Handling Blur

Image Restoration

Jan Flusser, Filip Sroubek, Barbara Zitova

Institute of Information Theory and Automation
Academy of Sciences of the Czech Republic
Prague

Contents

- Non-blind deconvolution
  - Wiener, Energy minimization, regularization
- Blind deconvolution
  - MAP, VB
- Multichannel (multiframe)
  - Super-resolution
- Space-variant blind deconvolution
  - Patch-wise, Parametric, Object motion, Conversion to space-invariant
- Hybrid methods
  - Fusion approach, High-speed cameras, Inertial sensors

Inverse Filter

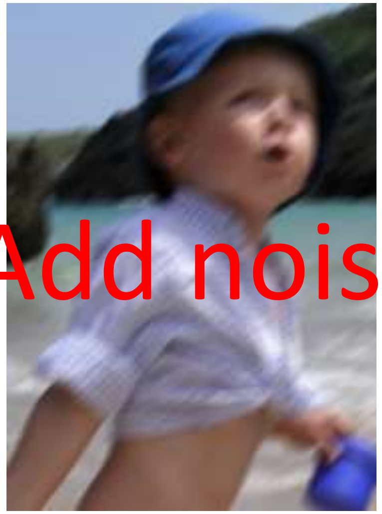

Add noise

$z(x)$

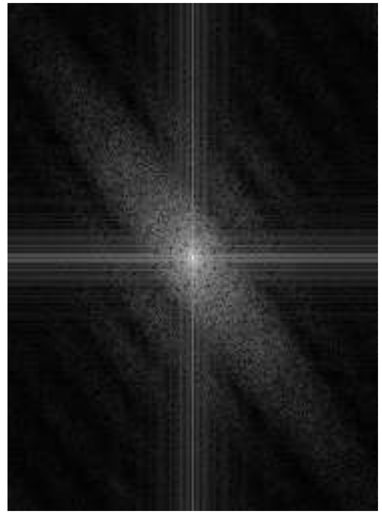

$Z(\omega)$

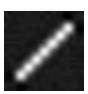

$h(x)$

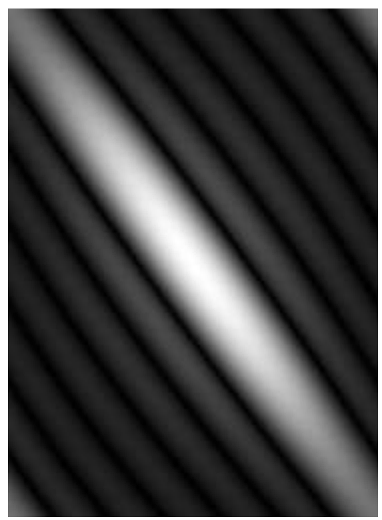

$H(\omega)$

$$\tilde{U}(\omega) = \frac{Z(\omega)}{H(\omega)}$$

equivalent

$$\tilde{u} = \arg\min_u \frac{1}{2}\|h * u - z\|^2$$

Wiener Filter

$z(x)$

Replace by a constant

$$\tilde{U} = \frac{H^*}{|H|^2 + \cancel{\frac{S_n}{S_u}}\lambda} Z$$

equivalent

$$\tilde{u} = \arg\min_u \frac{1}{2} \|h * u - z\|^2 + \lambda \|u\|^2$$

$\lambda = 0.001$

$\lambda = 0.1$

Deconvolution as energy minimization

$$\min_u E(u) = \min_u \underbrace{\boxed{\frac{1}{2}\|u * h - z\|^2}}_{\text{Data term}} + \underbrace{\boxed{\lambda Q(u)}}_{\text{Regularization}}$$

- Regularization: ill-posed problem --> well posed

  Enforce image smoothness
  - Wiener: $Q(u) = \int u(x)^2 dx$
  - Tikhonov: $\int |\nabla u(x)|^2$
  - Total Variation: $\int |\nabla u(x)|$
  - Non-convex L_p quasi-norm: $\int |\nabla u(x)|^p \,, \quad p < 1$

Regularization

$u(x)$

$z(x)$

Different regularization $Q(u)$

$|\nabla u|^2$

$|\nabla u|^1$

$|\nabla u|^0$

Estimated image $\tilde{u}(x)$

Regularization

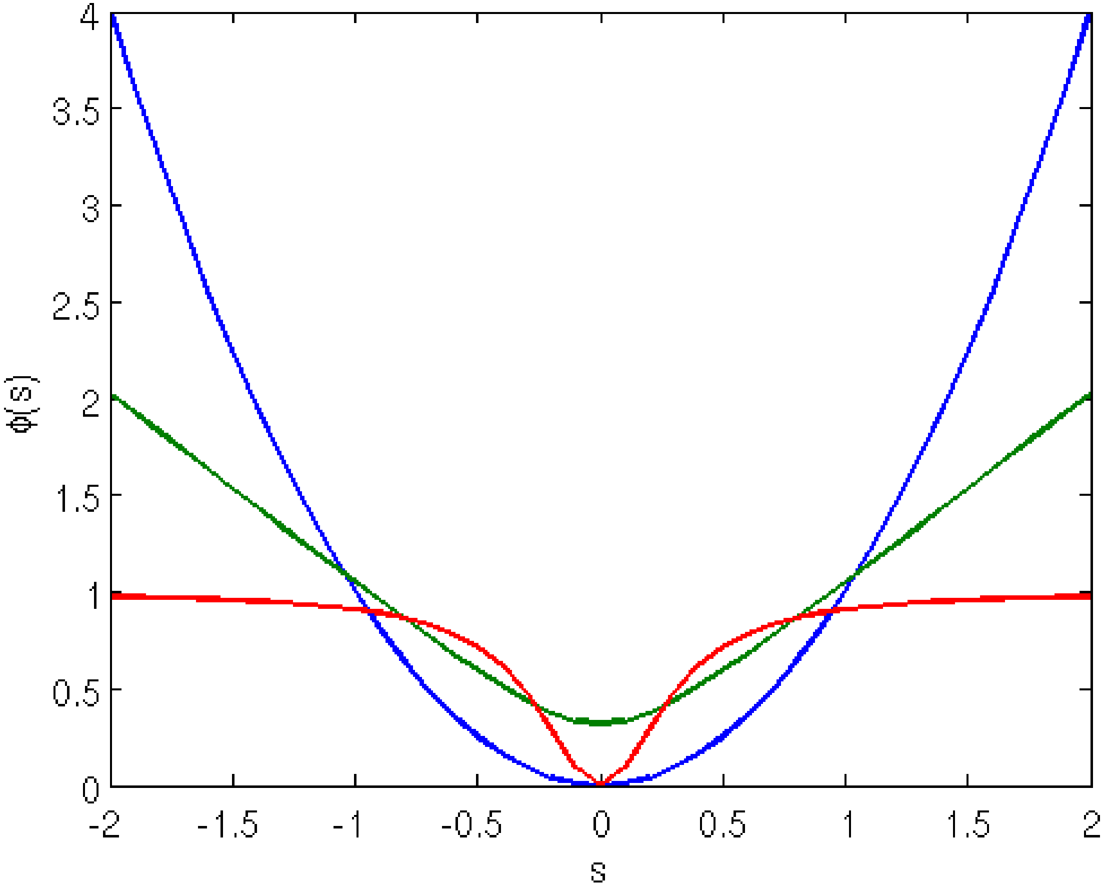

Tikhonov . . . $|s|^2$

TV . . . $|s| \approx \sqrt{|s|^2 + \epsilon}$

l_0-norm . . . $|s|_0 \approx \frac{|s|^2}{|s|^2 + \epsilon}$

Statistics of sharp images

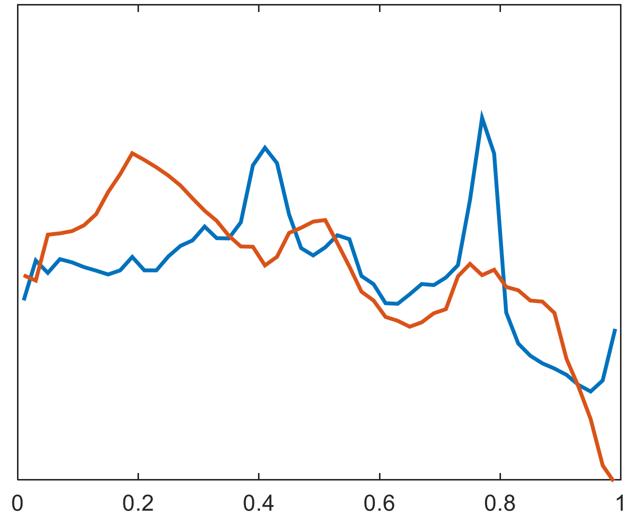

Intensity

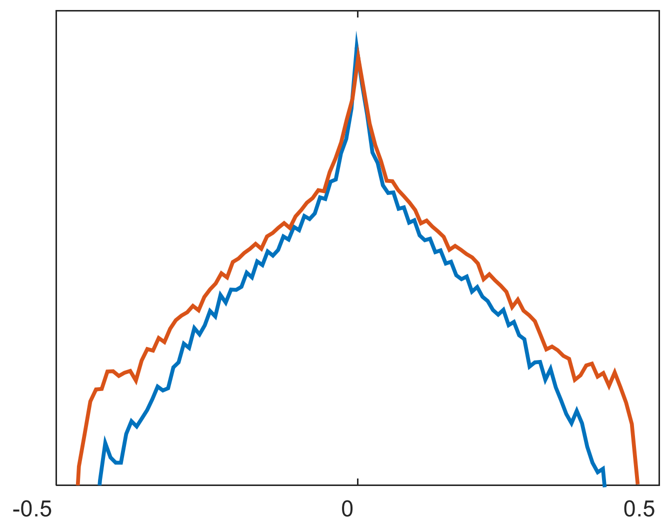

Gradient

Statistics of blurred images

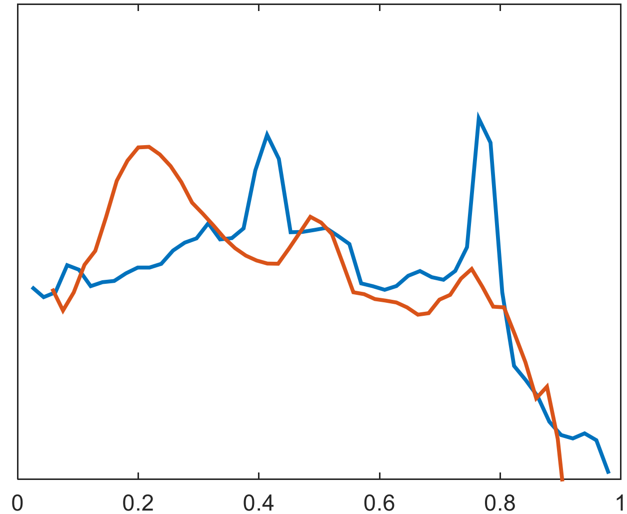

Intensity

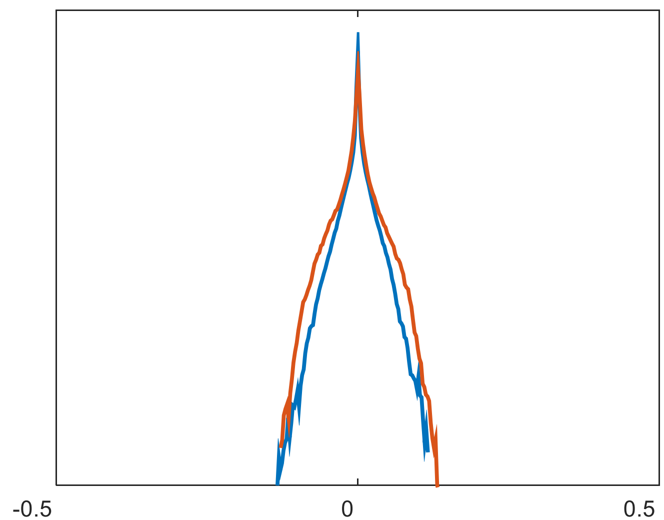

Gradient

How to tackle the blind case?

- When the blur kernel $h(x)$ is not known.

- One is tempted to:
  1) Add blur regularization
  2) Perform alternating minimization

Alternating Minimization

$$\min_{u,h} E(u,h) = \min_{u,h} \frac{1}{2}\|h * u - z\|^2 + \lambda Q(u) + \gamma R(h)$$

- Alternating Minimization

  1. *u-step:* $\tilde{u} = \arg\min_{u} E(u, \tilde{h})$

  2. *h-step:* $\tilde{h} = \arg\min_{h} E(\tilde{u}, h)$

  3. *repeat 1 and 2.*

No-blur solution

$$\min_{u,h} \frac{1}{2}\|u * h - z\|^2$$

- In theory any pair $(\tilde{u}, \tilde{h})$ is a solution → remember inverse filter
- common regularization → "NO-BLUR" solution

