Digital Image Restoration

Blur as a chance and not a nuisance

Filip Šroubek

sroubekf@utia.cas.cz
www.utia.cas.cz

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

CMP 2017

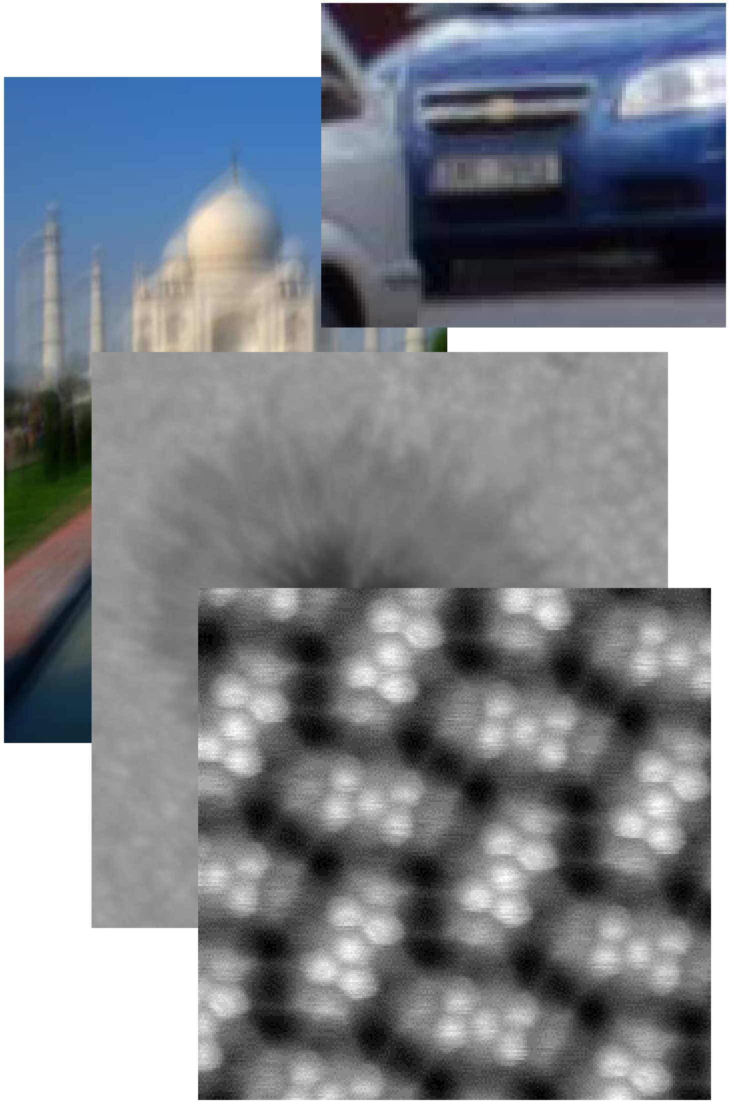

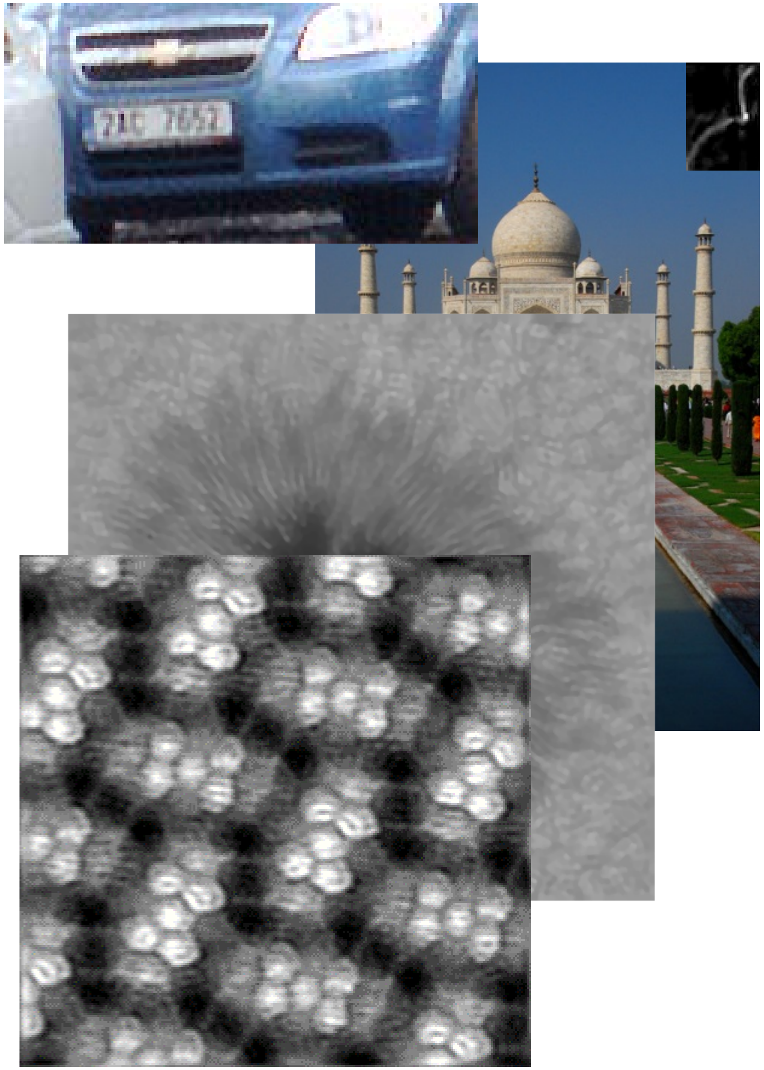

Image Degradation Model

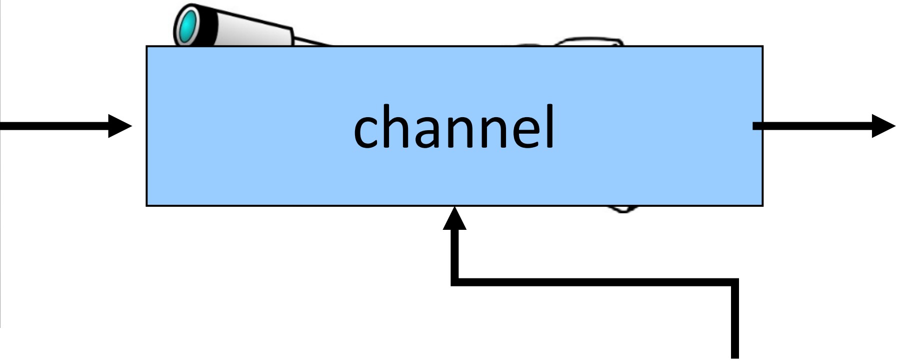

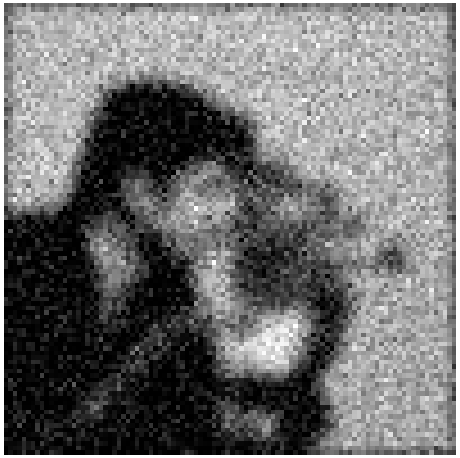

original image

noise

acquired image

$$\underbrace{[\, u * h \,]}_{\text{convolution}} + n = z$$

Energy Minimization

$$(\tilde{u}, \tilde{h}) = \arg\min_{u,h} E(u, h)$$

- 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.*

Energy Minimization

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

- Coupling of u and h → infinite number of solutions $(\tilde{u},\tilde{h})$
- Add regularization → well-posed problem
  → restrict solution vs. being general

Energy Minimization

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

Data term

Blur regularization

- Blur has different shapes
  - Compact support
  - Non-negative
  - Preserve energy

