCA

the orthonormal PCA transformation matrix for $\overline{\mathbf{X}}$ is set as

$$\mathbf{P}_{\overline{\mathbf{X}}} = \mathbf{\Phi}_{\overline{\mathbf{X}}}^{\mathrm{T}}$$

Applying $\mathbf{P}_{\overline{\mathbf{X}}}$ to dataset $\overline{\mathbf{X}}_{\upsilon}$, we have

$$\overline{\mathbf{Y}}_{\upsilon} = \mathbf{P}_{\overline{\mathbf{X}}}\overline{\mathbf{X}}_{\upsilon} = \mathbf{P}_{\overline{\mathbf{X}}}\overline{\mathbf{X}} + \mathbf{P}_{\overline{\mathbf{X}}}\mathbf{V} = \overline{\mathbf{Y}} + \mathbf{V}_{\mathbf{Y}}$$

LPG-PCA

Since $\overline{\mathbf{Y}}$ and noise $\mathbf{V_Y}$ are uncorrelated, we can easily derive that the covariance matrix of $\overline{\mathbf{Y}}_\upsilon$ is

$$\mathbf{\Omega}_{\overline{\mathbf{y}}_\upsilon} = \frac{1}{n}\overline{\mathbf{Y}}_\upsilon\overline{\mathbf{Y}}_\upsilon^T = \mathbf{\Omega}_{\overline{\mathbf{y}}} + \mathbf{\Omega}_{\upsilon_\mathbf{y}}$$

where $\mathbf{\Omega}_{\overline{\mathbf{y}}} = \mathbf{\Lambda}_{\overline{\mathbf{x}}}$ is the covariance matrix of decorrelated dataset $\overline{\mathbf{Y}}$ and $\mathbf{\Omega}_{\upsilon_\mathbf{y}} = \mathbf{P}_{\overline{\mathbf{x}}}\mathbf{\Omega}_\upsilon\mathbf{P}_{\overline{\mathbf{x}}}^T$ is the covariance matrix of noise dataset $\mathbf{V_Y}$.

LPG-PCA

Since $\overline{\mathbf{Y}}_{\upsilon}$ is centralized, the LMMSE of $\overline{Y}_k$, i.e. the kth row of $\overline{\mathbf{Y}}$, is obtained as

$$\hat{\overline{Y}}_k = w_k \cdot \vec{\overline{Y}}_{\upsilon}^{k}$$

where the shrinkage coefficient

$$w_k = \mathbf{\Omega}_{\overline{\mathbf{y}}}(k,k) / \mathbf{\Omega}_{\overline{\mathbf{y}}}(k,k) + \mathbf{\Omega}_{\mathbf{\upsilon_y}}(k,k)$$

LPG-PCA

In implementation we first calculate $\boldsymbol{\Omega}_{\bar{\mathbf{y}}_\upsilon}$ from the available noisy dataset $\bar{\mathbf{Y}}_\upsilon$ and then estimate $\boldsymbol{\Omega}_{\bar{\mathbf{y}}}(k,k)$ by $\boldsymbol{\Omega}_{\bar{\mathbf{y}}}(k,k) = \boldsymbol{\Omega}_{\bar{\mathbf{y}}_\upsilon}(k,k) - \boldsymbol{\Omega}_{\upsilon_\mathbf{y}}(k,k)$.

In flat zones, it is often that $\boldsymbol{\Omega}_{\bar{\mathbf{y}}_\upsilon}(k,k) - \boldsymbol{\Omega}_{\upsilon_\mathbf{y}}(k,k) \le 0$, and then we set $\boldsymbol{\Omega}_{\bar{\mathbf{y}}}(k,k) = 0$. In this case w_k will be exactly 0 and all the noise in $\bar{Y}_\upsilon$ will be removed.