$$\tilde{u}(x) = z(x), \quad \tilde{h}(x) = \delta(x)$$

Regularization favors blur

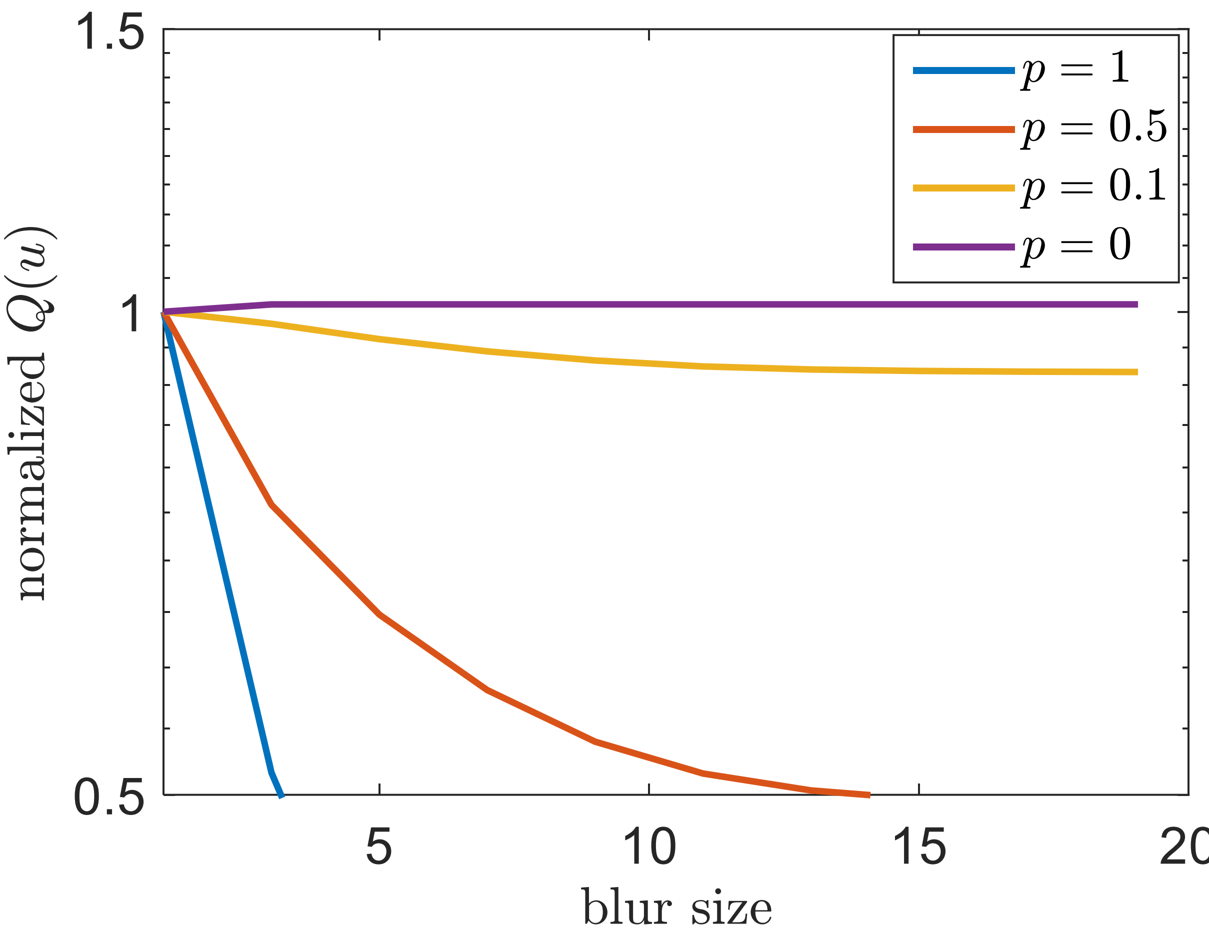

$$Q(u) = \int |\nabla u|^p$$

Regularization favors blur

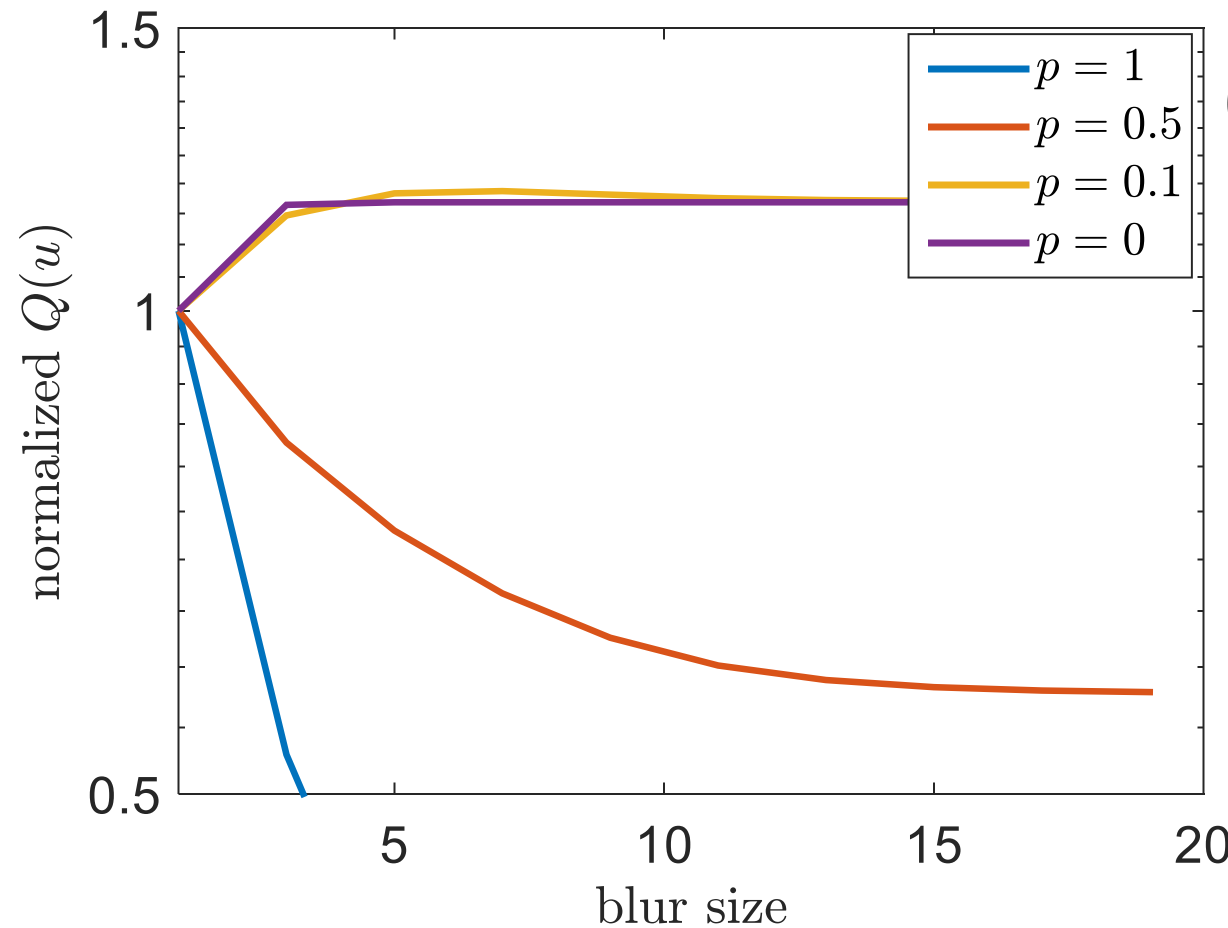

$$Q(u) = \int |\nabla u|^p$$

Artificially sparsify images

We need tricks!

- To avoid “no-blur” solution:
  - Artificial sharpening
  - Remove spikes
  - Adjusting priors on the fly
  - Hierarchical approach

Chan TIP1998
Shan SigGraph08
Cho SigGraph 09
Xu ECCV09, 13
Almeida TIP10
Krishnan 11
Zhong 13
Sun 13
Michael 14
Perrone 15
Pan 16

Artificial Sharpening

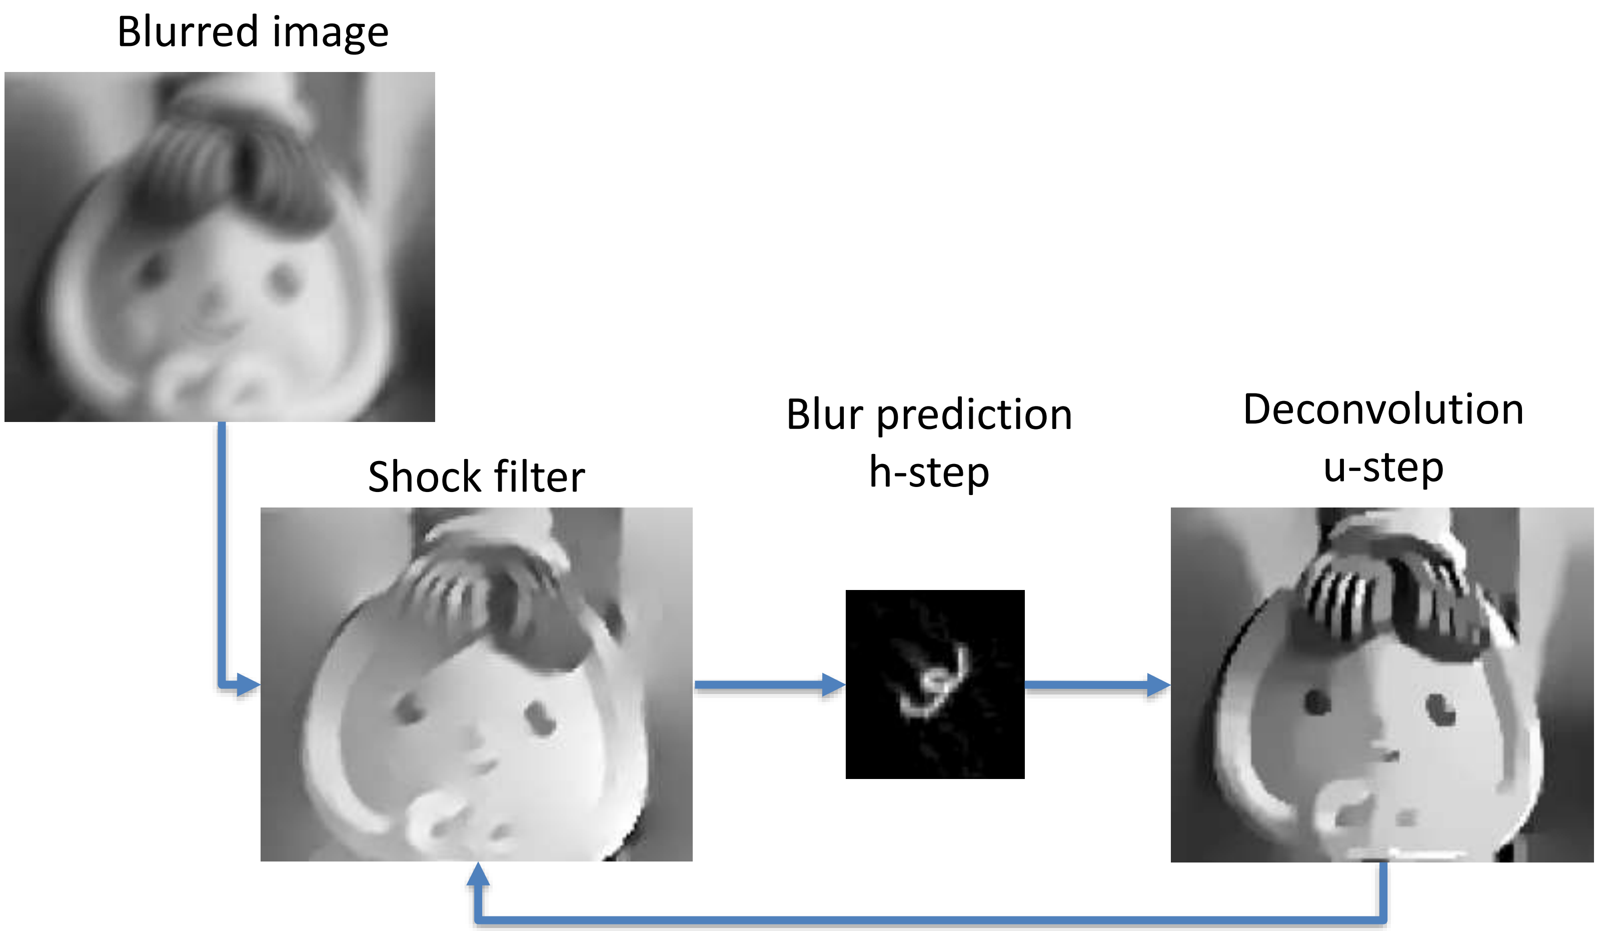

Cho et al., SIGGRAPH 2009

Remove Spiky Objects

Xu et al., ECCV 2010

Remove Spiky Objects

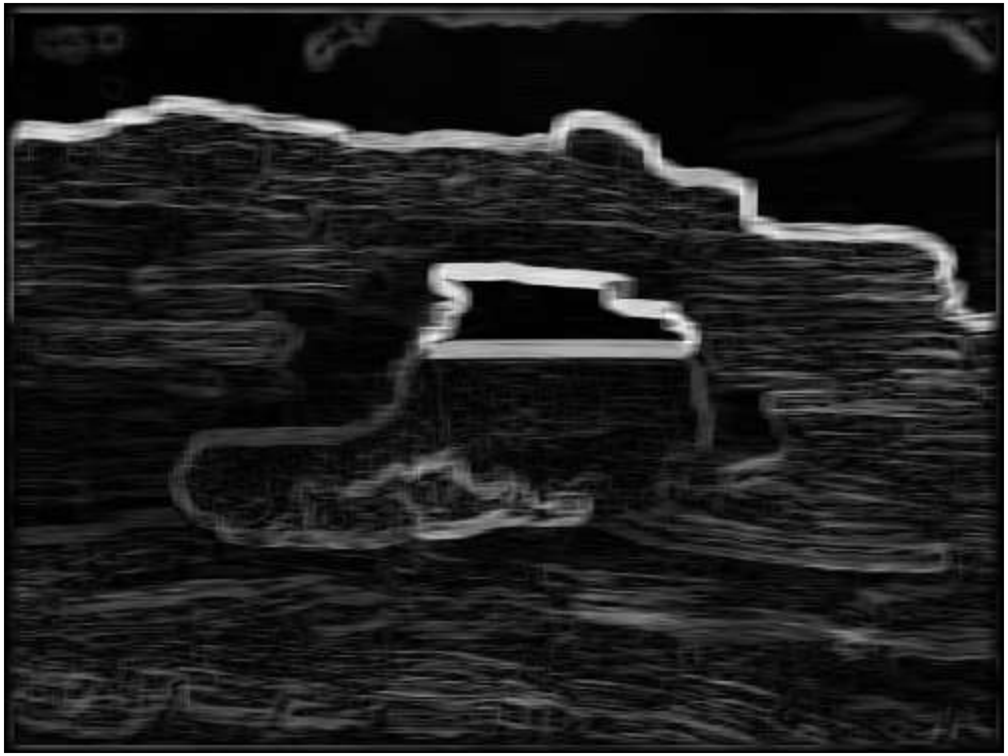

Mask out small objects

Xu et al., ECCV 2010

Remove Spiky Objects

Reconstructed image with
small objects removed

Xu et al., ECCV 2010

Adjusting priors

u $\quad Q(u) = \sum_i |\nabla u_i|^2 \quad \sum_i |\nabla u_i|^1 \quad \sum_i |\nabla u_i|^{0.5}$

- Start with an overestimated noise level and slowly decrease it to the correct level.
- Start with p<<1 and slowly increase it to p=1.

Hierarchical Deconvolution

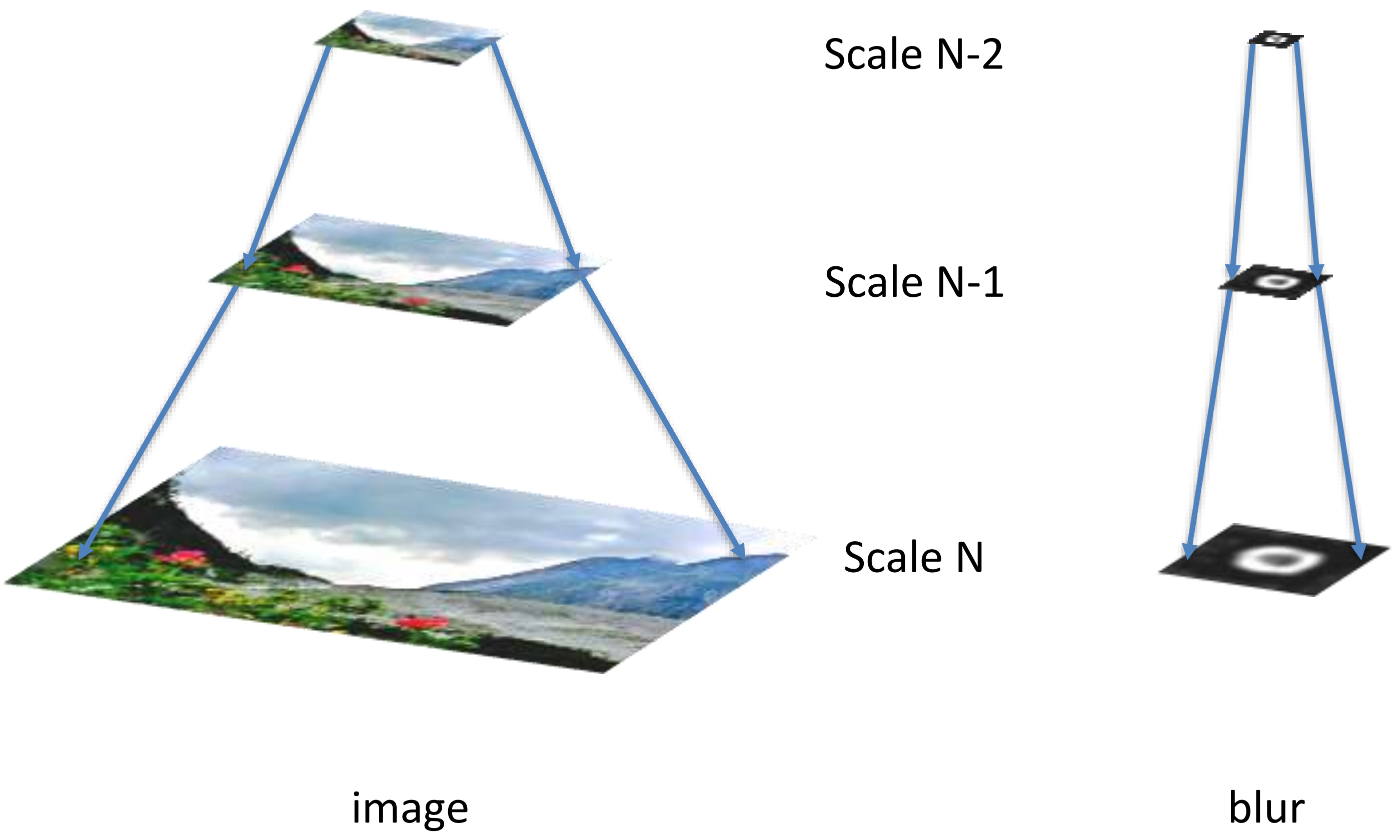

Bayesian Paradigm

$$p(u,h|z) = \frac{p(z|u,h)\,p(u)\,p(h)}{p(z)}$$

a posteriori **distribution**
unknown

likelihood
given by our problem

a priori **distribution**
our prior knowledge

- Maximum a posteriori (MAP): $\max p(u,h|z)$

Blind deconvolution with MAP

- max *a posteriori* probability $p(u,h|z)$

  ==> $\min -\log p(u,h|z)$

$$-\log p(u,h|z) \propto \boxed{-\log p(z|u,h)} \boxed{-\log p(u) - \log p(h)}$$

- Exponential family

$$E(u,h) = \frac{\lambda}{2}\|u * h - z\|^2 + Q(u) + R(h)$$

NOTHING NEW!