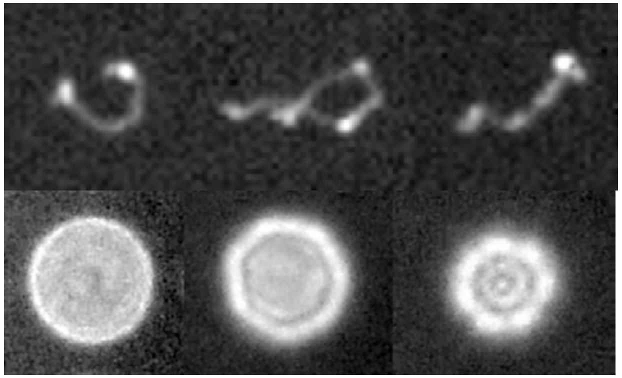
Motion blur

Out-of-focus & Abberations

Energy Minimization

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

Data term

Image regularization

- 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: $\int |\nabla u(x)|^p \,, \quad p < 1$

Statistics of sharp images

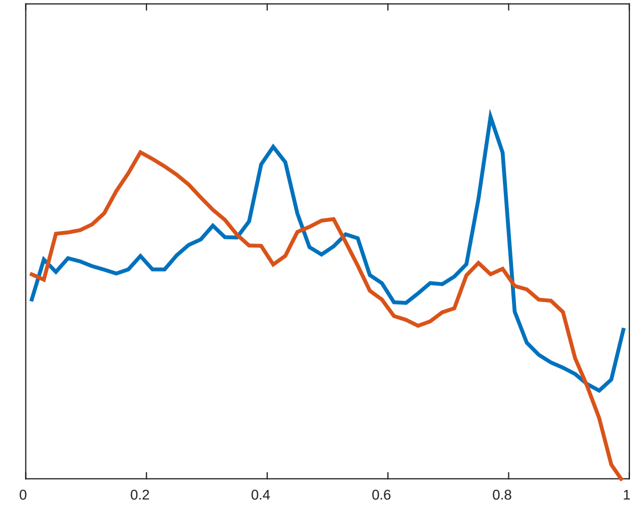

Intensity

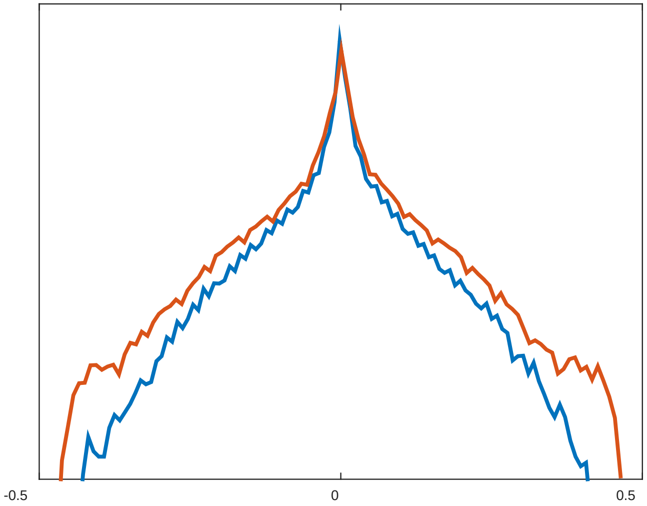

Gradient

Statistics of blurred images

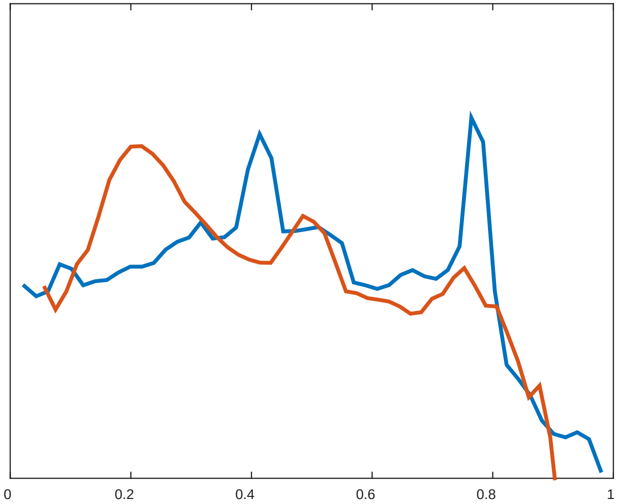

Intensity

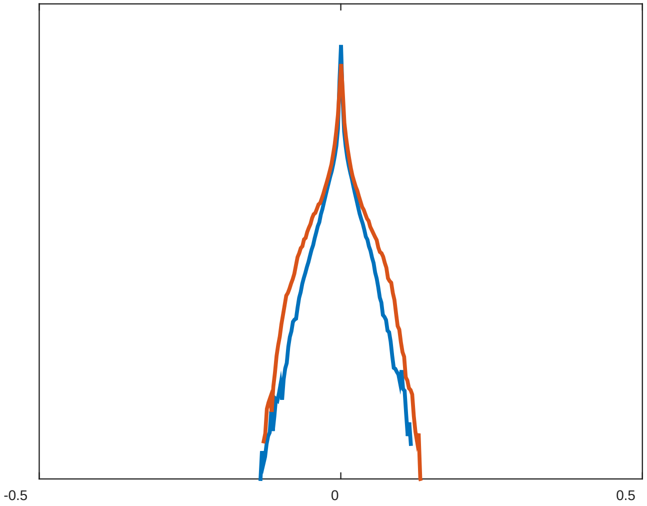

Gradient

Regularization

$u(x)$

$z(x)$

Different regularization $Q(u)$

$|\nabla u|^2$ $|\nabla u|^1$ $|\nabla u|^0$

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

Energy Minimization

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

- “NO-BLUR” solution: $\tilde{u}(x) = z(x), \quad \tilde{h}(x) = \delta(x)$
- WHY?
- Regularization favors blurred images