LPG-PCA

Denote by $\hat{\overline{\mathbf{Y}}}$ the matrix of all $\hat{\overline{Y}}_k$. By transforming $\hat{\overline{\mathbf{Y}}}$ back to the time domain, we obtain the denoised result of $\overline{\mathbf{X}}_\upsilon$ as

$$\hat{\overline{\mathbf{X}}} = \mathbf{P}_{\overline{\mathbf{x}}}^T \cdot \hat{\overline{\mathbf{Y}}}$$

LPG-PCA

Adding the mean values μ_k back to $\overline{\mathbf{X}}$ gives the denoised dataset $\hat{\mathbf{X}}$. The estimation of the central block $\vec{x}_0$, denoted as $\hat{\vec{x}}_0$, can then be extracted from $\hat{\mathbf{X}}$ and finally the denoised result of the underlying central pixel can be extracted from $\hat{\vec{x}}_0$.

Applying the above procedure to each pixel leads to the full denoised image of I_v.

LPG-PCA

Denote by $\hat{I}$ the denoised version of noisy image I_v in the first stage. We can write $\hat{I}$ as $\hat{I} = I + \upsilon_s$, where υ_s is the residual in the denoised image.

LPG-PCA

need to estimate the level of υ_s, denoted by $\sigma_s = \sqrt{E[\upsilon_s^2]}$, and input it to the second stage of LPG-PCA denoising. Here we estimate σ_s based on the difference between $\hat{I}$ and I_υ. Let

$$\tilde{I} = I_\upsilon - \hat{I} = \upsilon - \upsilon_s$$

We have:

$$\begin{aligned} E[\tilde{I}^2] &= E[\upsilon^2] + E[\upsilon_s^2] - 2E[\upsilon \cdot \upsilon_s] \\ &= \sigma^2 + \sigma_s^2 - 2E[\upsilon \cdot \upsilon_s] \end{aligned}$$

LPG-PCA

We approximately view υ_s as the smoothed version of noise υ, and it contains mainly the low frequency component of υ. Let $\tilde{\upsilon} = \upsilon - \upsilon_s$ be their difference and $\tilde{\upsilon}$ contains mainly the high frequency component of υ. There is $E[\upsilon \cdot \upsilon_s] = E[\tilde{\upsilon} \cdot \upsilon_s] + E[\upsilon_s^2]$.

LPG-PCA

$$E[\upsilon \cdot \upsilon_s] = E[\tilde{\upsilon} \cdot \upsilon_s] + E[\upsilon_s^2] \approx E[\upsilon_s^2] = \sigma_s^2$$

$$\sigma_s^2 \approx \sigma^2 - E[\tilde{I}^2]$$

LPG-PCA

In practice, υ_s will include not only the noise residual but also the estimation error of noiseless image *I*. Therefore, in implementation we let

$$\sigma_s = c_s \cdot \sqrt{\sigma^2 - E[\tilde{I}^2]}$$

where $c_s < 1$ is a constant. We experimentally found that setting c_s around 0.35 can lead to satisfying denoising results for most of the testing images.