Bayesian Paradigm revisited

- Marginalize the posterior

$$p(h|z) = \int p(u, h|z) du$$

- Maximize the marginalized prob.

$$\hat{h} = \arg\max_h p(h|z)$$

- and then maximize the posterior

$$\hat{u} = \arg\max_u p(u, \hat{h}|z)$$

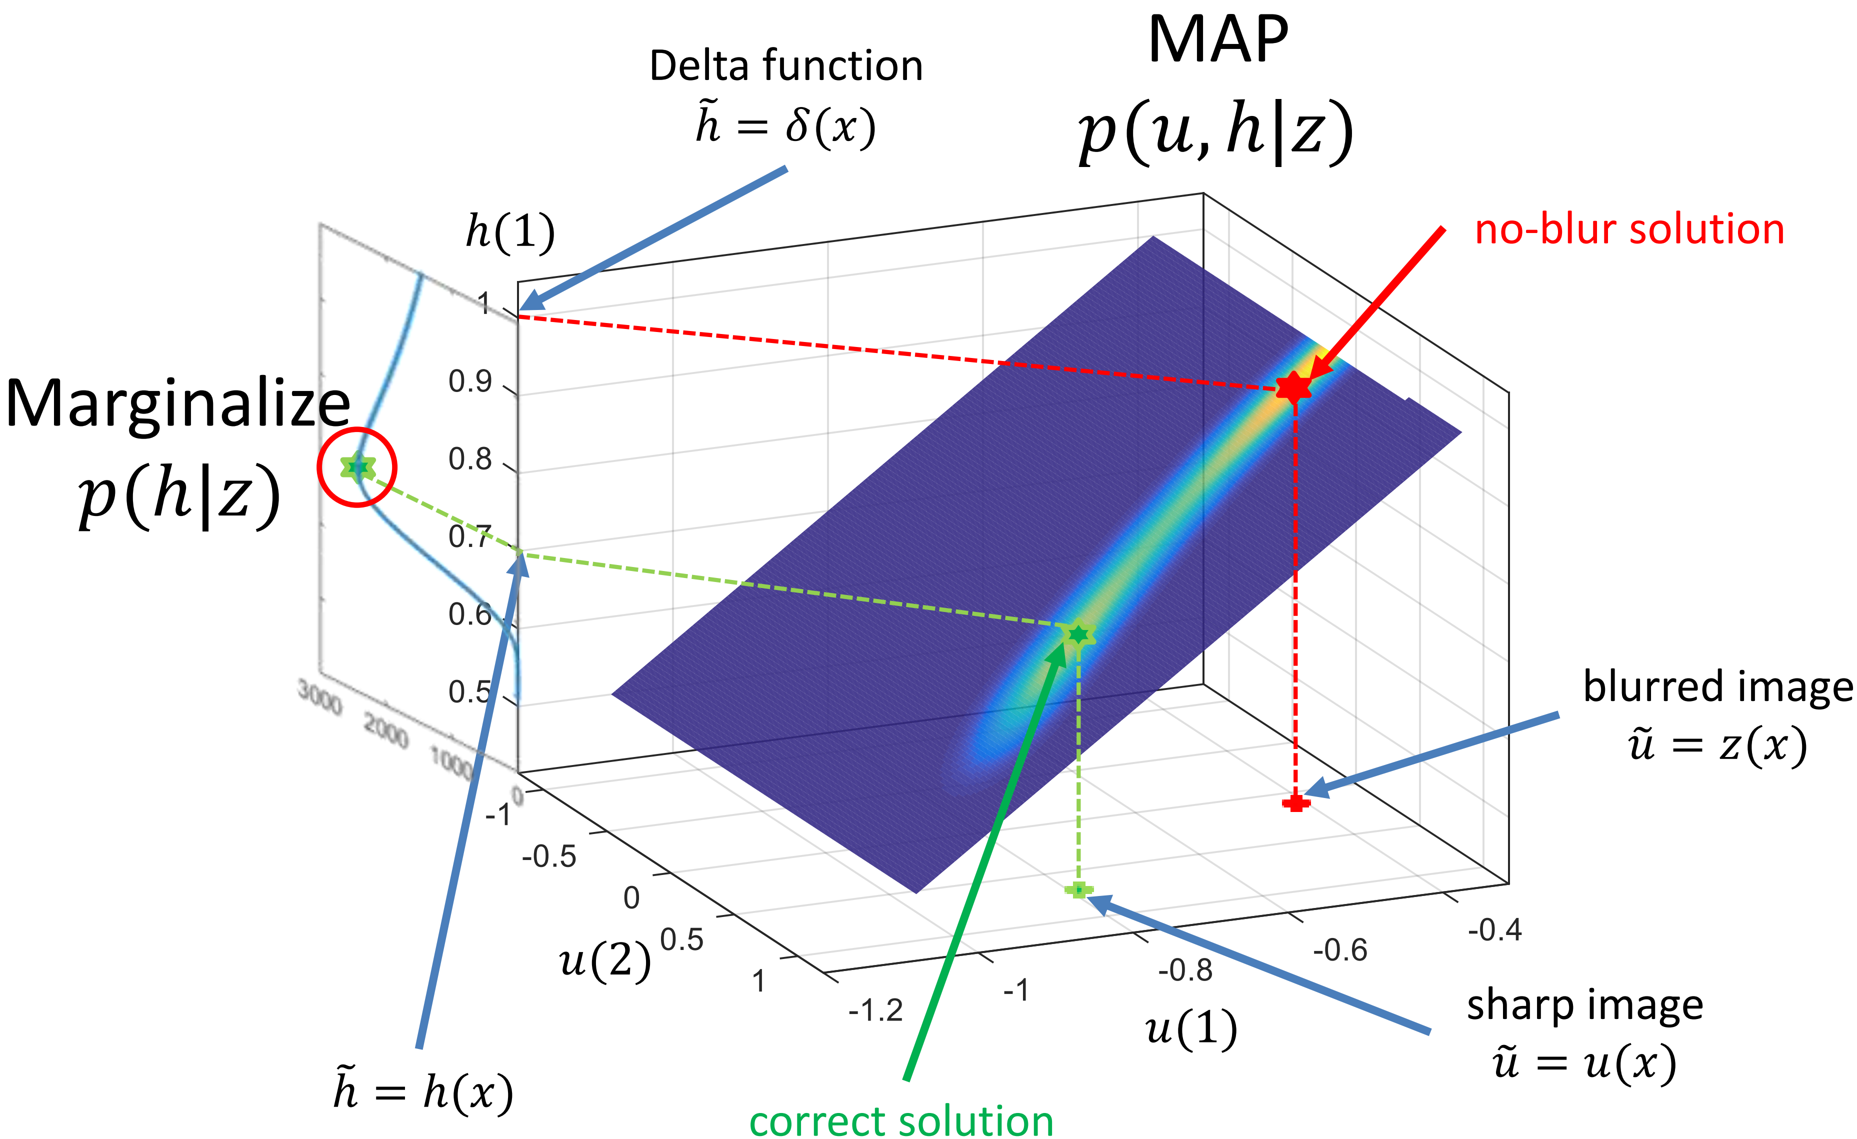
MAP
$p(u,h|z)$
Delta function
$\tilde{h} = \delta(x)$
no-blur solution
Marginalize
$p(h|z)$
$h(1)$
1
0.9
0.8
0.7
0.6
0.5
3000
2000
1000
0
-1
-0.5
0
0.5
1
$u(2)$
-1.2
-1
-0.8
-0.6
-0.4
$u(1)$
blurred image
$\tilde{u} = z(x)$
sharp image
$\tilde{u} = u(x)$
$\tilde{h} = h(x)$
correct solution

How to marginalize?

$$p(h|z) = \int p(u,h|z)du$$

- If Gaussian distributions → analytic solution exists in the form of Gaussian distribution
- If not (our case) → approximation
  - Laplace approximation
  - Factorization with Variational Bayes

Variational Bayes

- Factorization of the posterior

$$p(u, h|z) \approx q(u)q(h)$$

and then marginalization is trivial.

Miskin 01
Fergus 06
Whyte 10
Levin 11
Babacan 12
Wipf 14

- Every factor q depends on moments of other variables => must be solved iteratively.

Example of blind deconvolution

Blurred image
$z(x)$

Reconstructed image
$\tilde{u}(x)$

Multichannel Model

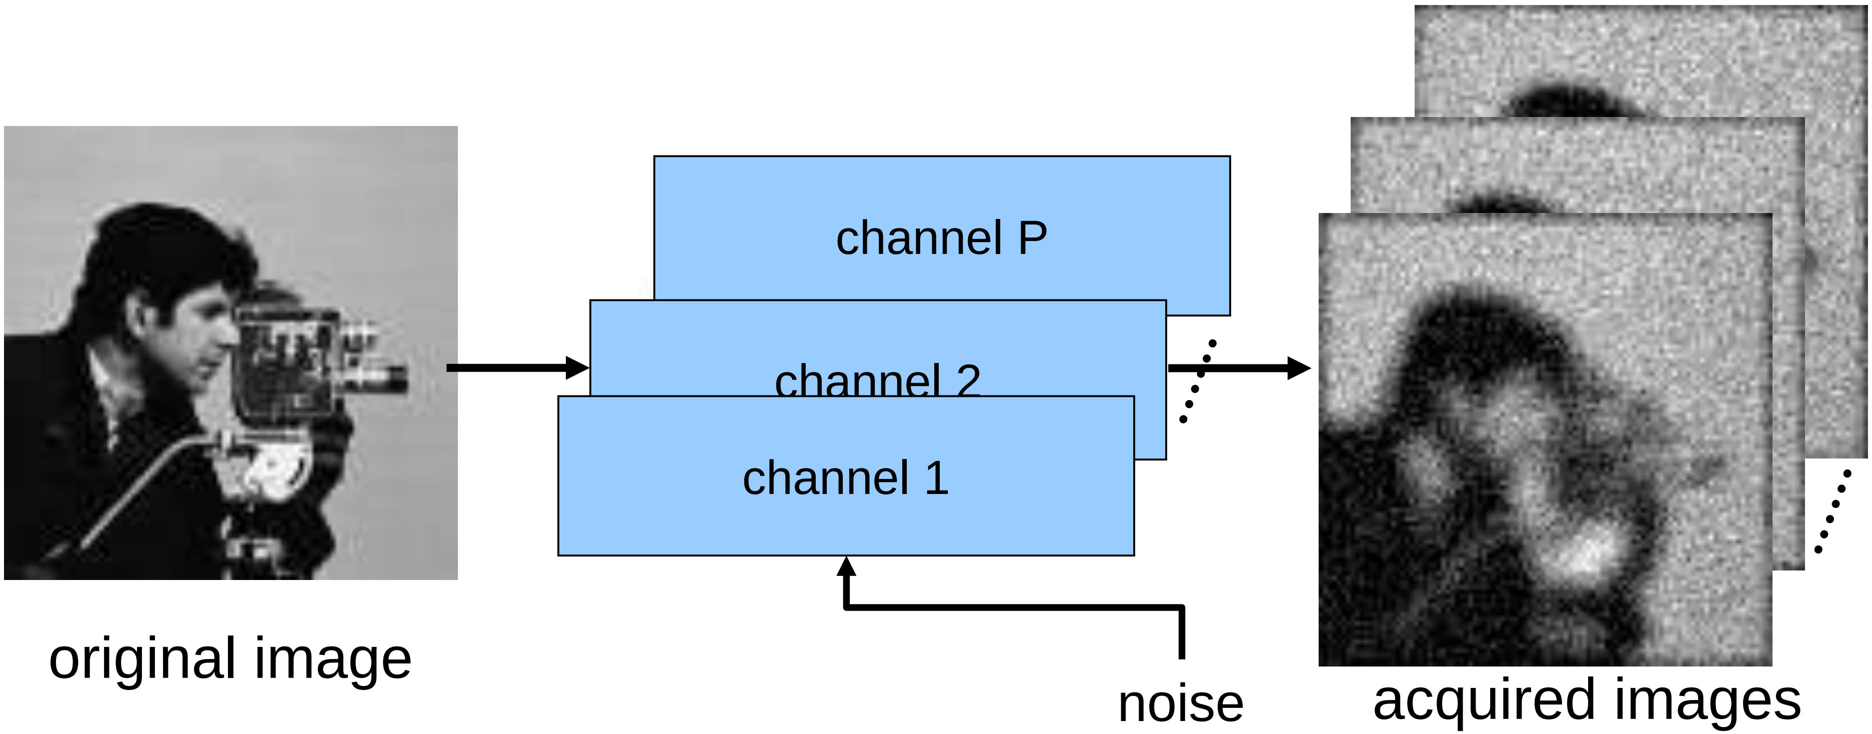

$$[u * h_p] + n_p = z_p$$

Image Restoration

- Acquisition model:

$$z_1 = (h_1 * u) + n_1$$
$$\vdots$$
$$z_P = (h_P * u) + n_P$$

- Optimization problem

$$E(u, \{h_p\}) = \boxed{\frac{1}{2} \sum_p \|(h_p * u) - z_p\|^2} + \boxed{\lambda Q(u)} + \boxed{\gamma R(\{h_p\})}$$

Data term

Image regularization term

Blur Regularization term

Gaussian noise L2 norm

$$Q(u) = \int \phi(|\nabla u(x)|)$$

Multichannel Blur Regularization

$$u$$

$$z_1 = u * h_1 \qquad\qquad h_2 * u = z_2$$

$$\boxed{z_1 * \tilde{h}_2} = u * h_1 * \tilde{h}_2 \qquad - \qquad \tilde{h}_1 * h_2 * u = \boxed{\tilde{h}_1 * z_2}$$

$$=$$

$$\mathbf{0}$$

$$R(\{h_p\}) = \frac{1}{2} \sum_{1 \le p,q \le P} \| z_p * h_q - z_q * h_p \|^2$$

Out-of-focus Camera

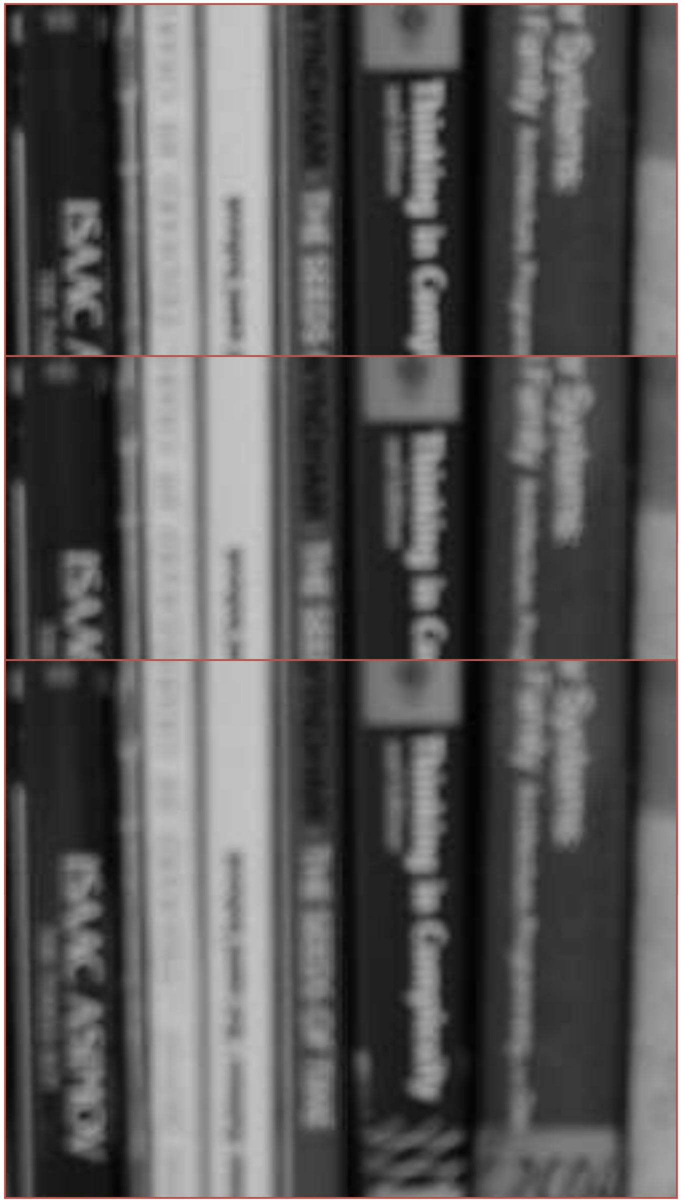

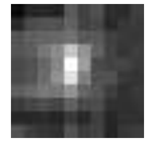

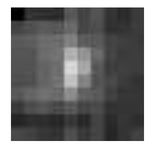

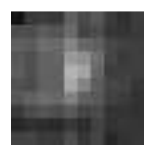

reconstructed

in focus

Camera motion

Astronomical images

Degraded images

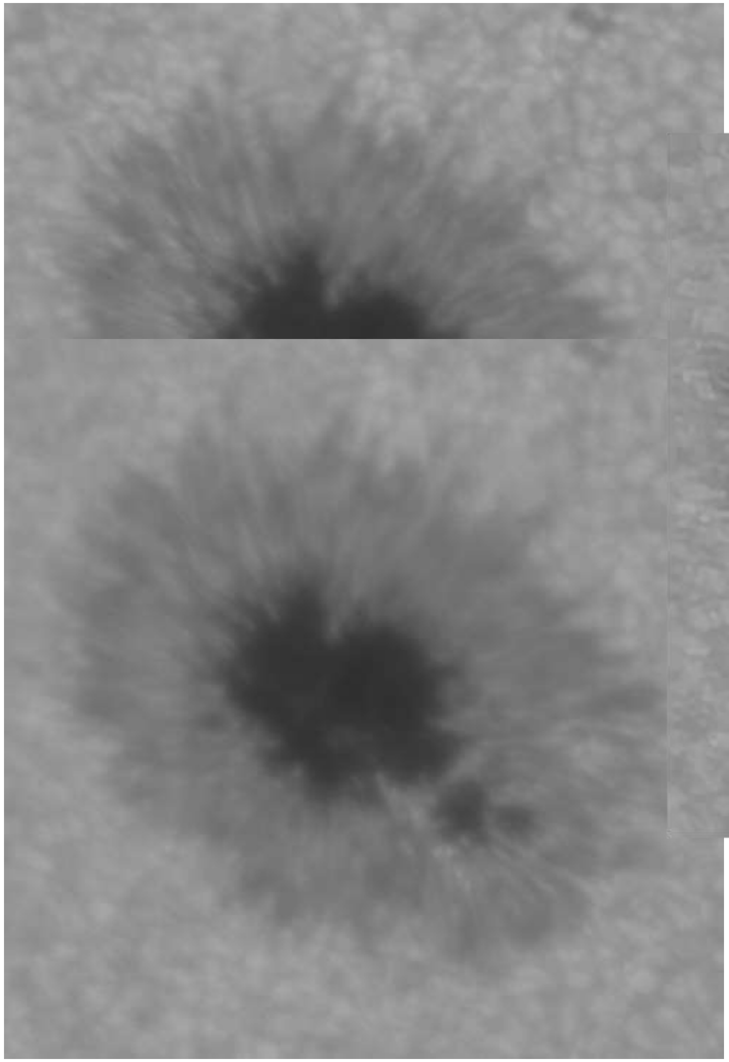

Reconstructed image

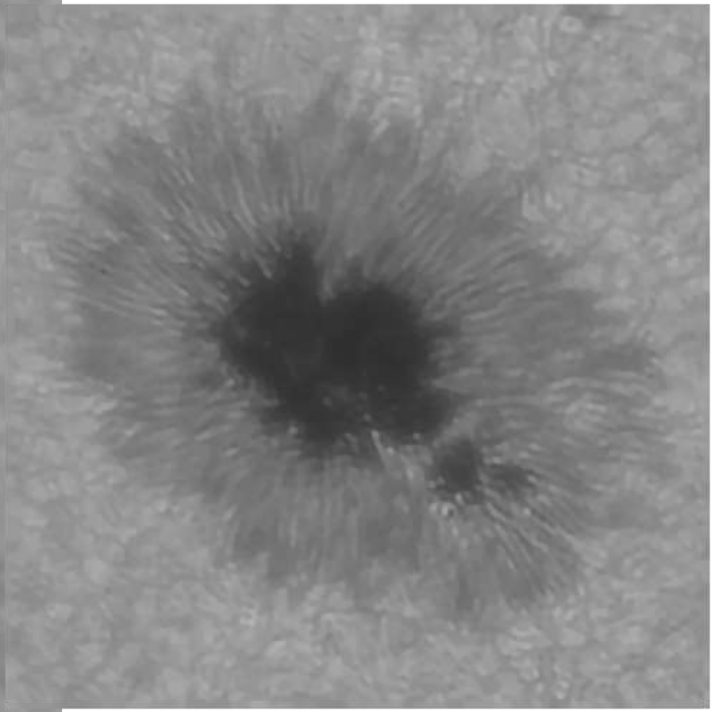

PSF estimation

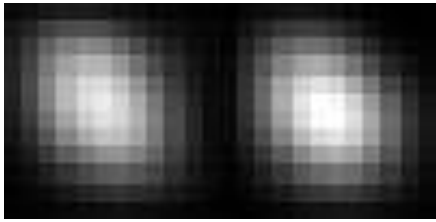

Misregistration of the channels

   