Regularization favors blur

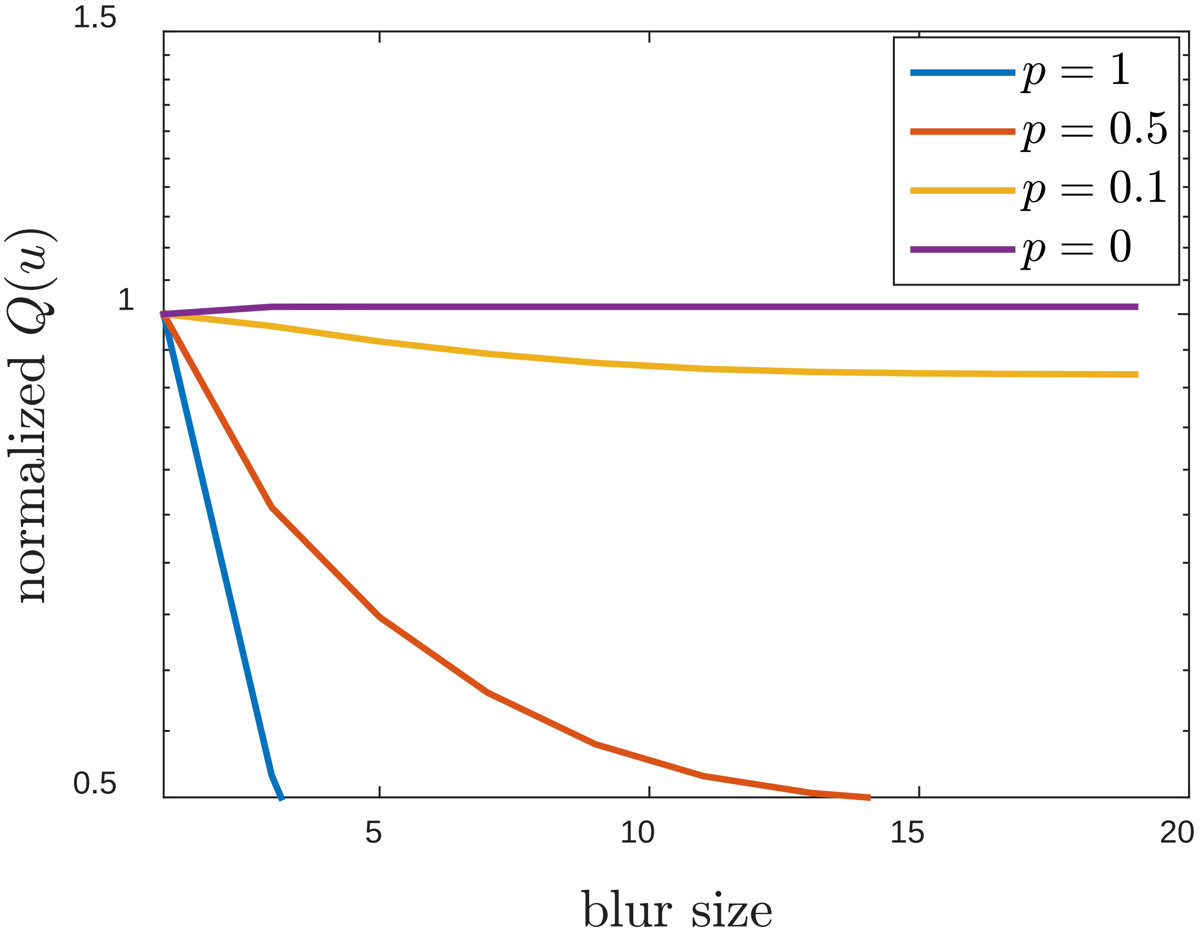

$$Q(u) = \int |\nabla u|^p$$

Regularization favors blur

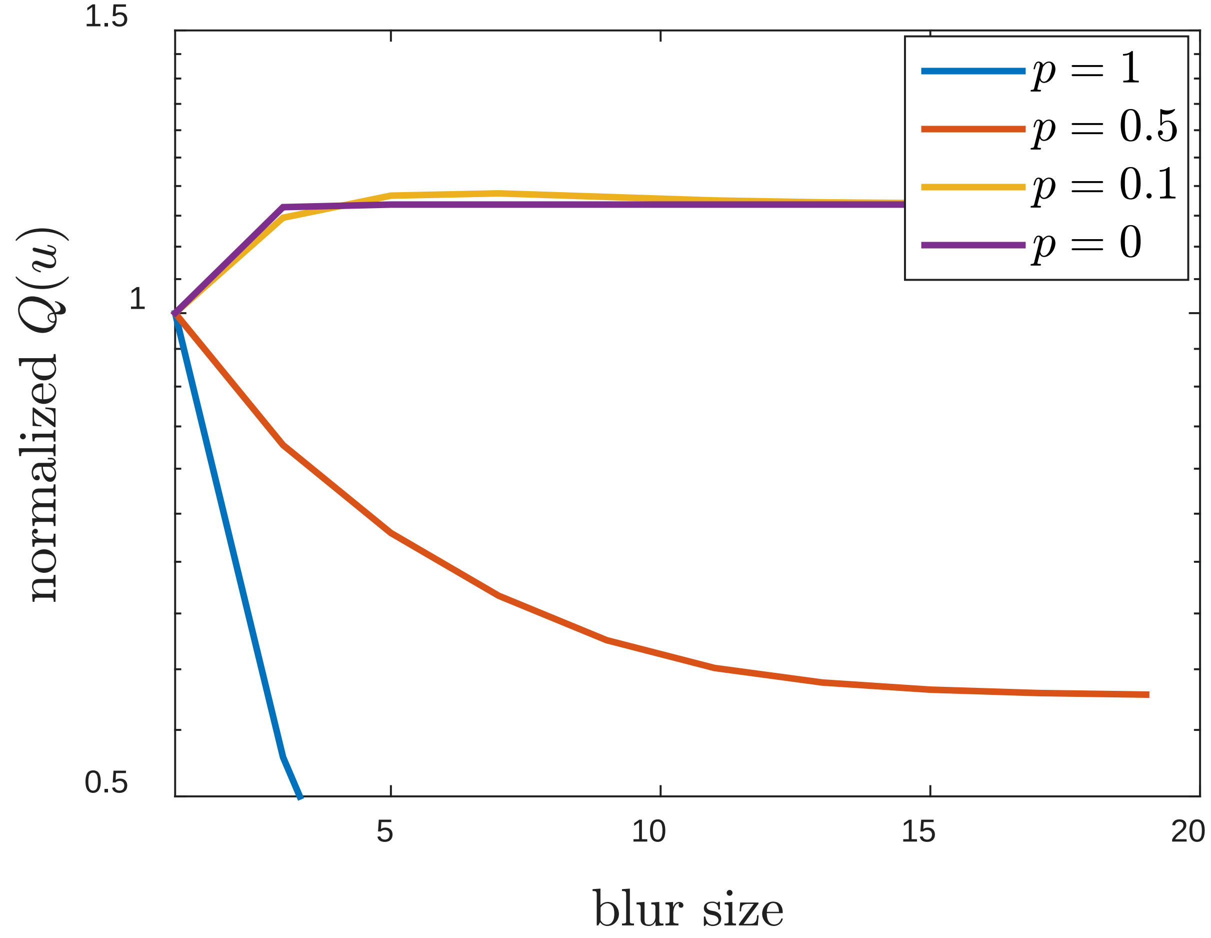

$$Q(u) = \int |\nabla u|^p$$

Artificially sparsify images

Ad-hoc steps!

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

Chan TIP98
Shan SigGraph08
Cho SigGraph09
Xu ECCV09, CVPR13
Almeida TIP10
Krishnan CVPR11
Zhong CVPR13
Sun ICCP13
Michaeli ECCV14
Perrone PAMI15
Pan CVPR16