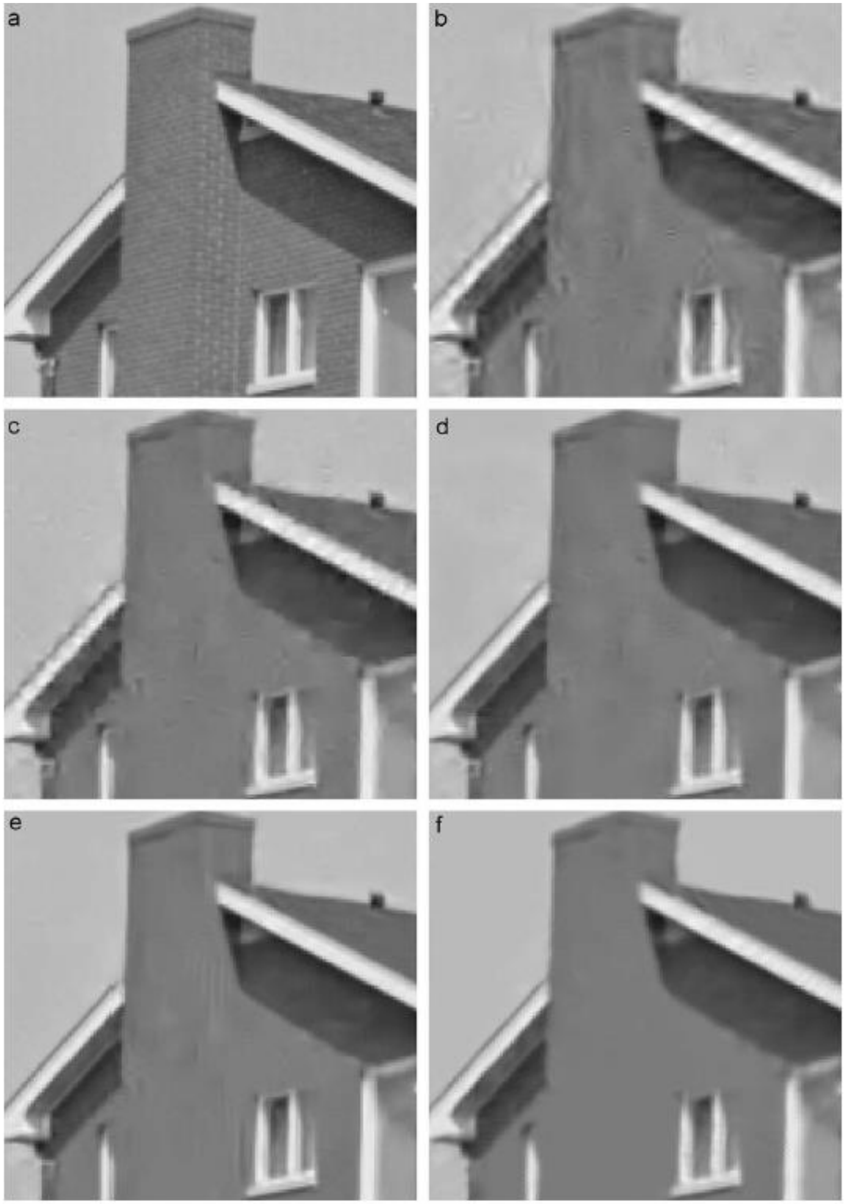
a
b
c
d
e
f