... leads to artefacts if not handled properly

Robustness to misalignment

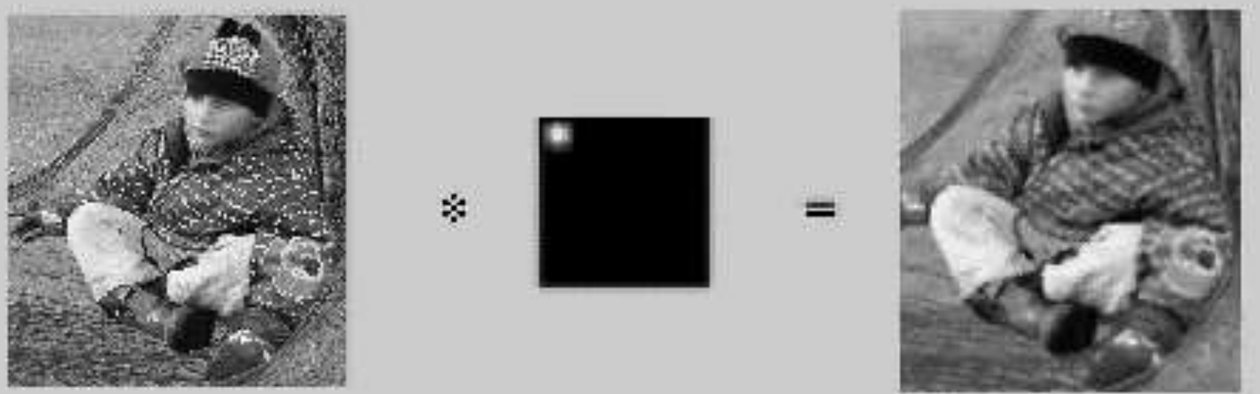

original image PSF degraded image

$$(u * h_p)[\tau_p(x)] + n_p(x) = z_p(x)$$

$$(u * \hat{h}_p)(x) + n_p(x) = z_p(x)$$

Slight misalignment

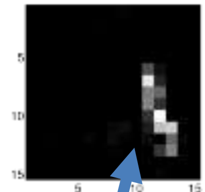

Misalignment compensated

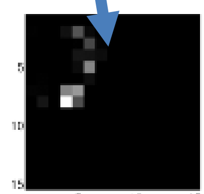

No compensation

With compensation

Image acquisition with aliasing

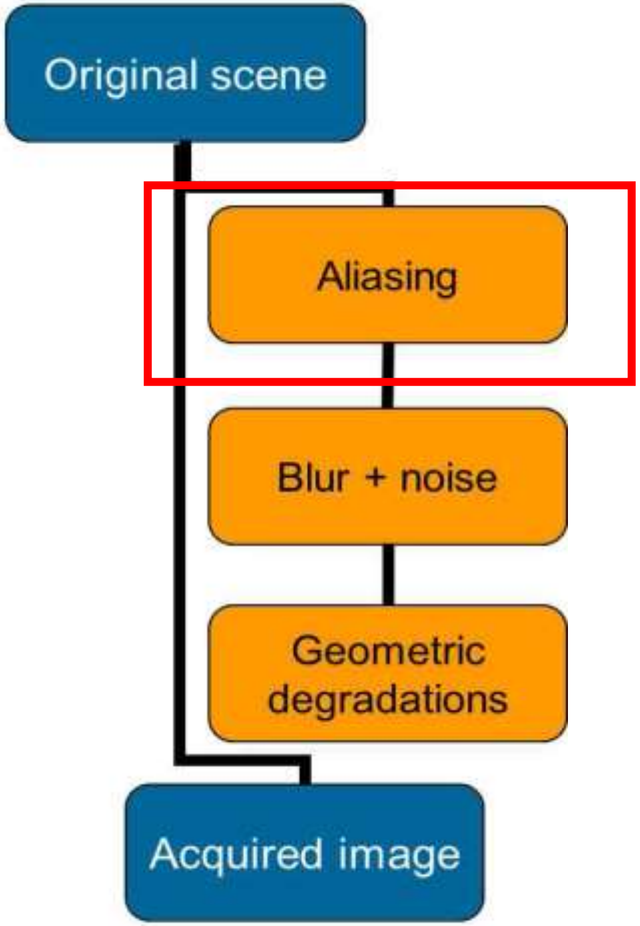

Aliasing

- The loss of the high frequencies (details)

due to an insufficient resolution of the camera

To suppress aliasing, apply multiple acquisitions with sub-pixel shifts

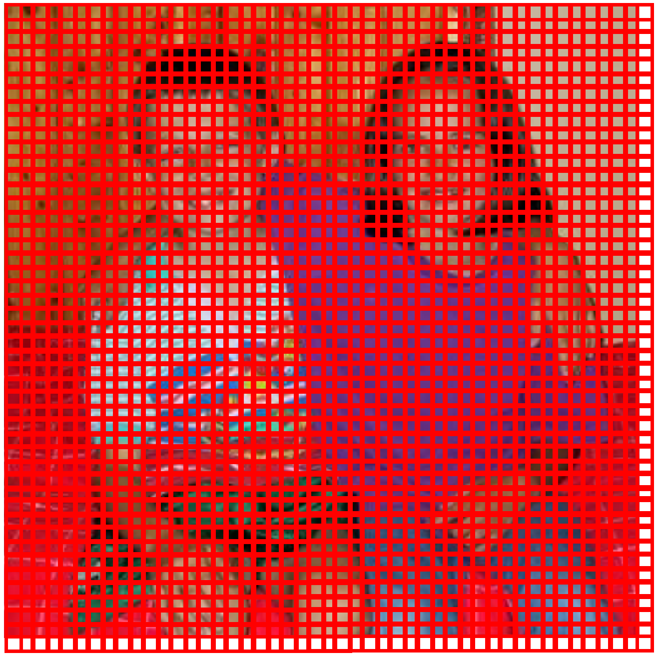

The image resolution has increased

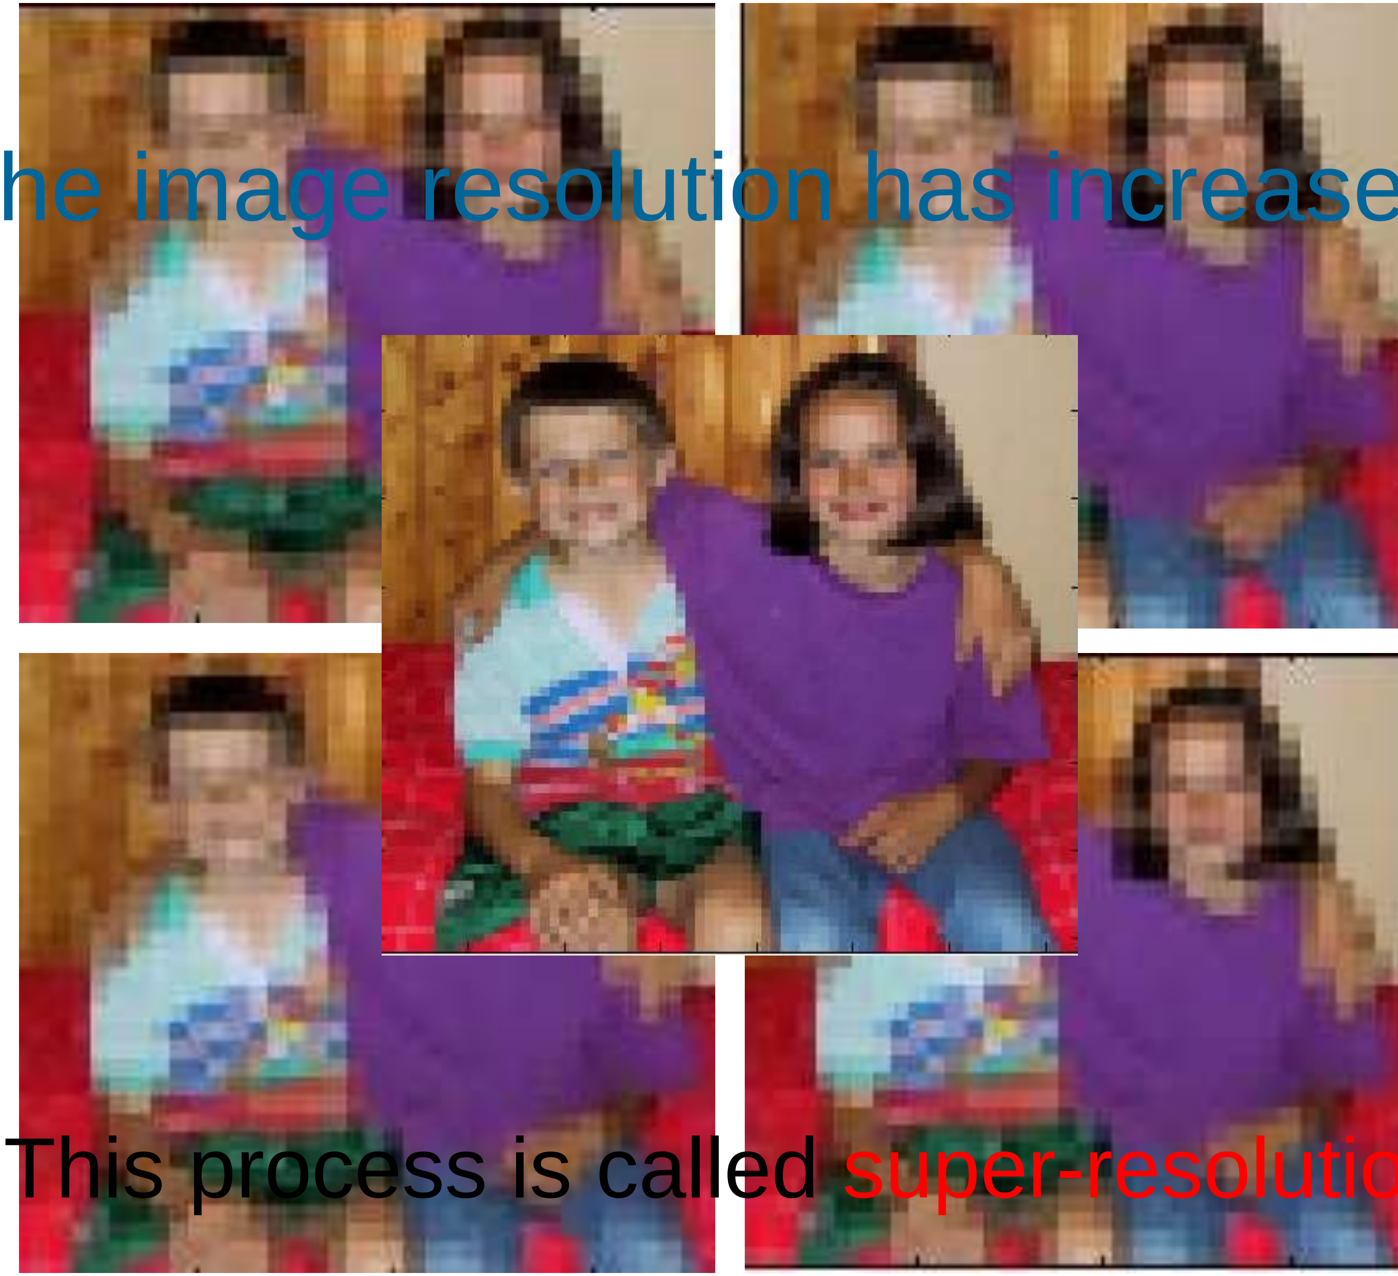

This process is called super-resolution

Realistic superresolution

SR must include also de-blurring

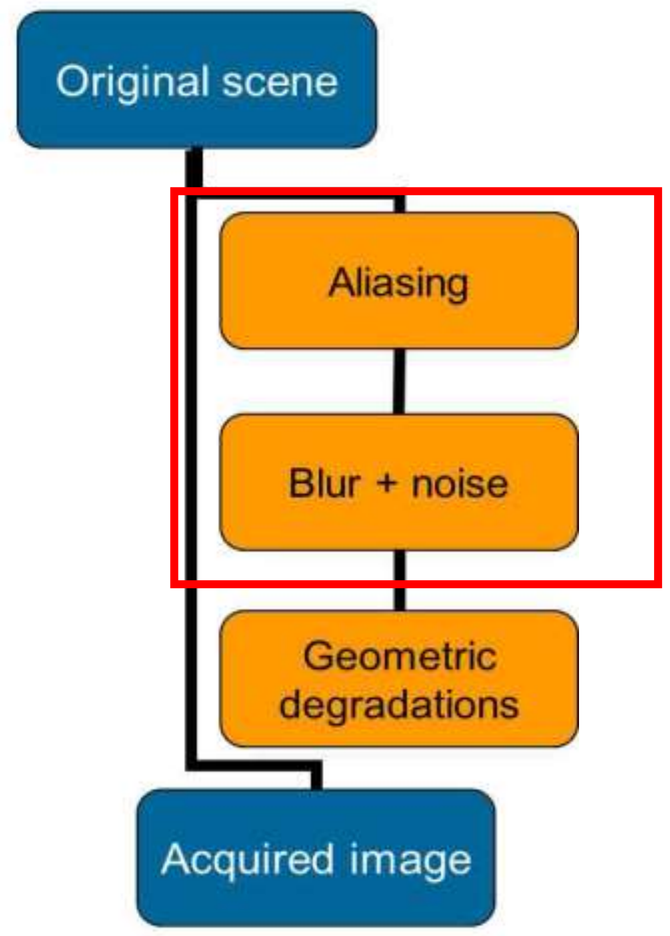

MC Model with Decimation

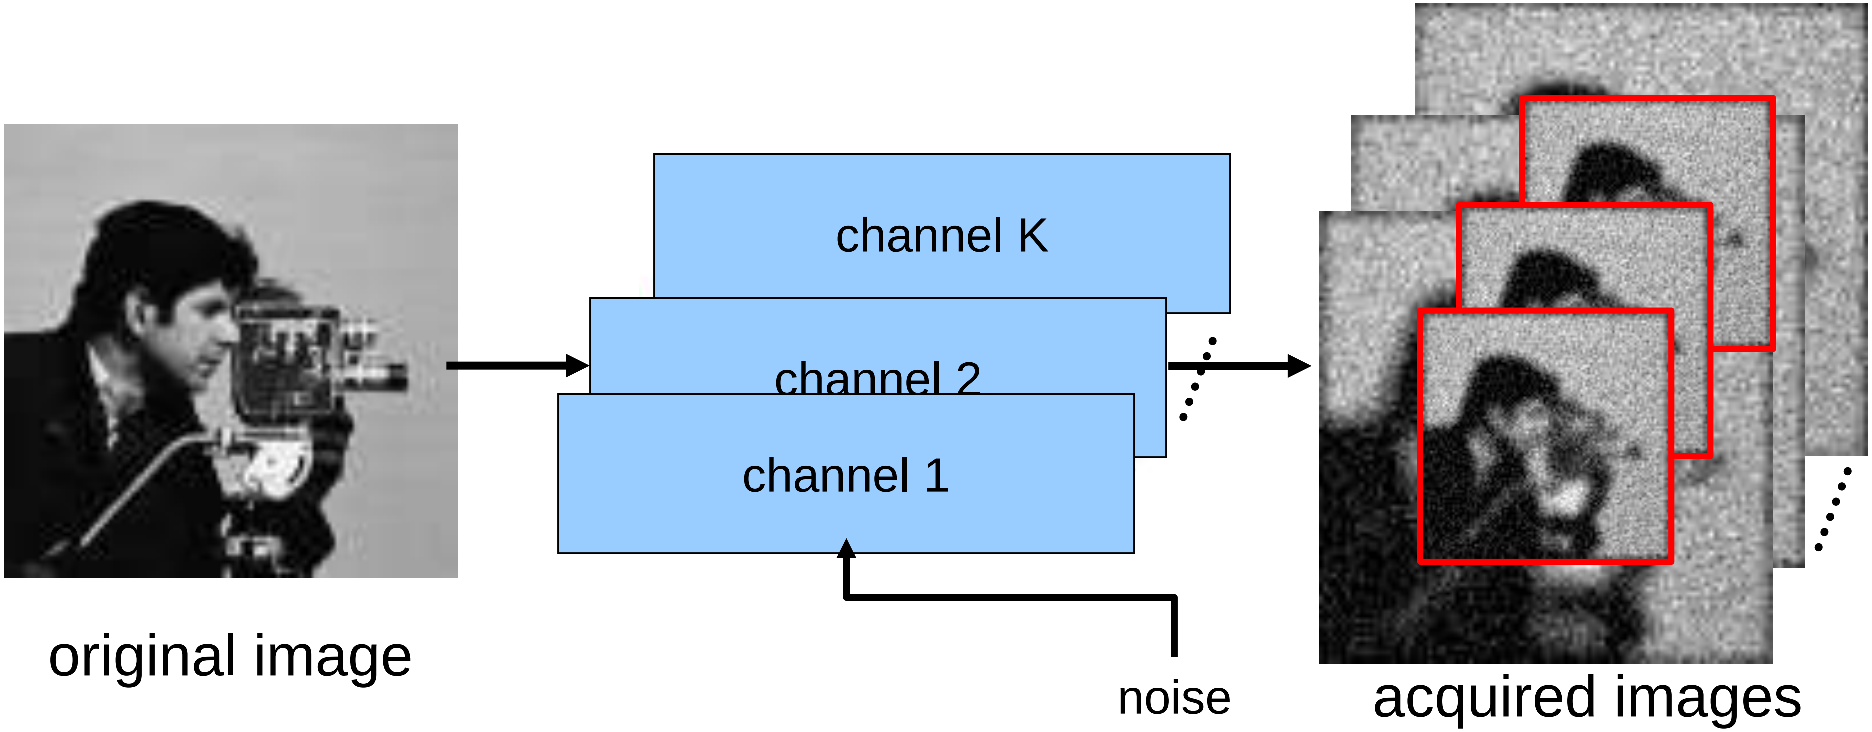

$$D[u * h_k] + n_k = z_k$$

$$\min_{u,h} E(u,h) = \min_{u,h} \frac{1}{2} \sum_k \| D(h_k * u) - z_k \|^2 + \lambda Q(u) + \gamma R(h)$$