Artificial Sharpening

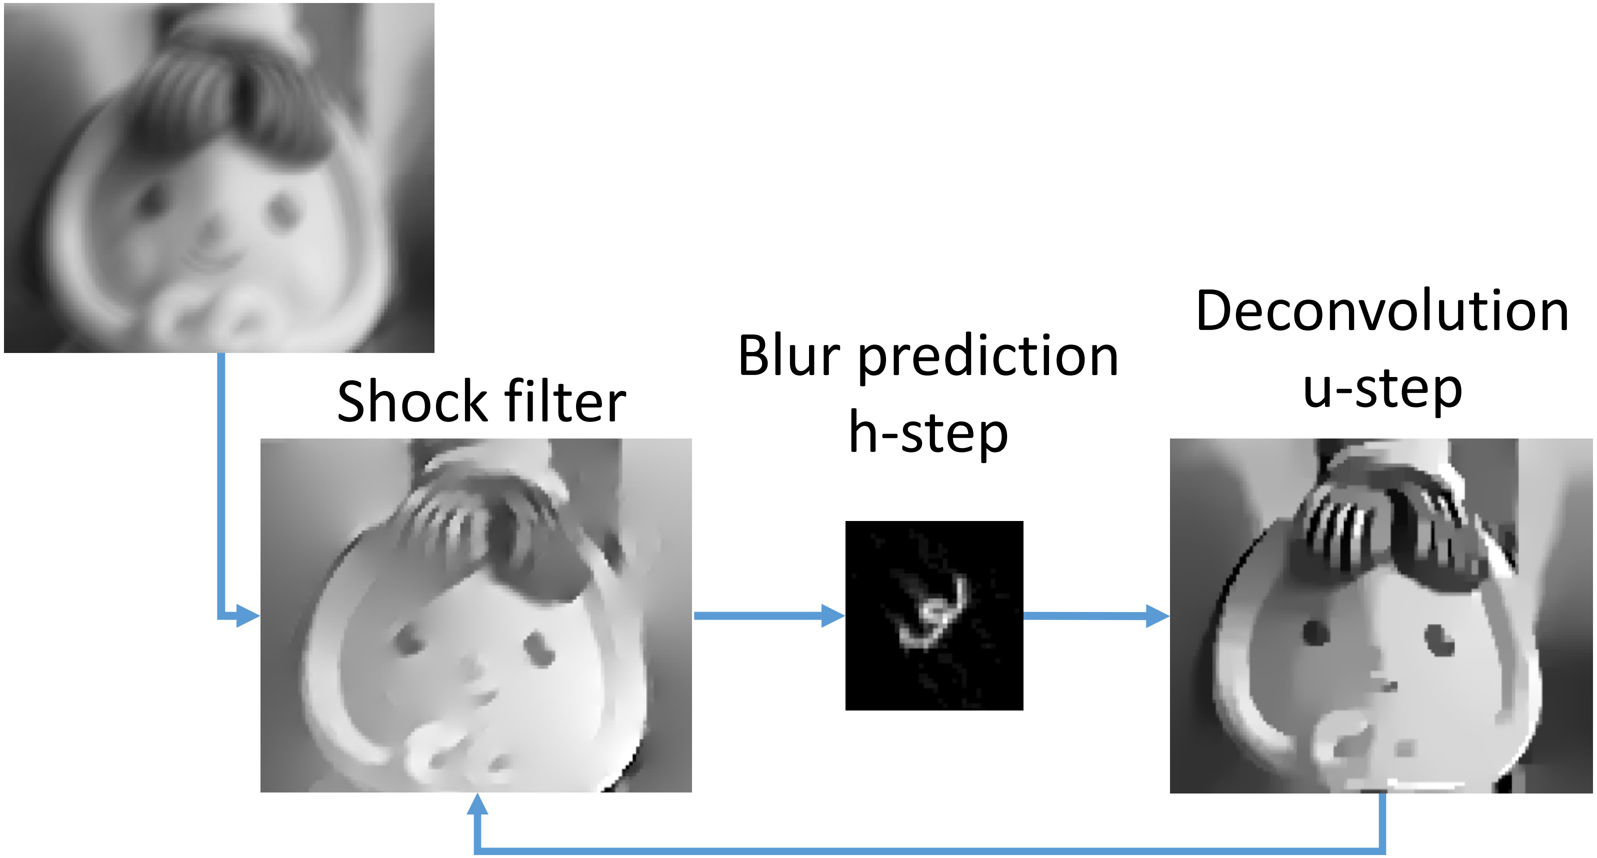

Blurred image

Shock filter

Blur prediction
h-step

Deconvolution
u-step

Cho et al., SIGGRAPH 2009

Remove Spiky Objects

Xu et al., ECCV 2010