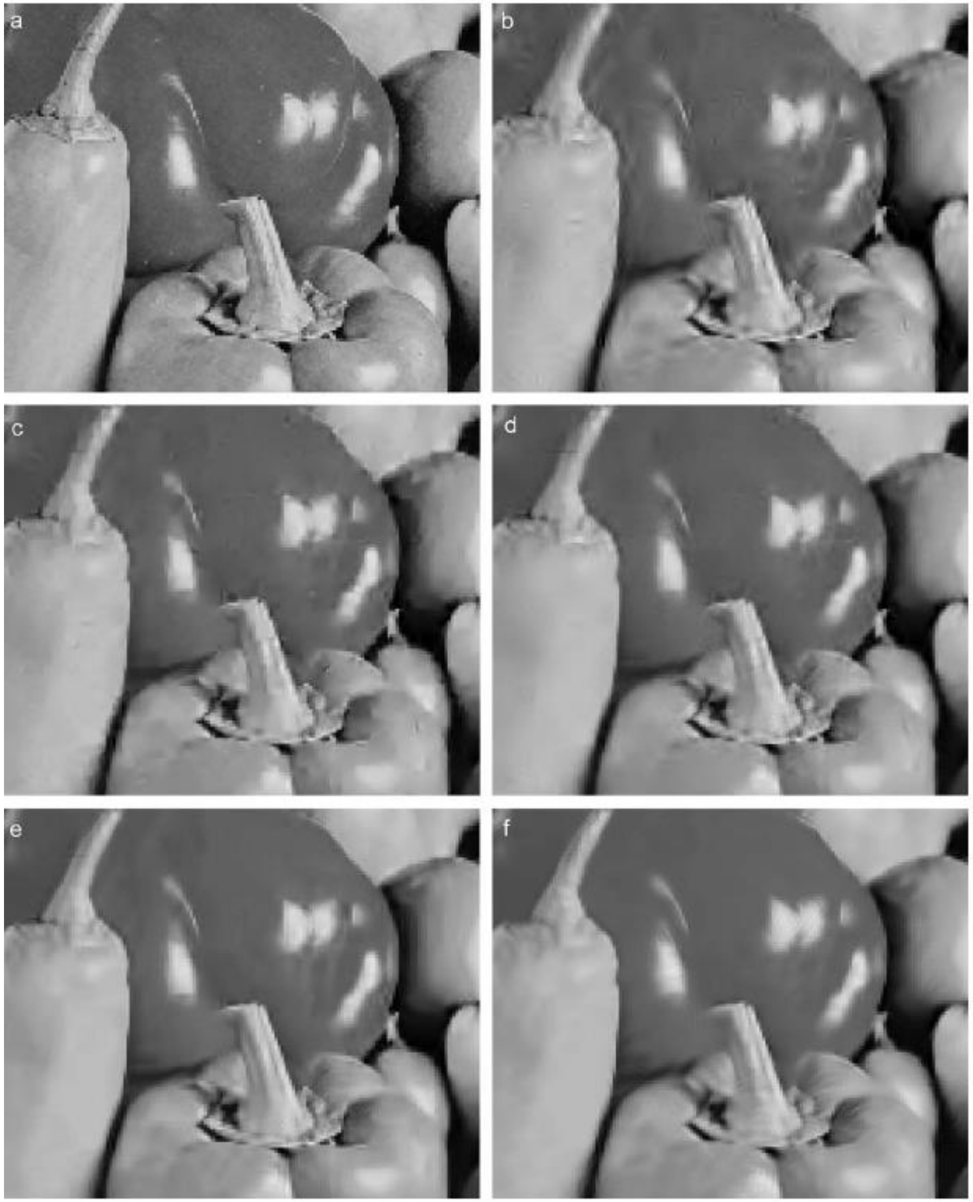
a
b
c
d
e
f

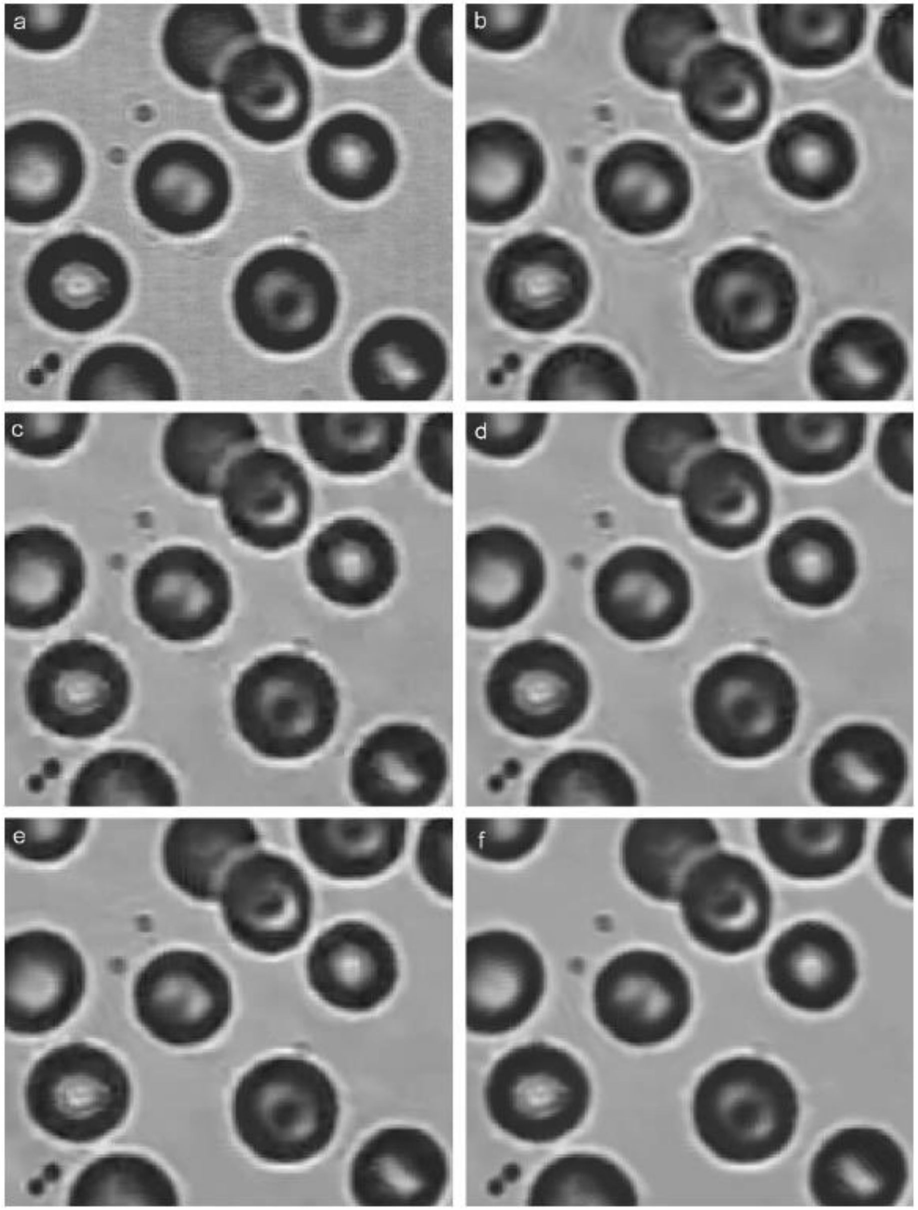