Superresolution

rough registration

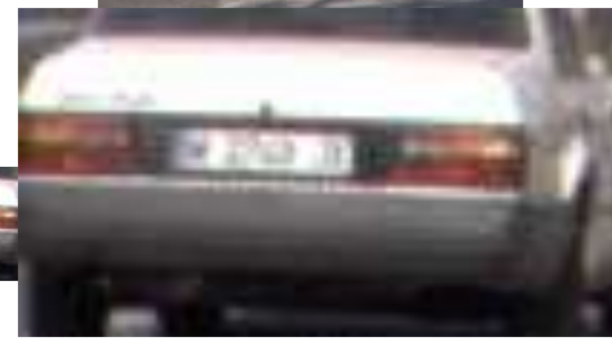

Superresolved image (2x)

Optical zoom (ground truth)

8 images

original

SR 2x

SR 3x

Video Restoration

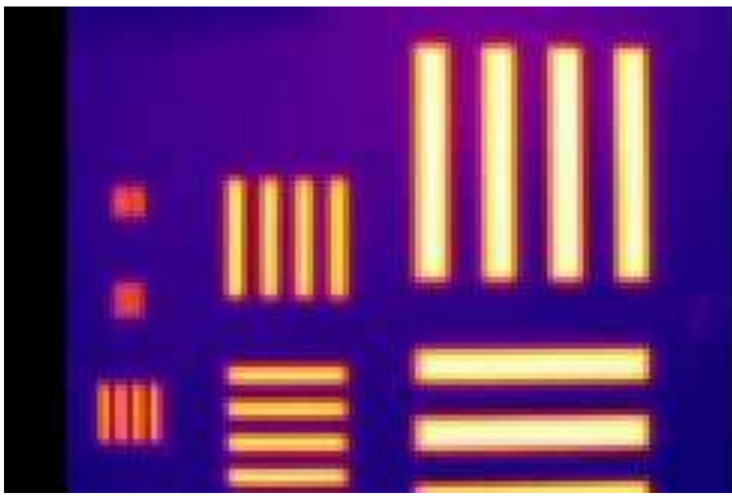

Input video

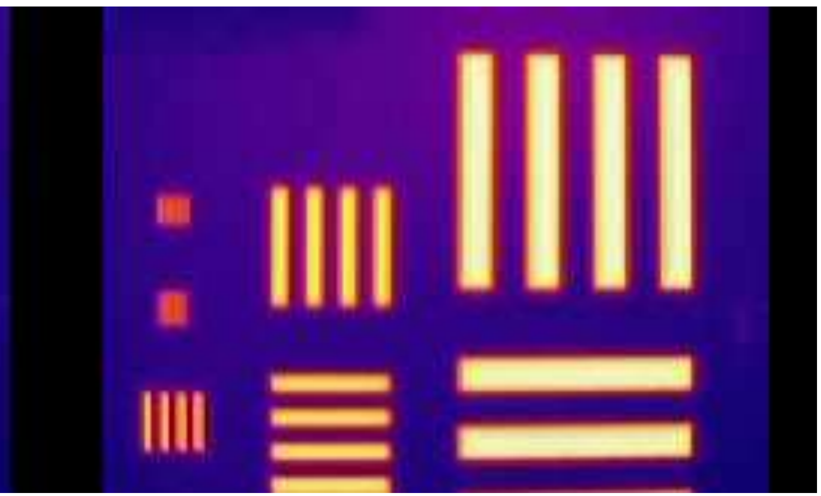

Super-resolution

Space-variant Blurs

Camera Motion

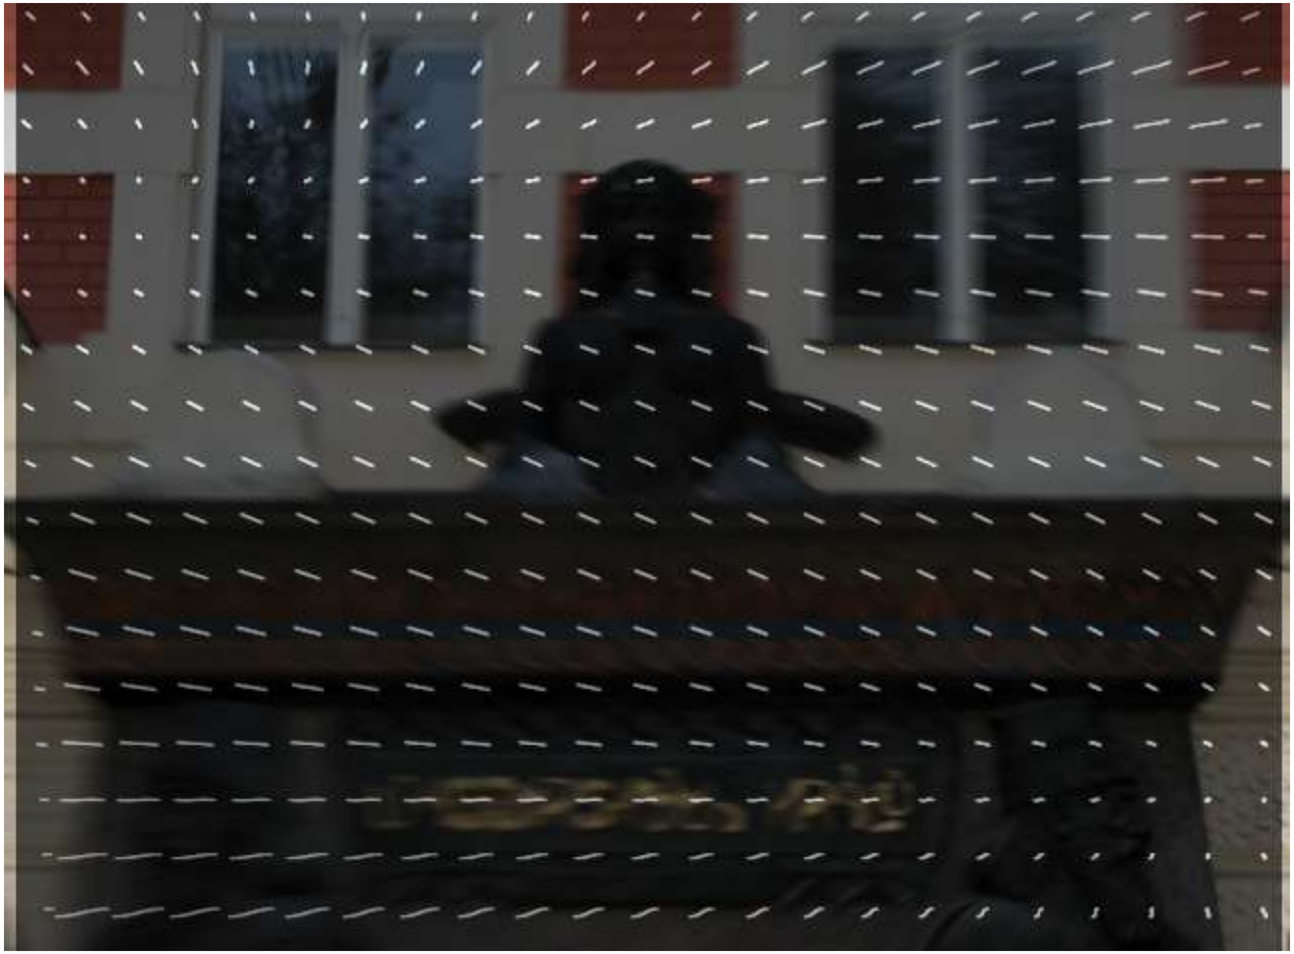

Object Motion

Optical aberrations

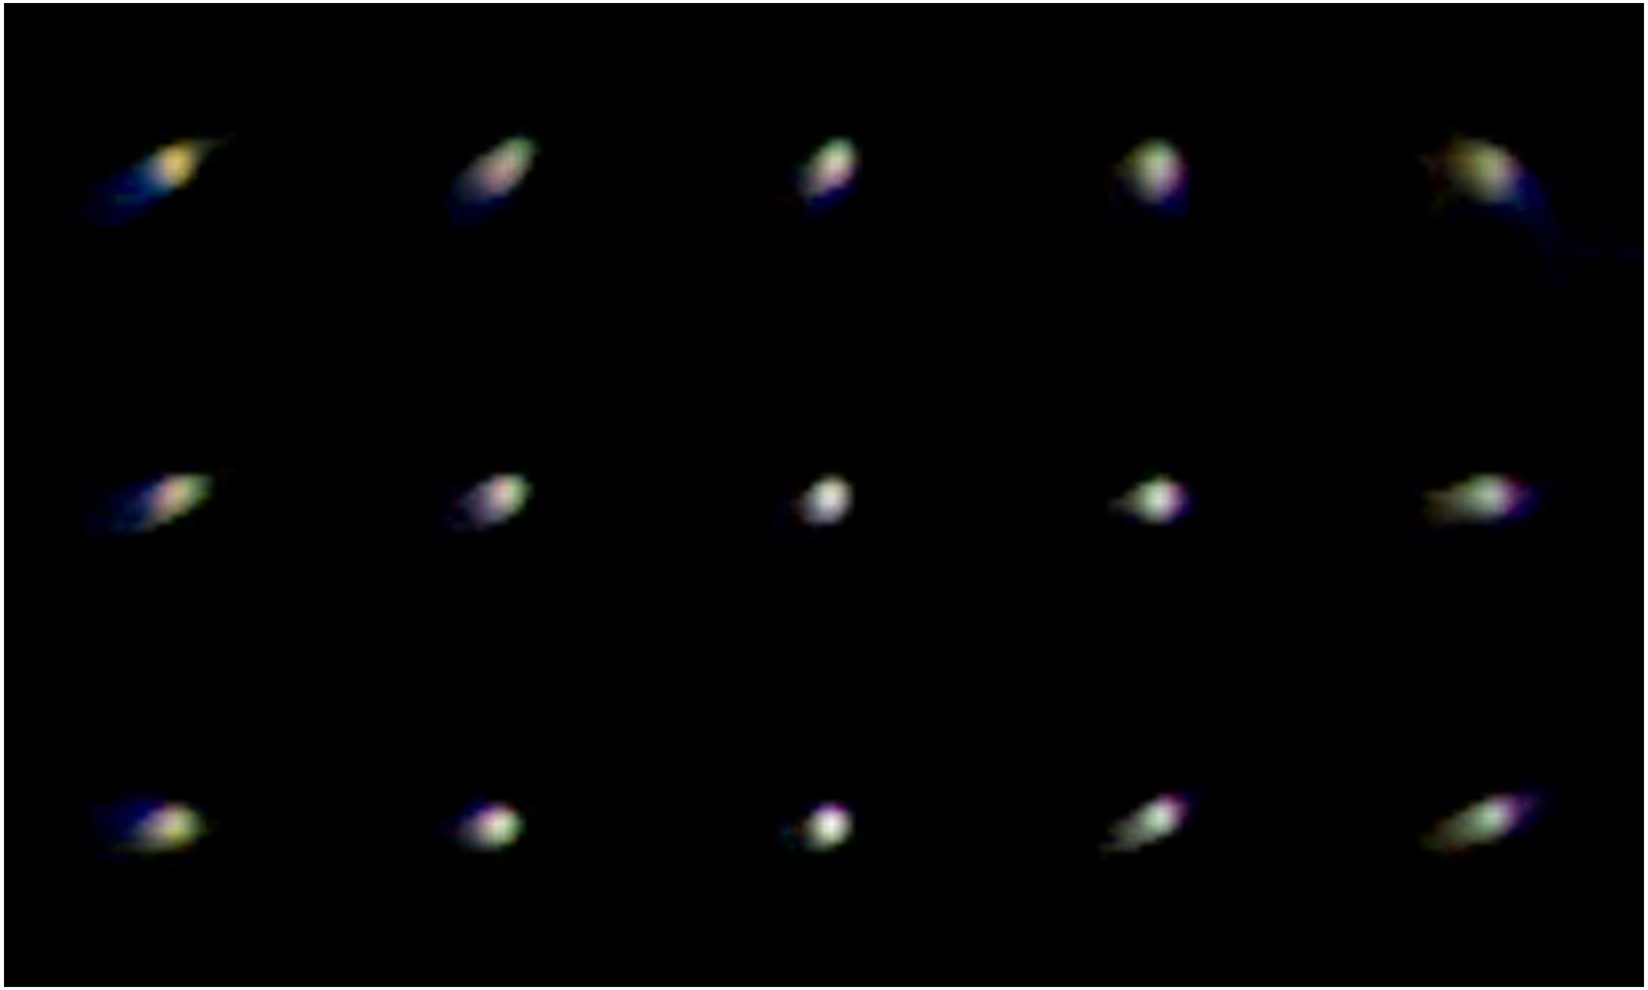

Space-variant Out-of-focus Blur

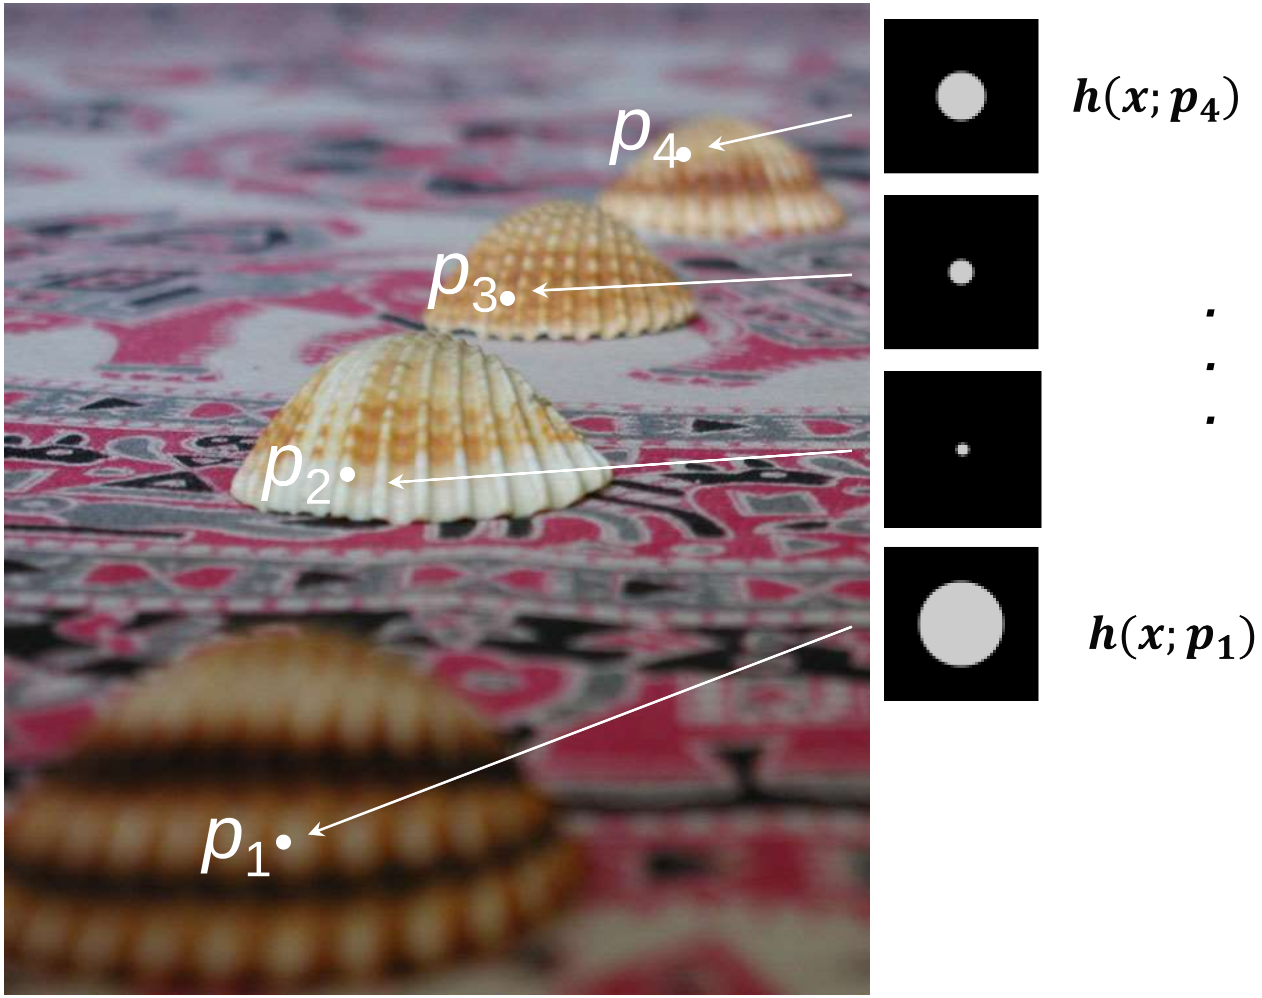

Approximation of SV Blur

- Patch-wise convolution
  - General SV PSF
  - Locally convolution may not hold

Joshi CVPR08, Sorel ICIP09,
Ji CVPR12,Sun CVPR15

- Parametric model (Blur Basis)
  - More accurate
  - Model may not hold

Whyte CVPR10, Gupta ECCV10,
Hirsch ICCV11, Zhang NIPS13

- Local object motion
  - Segmentation
  - Line blur

Levin NIPS06, Shan ICCV07, Dai CVPR08,
Chakrabarti CVPR10, Kim CVPR14

- Conversion to space-invariant
  - HW design

Levin SIGGRAPH08

Patch-wise restoration

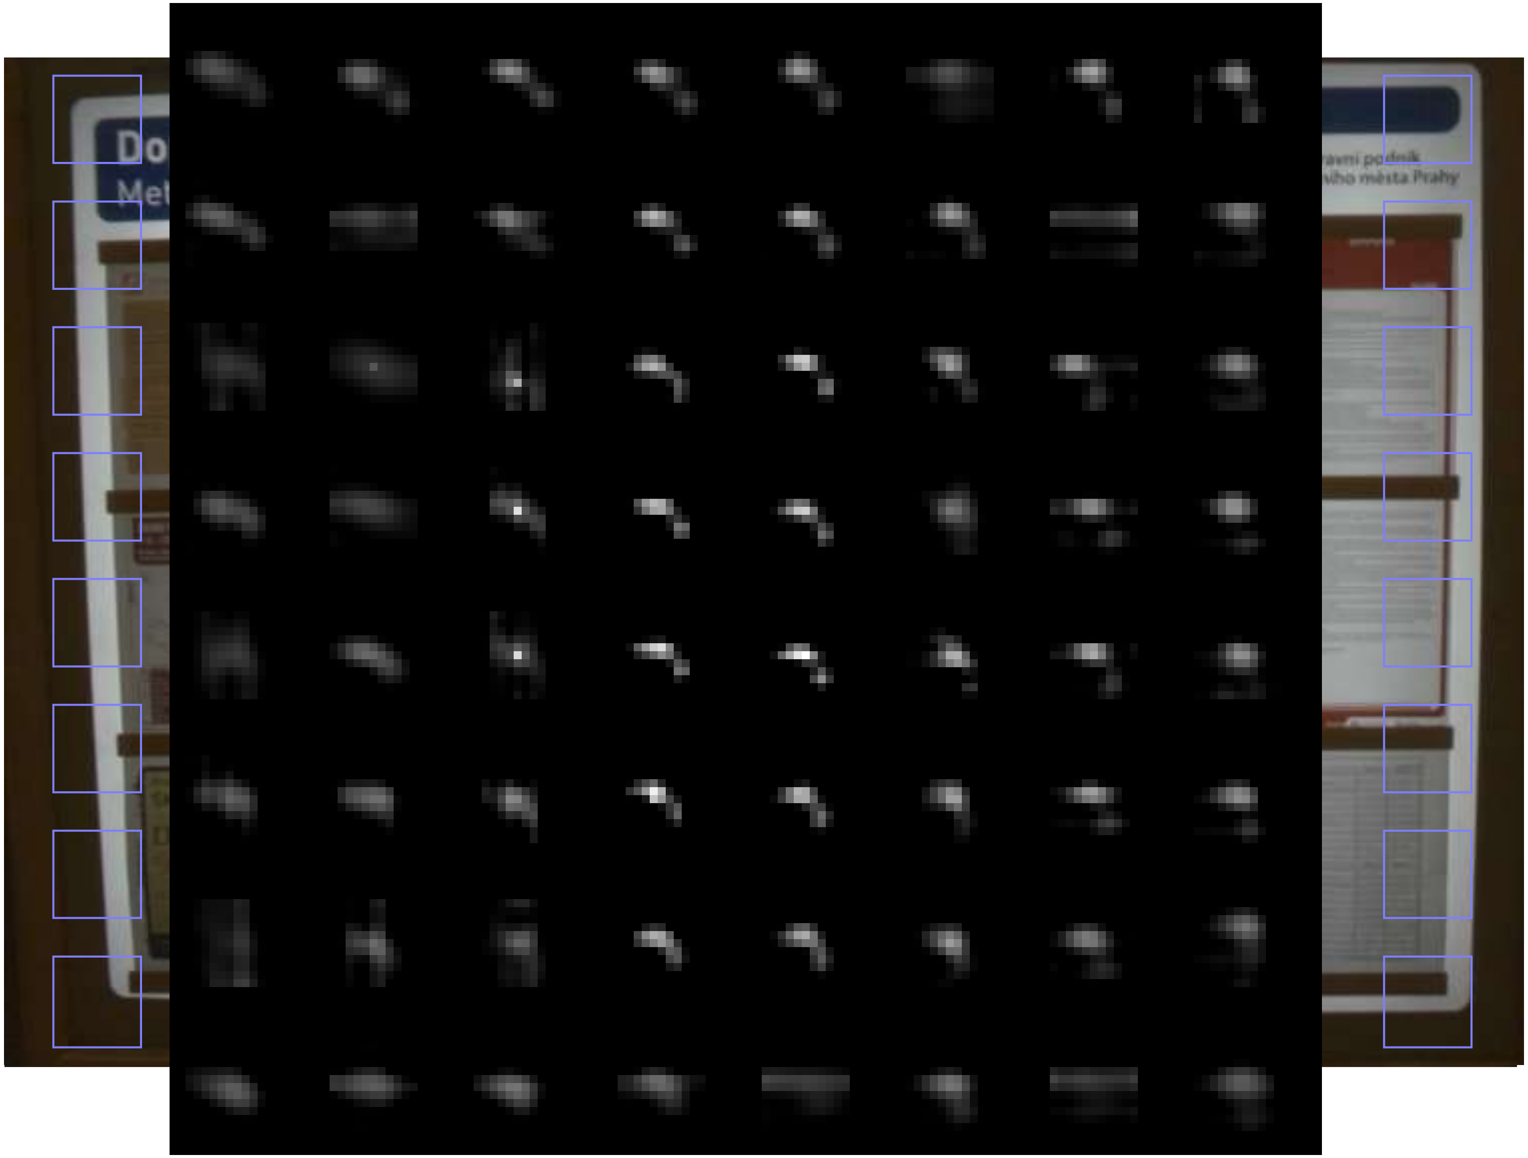

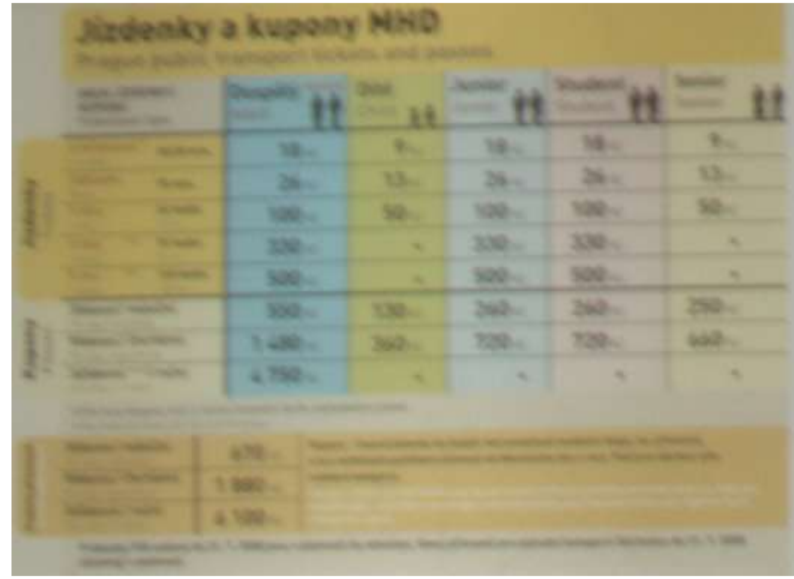
Jízdenky a kupony MHD

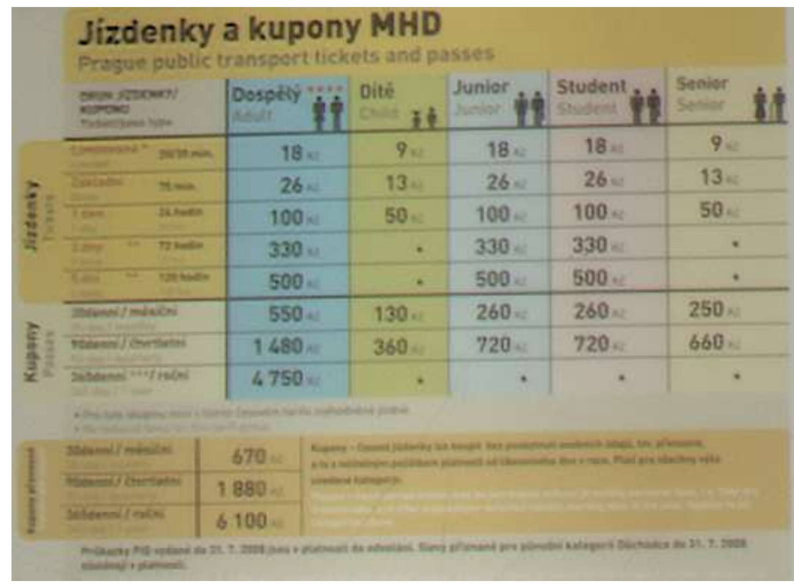
Jízdenky a kupony MHD
Prague public transport tickets and passes
Dospělý
Dítě
Junior
Student
Senior
Jízdenky
Kupony
18 9 18 18 9
26 13 26 26 13
100 50 100 100 50
330 330 330
500 500 500
550 130 260 260 250
1 480 360 720 720 660
4 750
670
1 880
6 100

Parametric Model (Blur Basis)

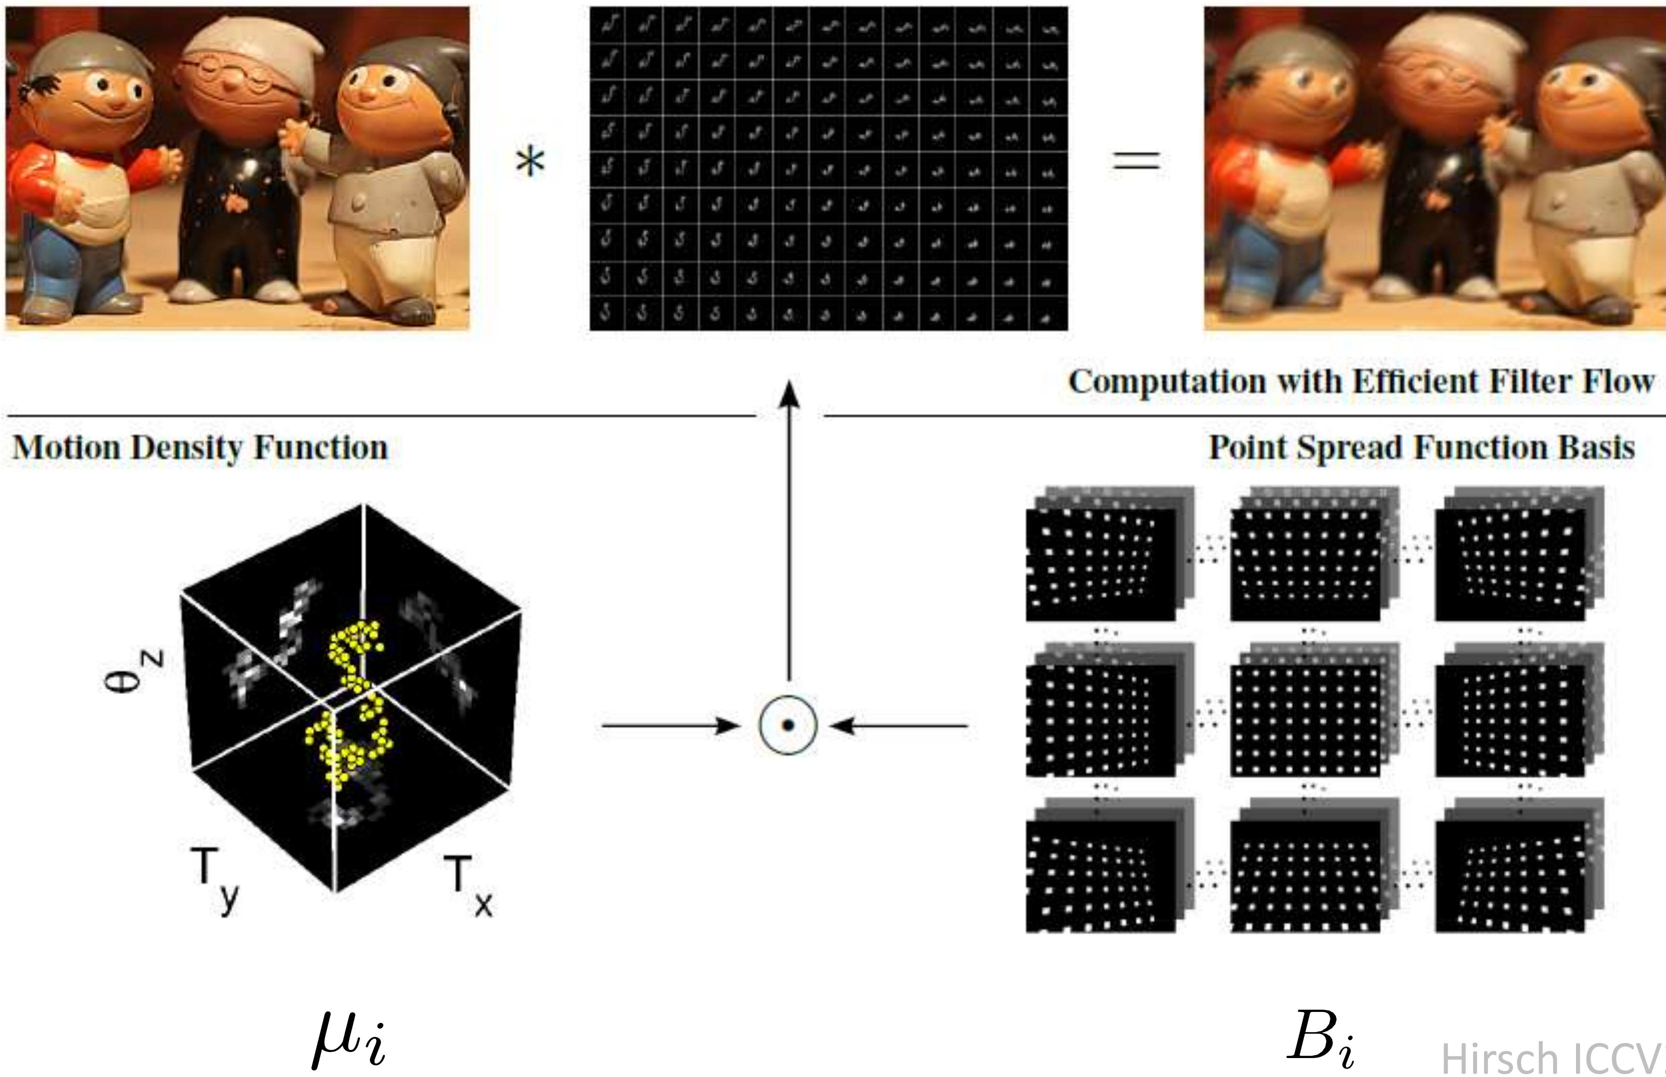

Hirsch ICCV2011

Object motion

- Estimate locally linear motion blur
  → line (2 parameters)

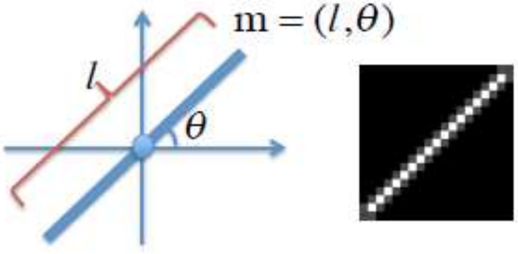

- Ignore occlusion

- Consider occlusion

Levin NIPS06
Chakrabarti CVPR10
Oliveira TIP14
Sun CVPR15

Shan ICCV07
Dai CVPR08

Ignore occlusion

Levin NIPS06
Chakrabarti CVPR10
Oliveira TIP14

- Local image statistics:

  Image autocorrelation increases in the blur direction

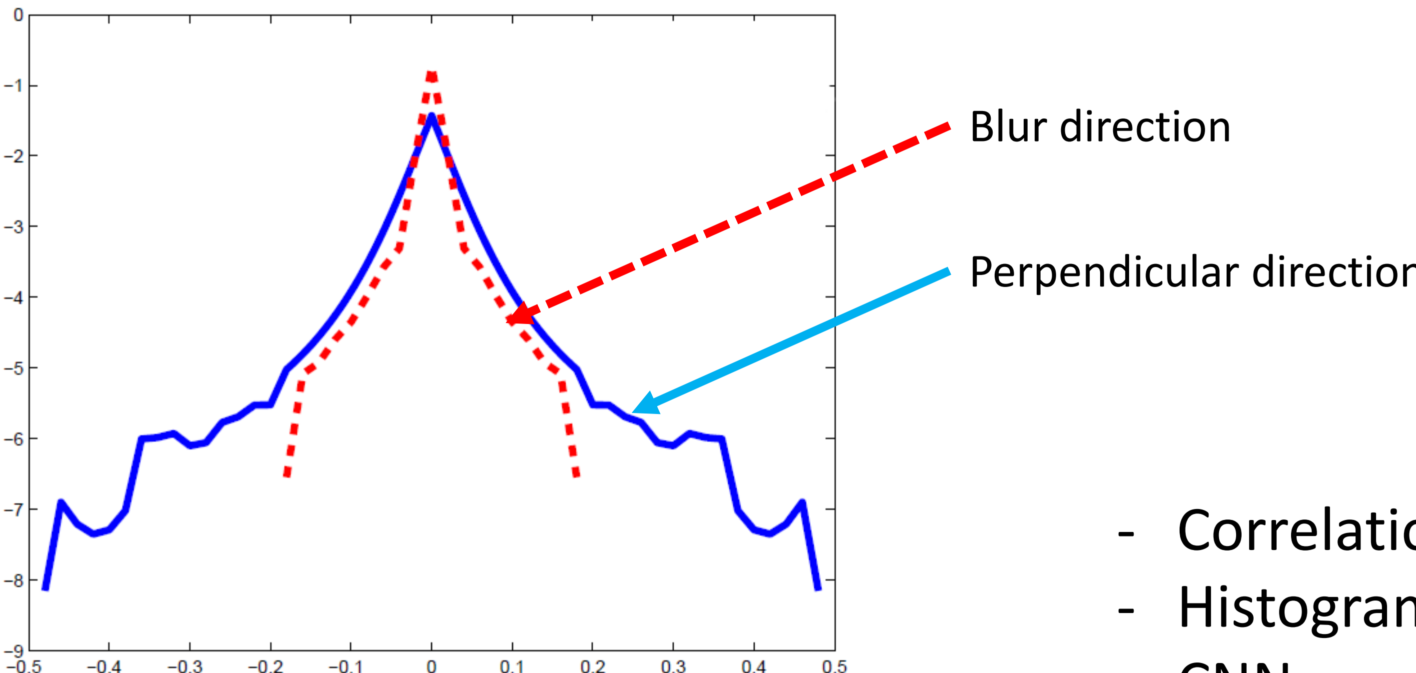

Derivatives histogram

- Correlation
- Histograms
- CNN Sun CVPR15

Ignore occlusion

Kim CVPR14

- Blind deconvolution with linear motion constraint

$$\min_{u,\mu} \|(h(\mu) * u) - z\|_1 + Q(u(x)) + R(\mu)$$

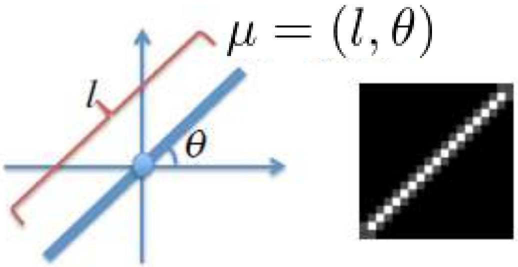

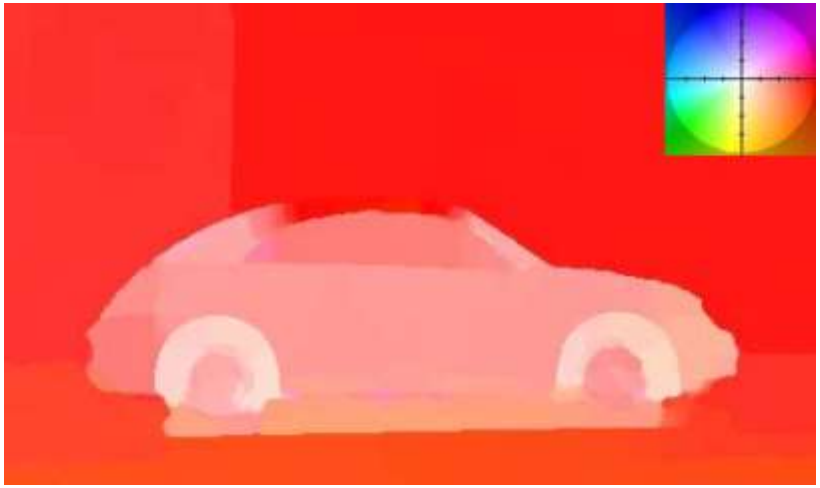

Occlusion

$$z = (1 - h * a)b + h * f$$

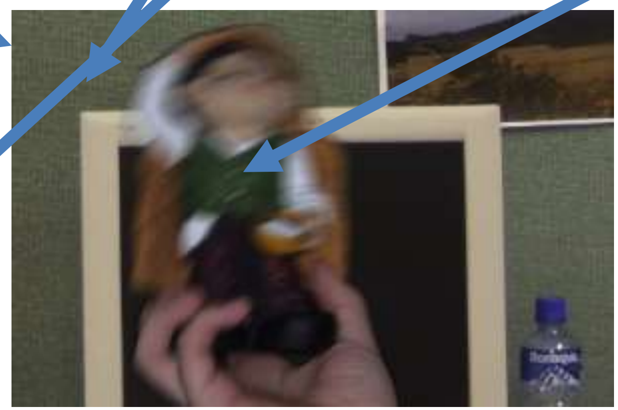

Alpha matting

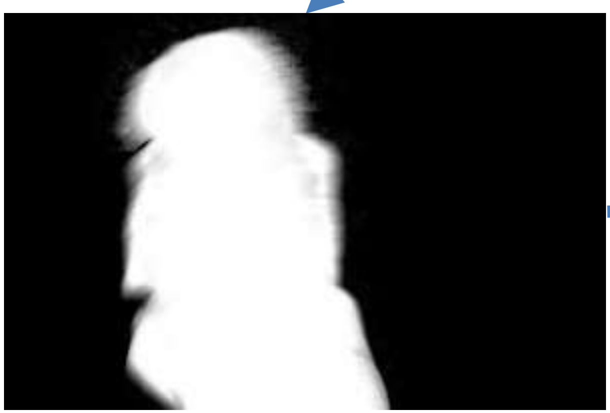

Blind deconvolution

Easier to solve
result is binary

h

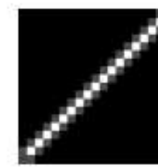

a

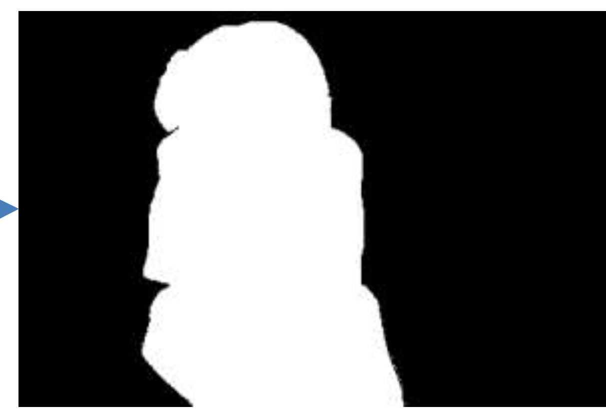

Conversion to space-invariant

- Wavefront coding – out-of-focus blur

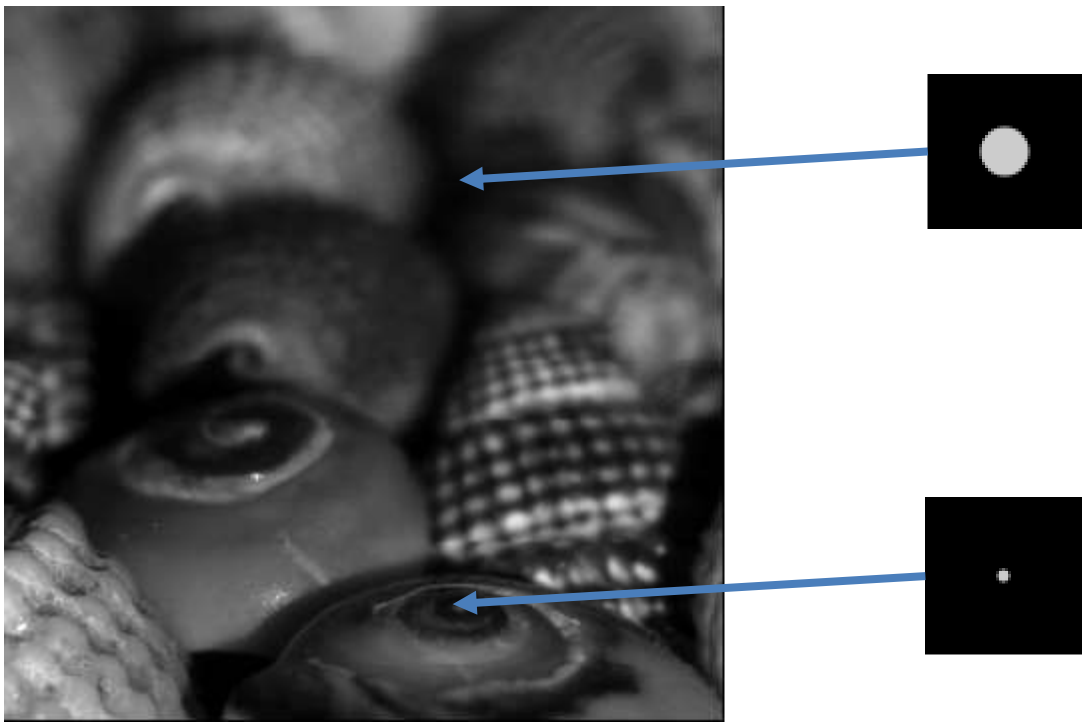

Conversion to space-invariant

- Wavefront coding

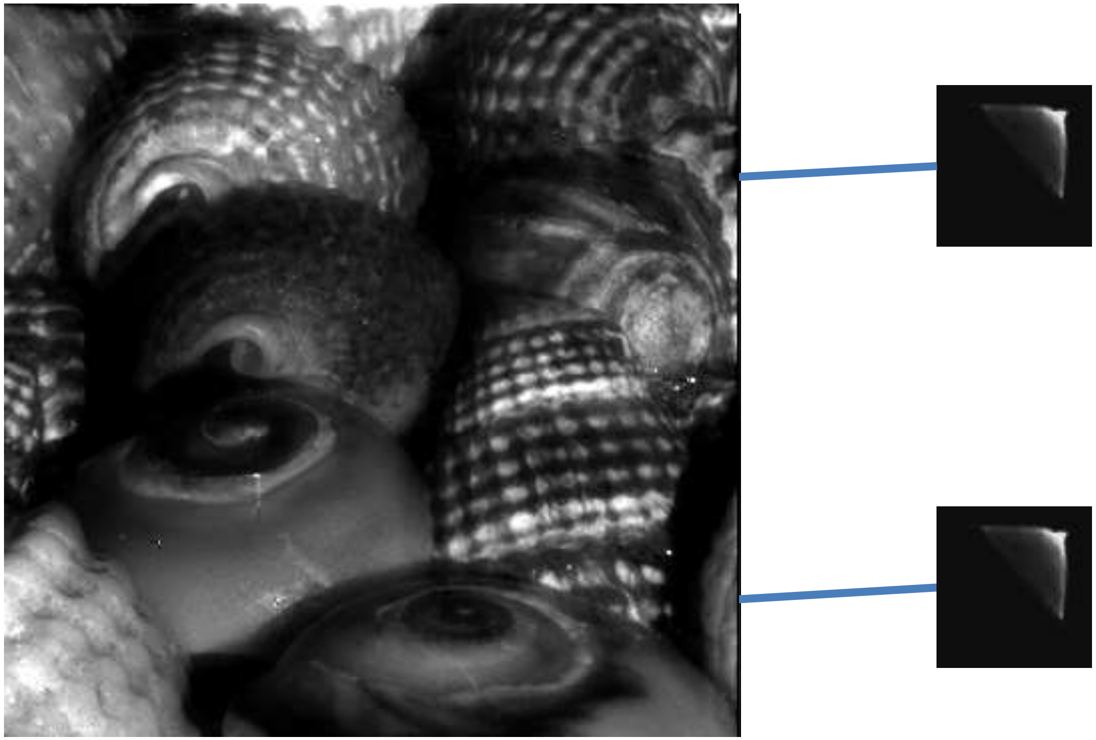

Space-invariant
non-blind
deconvolution

Conversion to space-invariant

- Motion-invariant photography
  - Space-variant linear motion blur

Levin SIGGRAPH08