a
b
c
d
e
f


a
b
c
d
e
f

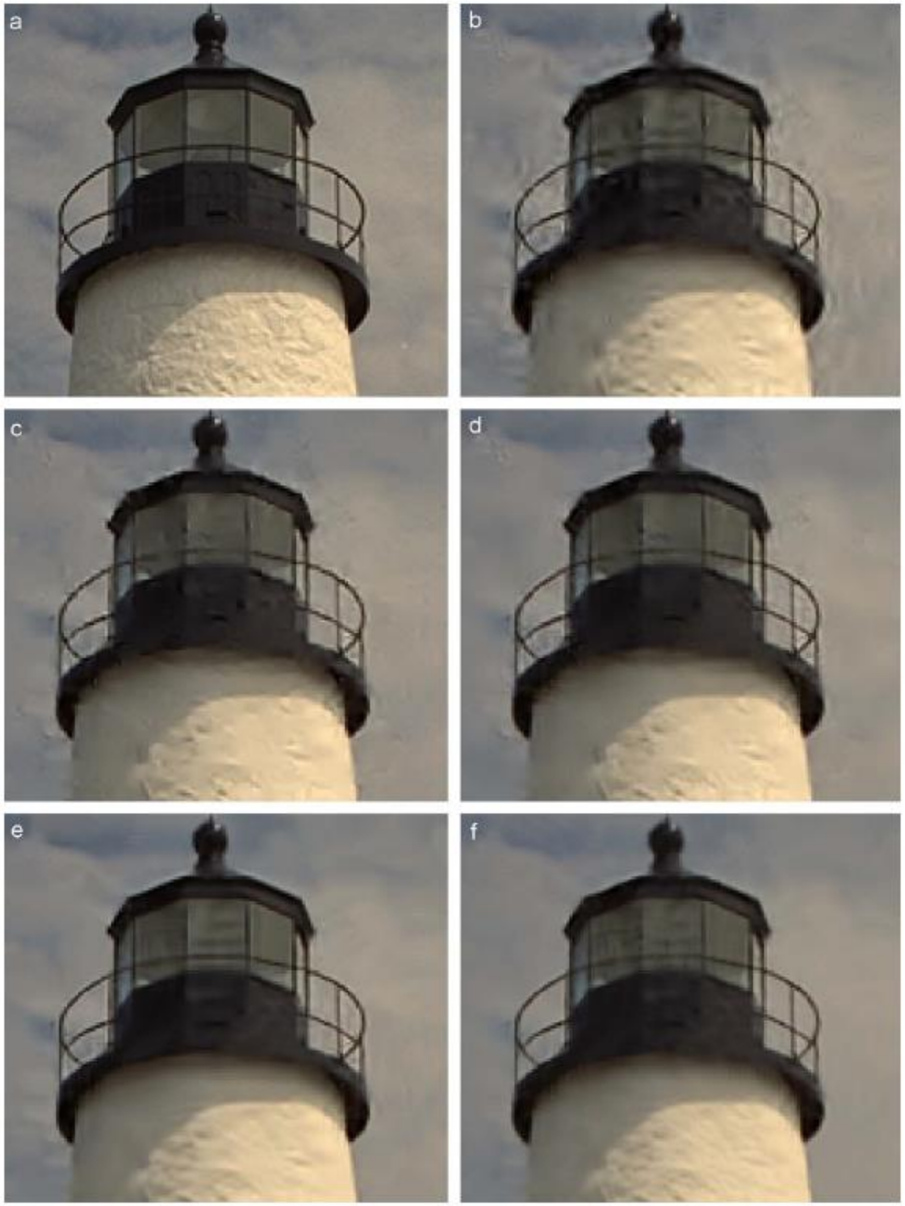