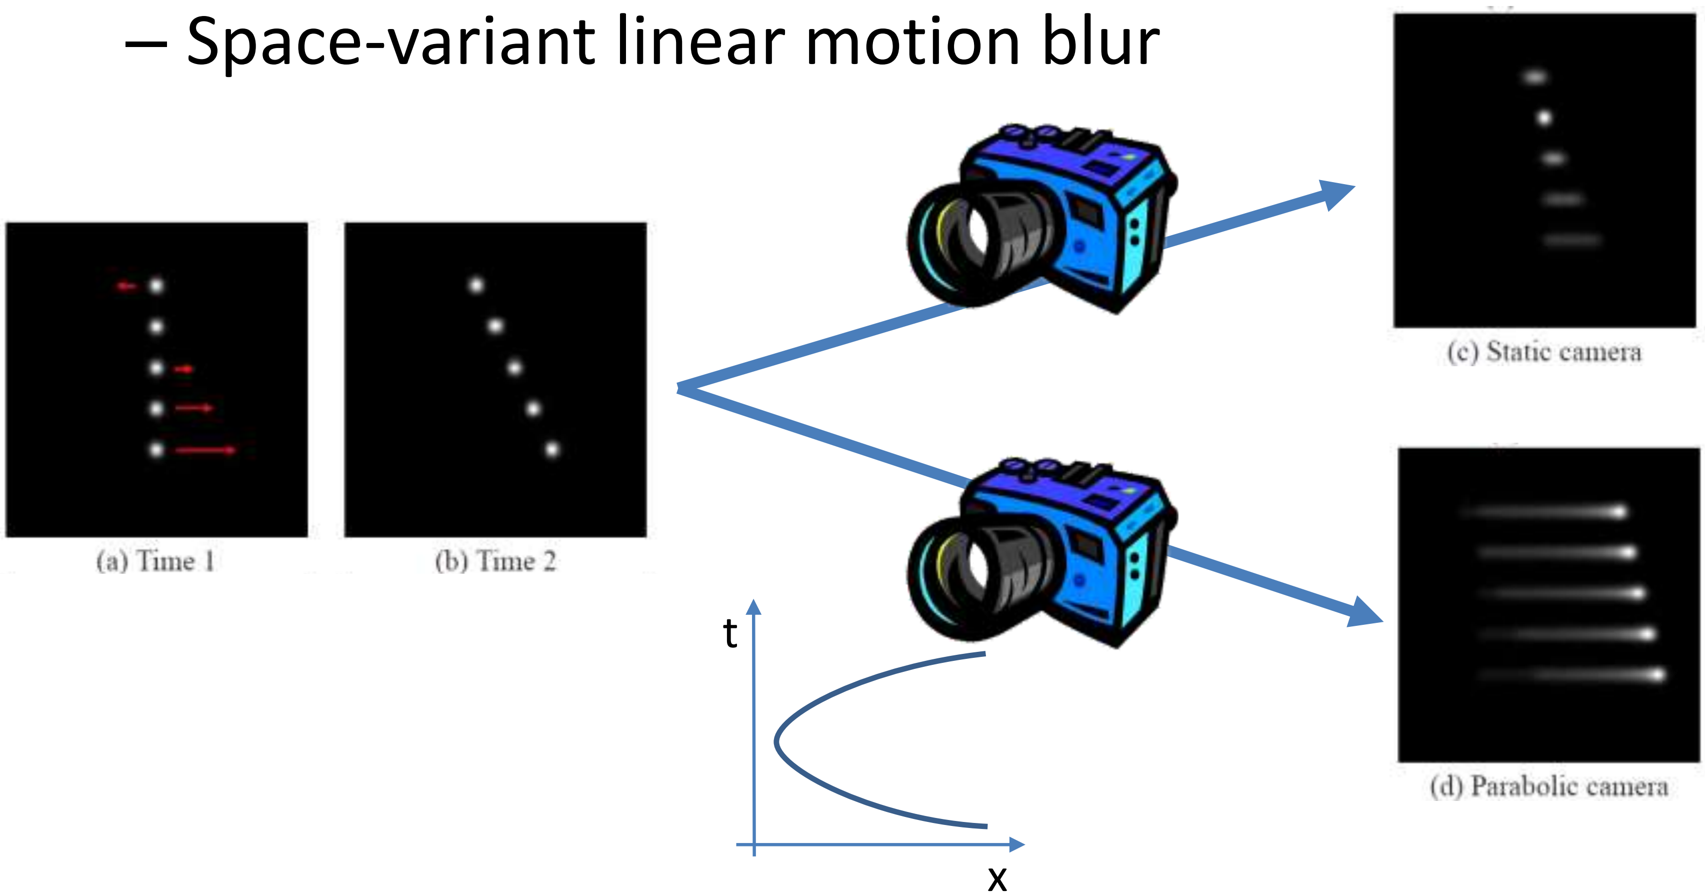

High-speed cameras

Ben-Ezra PAMI04
Tai PAMI10

- Low FPS HR camera with high FPS LR camera

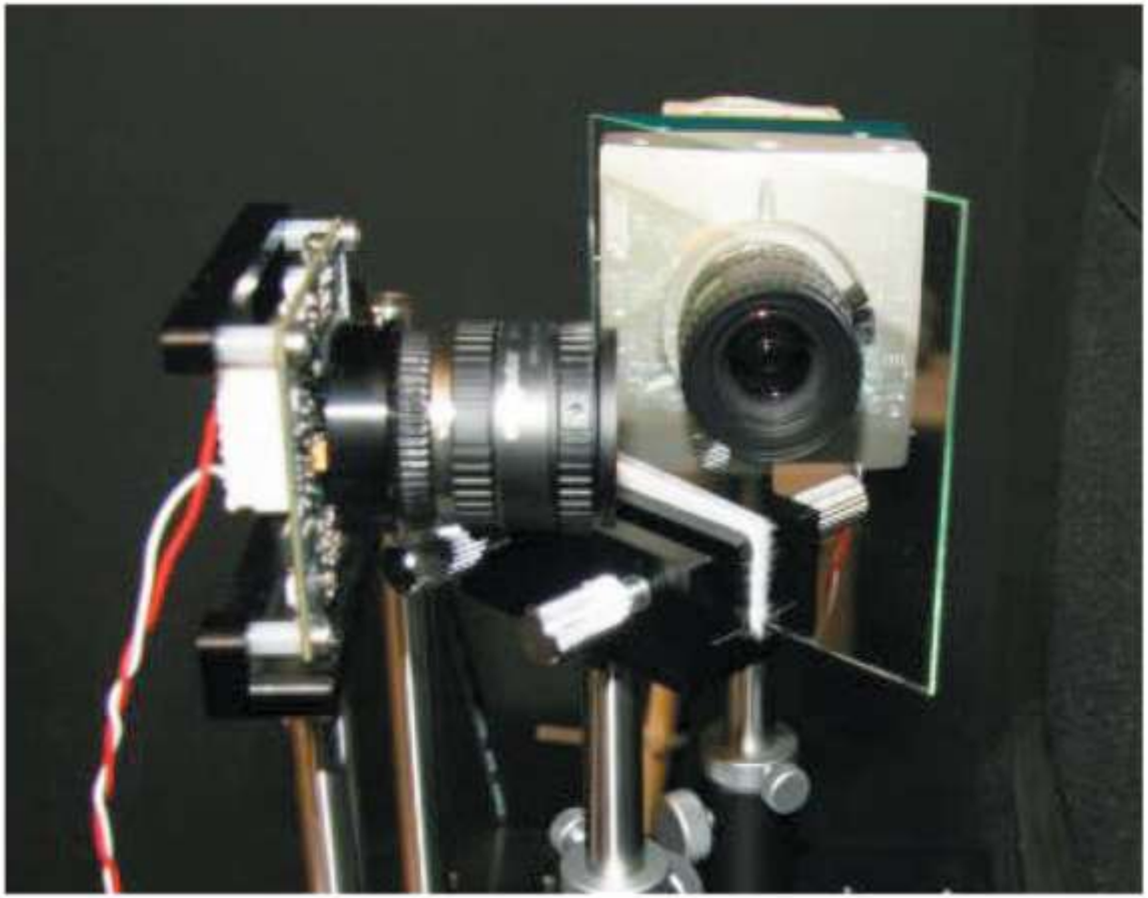

High-speed camera

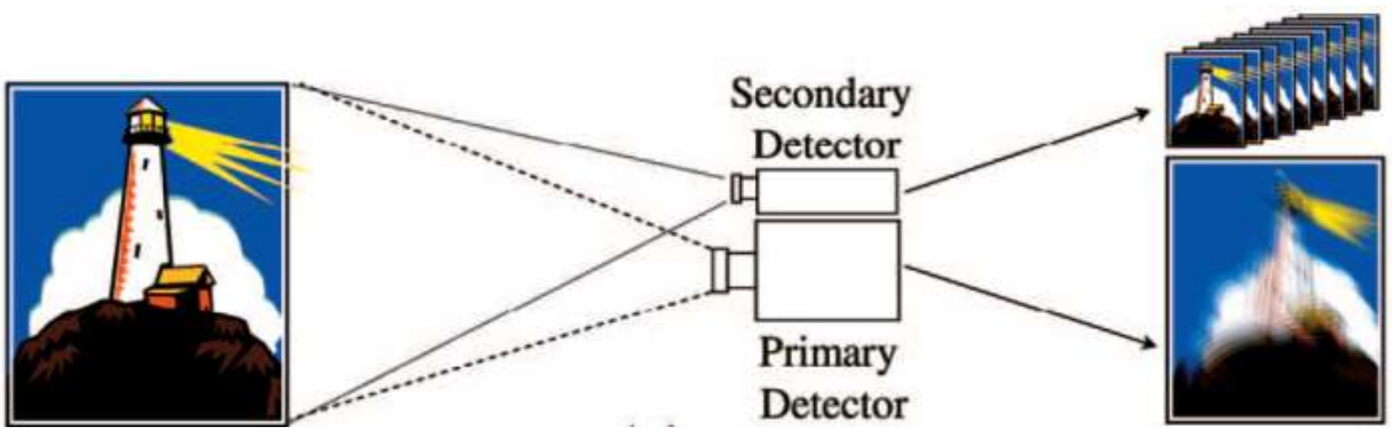

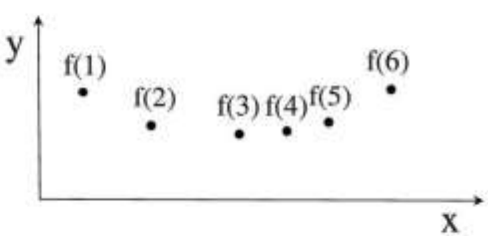

Sampled trajectory
(secondary detector)

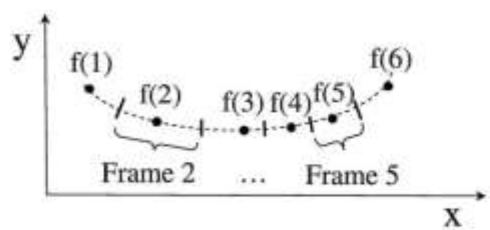

Interpolated trajectory

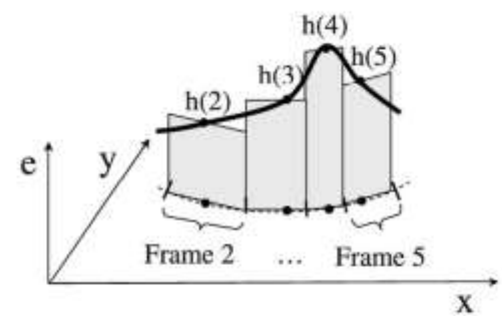

Estimated PSF

Space-time Processing

Takeda TIP11

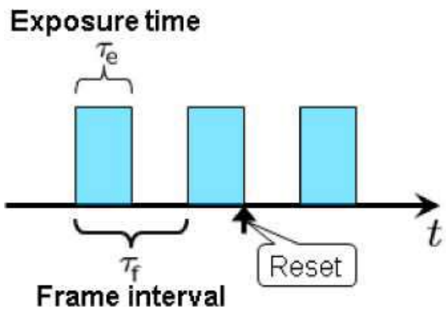

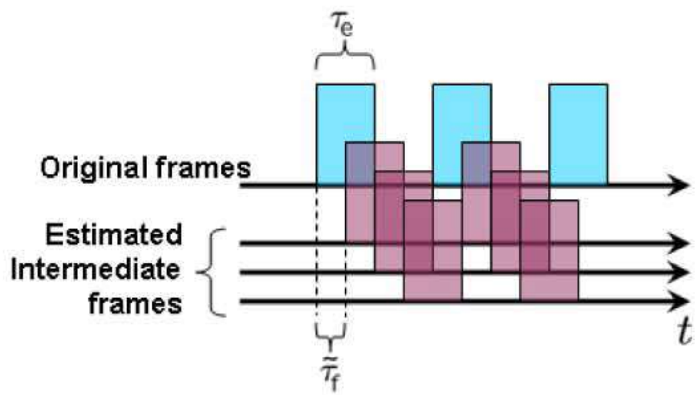

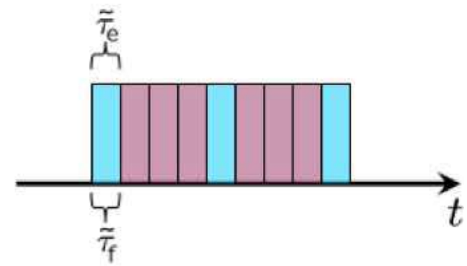

Simple deconvolution
in time

3D steering
kernel regression

Gyroscopes

Joshi ACM Trans.Graph.10
Sindelar JEI11
Levoy

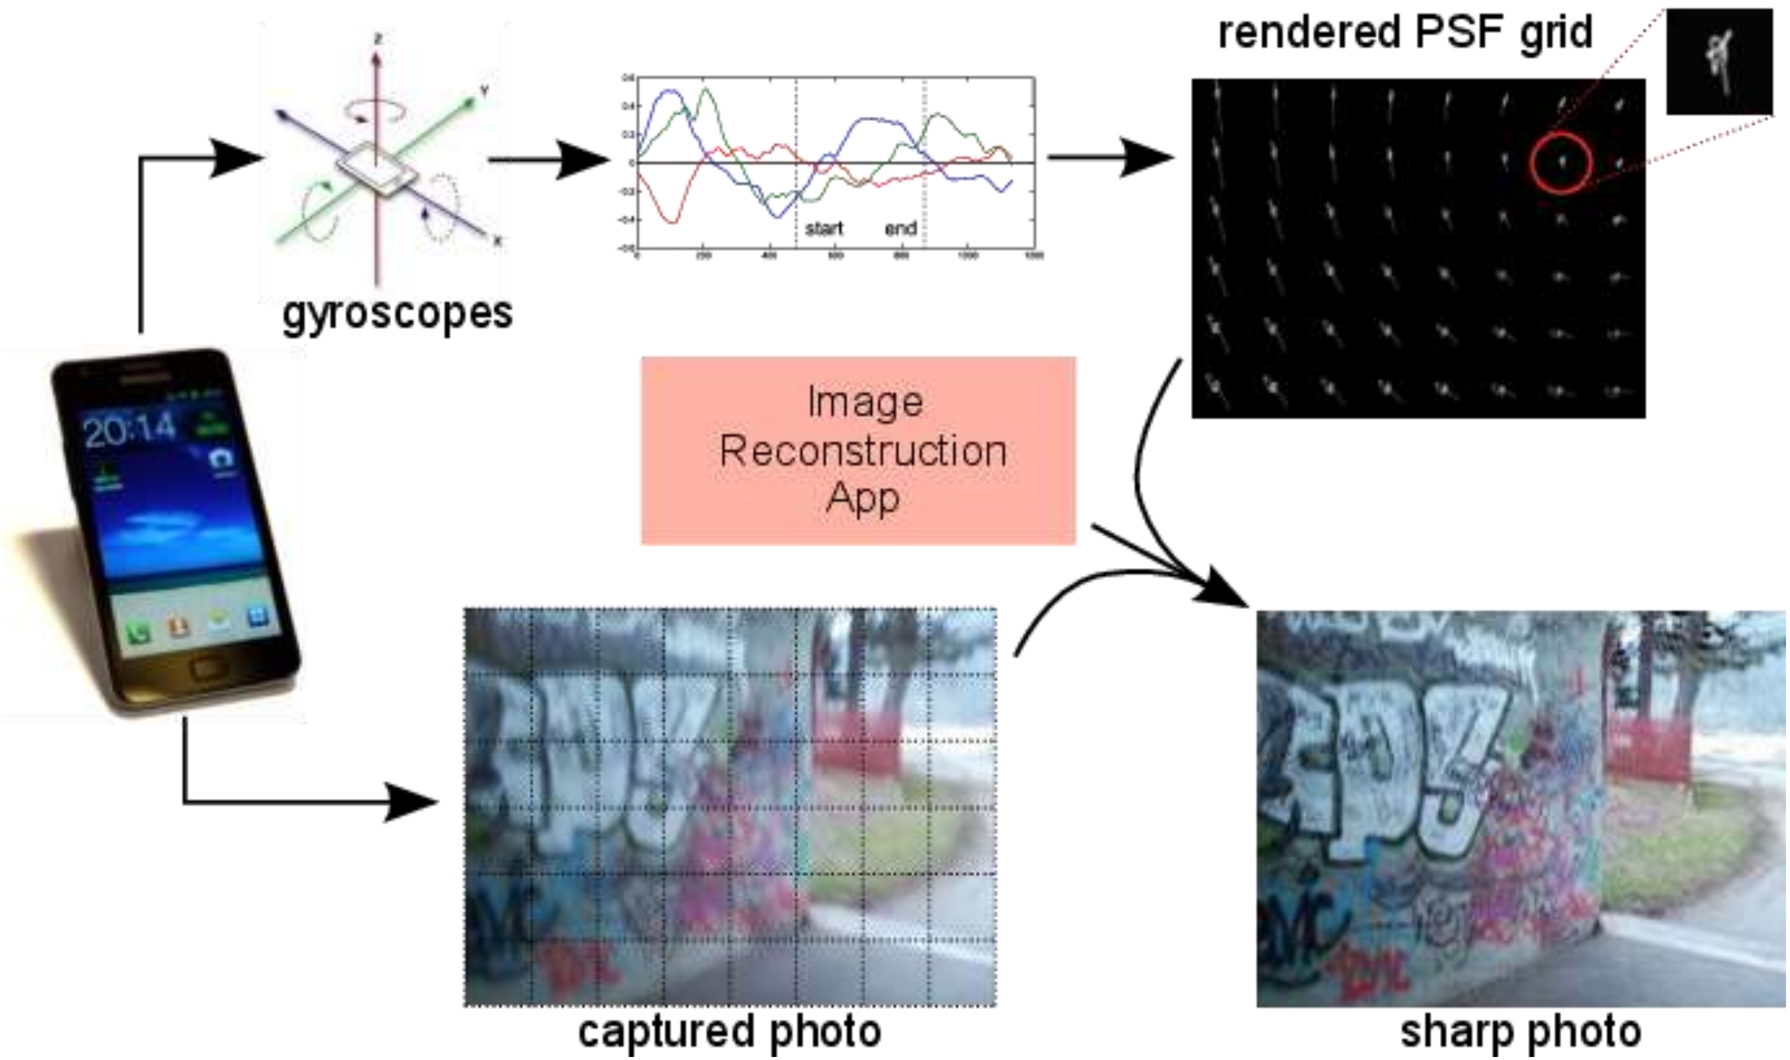

Acquired blurred image

Blur estimation

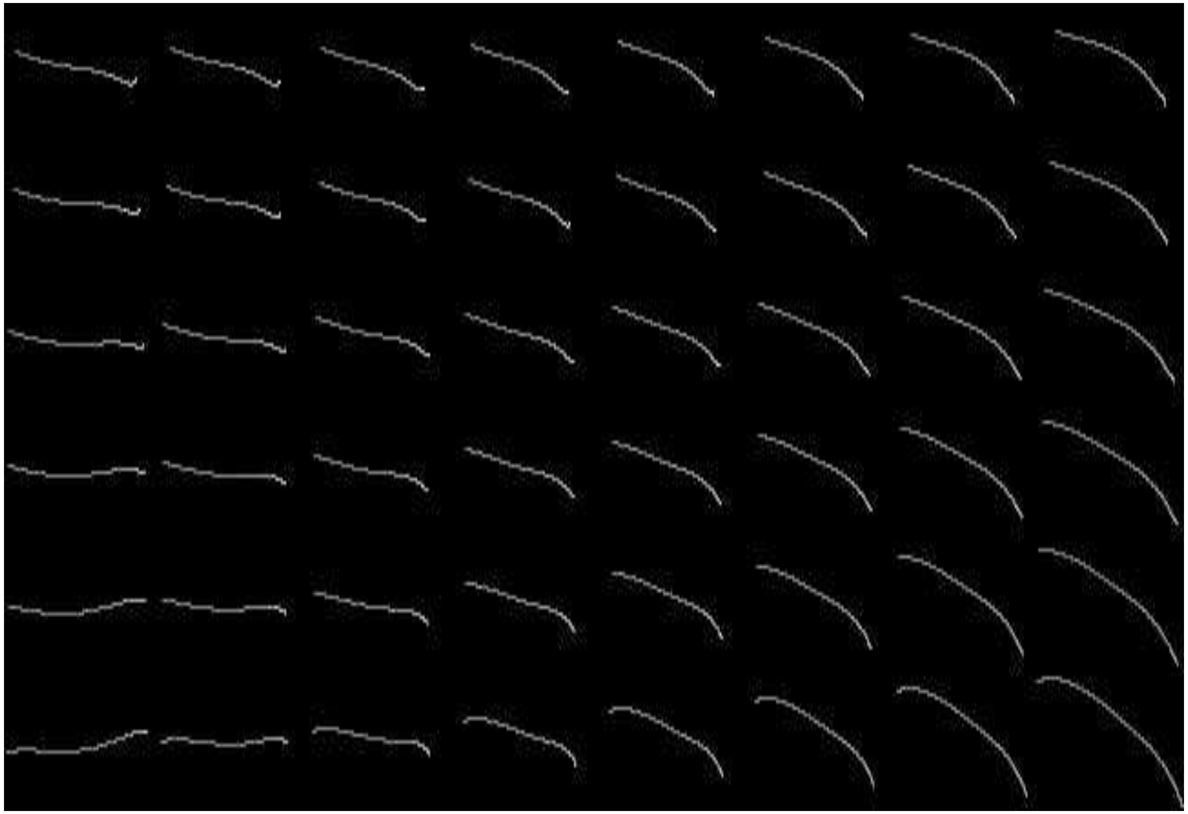

Patch-wise Deconvolution

Final Remarks

- MC is far more stable than SC BD.
- How to make SC more robust?
- Can Deep Learning make classical BD methods obsolete?

Thank You…