a
b
c
d
e
f

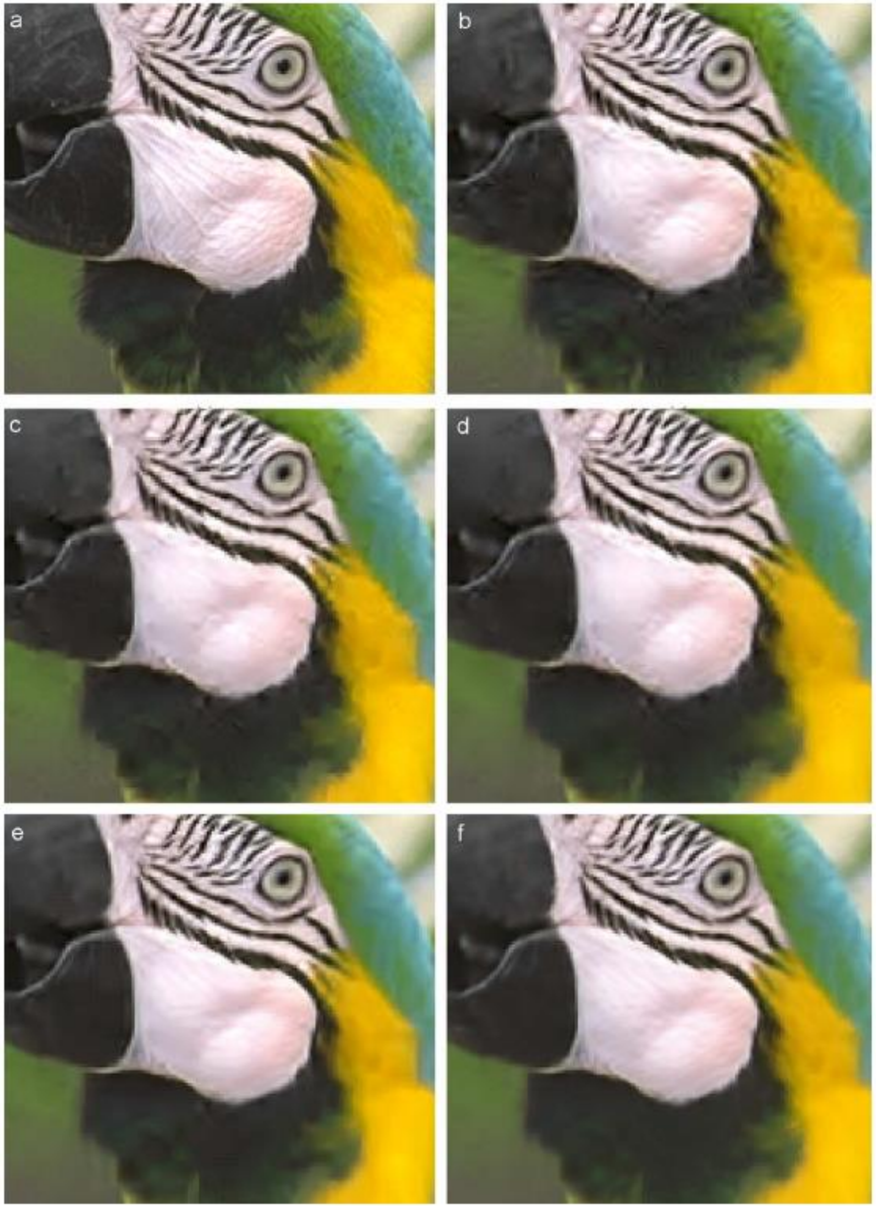
a
b
c
d
e
f

Q&A