Remove Spiky Objects

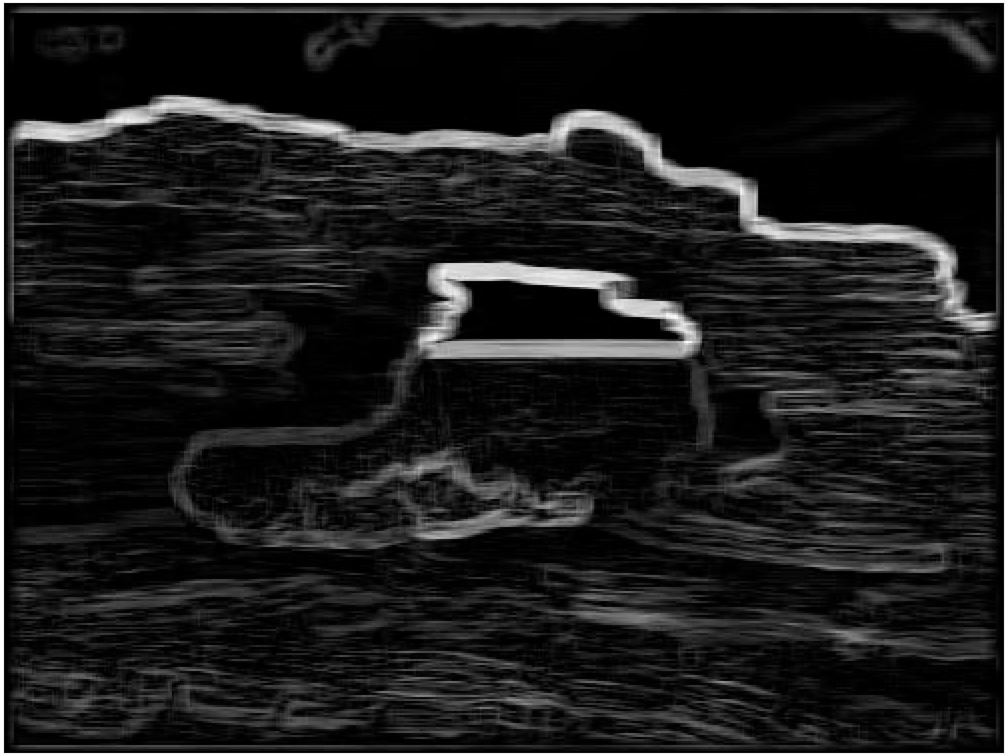

Mask out small objects

Xu et al., ECCV 2010

Remove Spiky Objects

Reconstructed image with small objects removed

Xu et al., ECCV 2010

Hierarchical Deconvolution

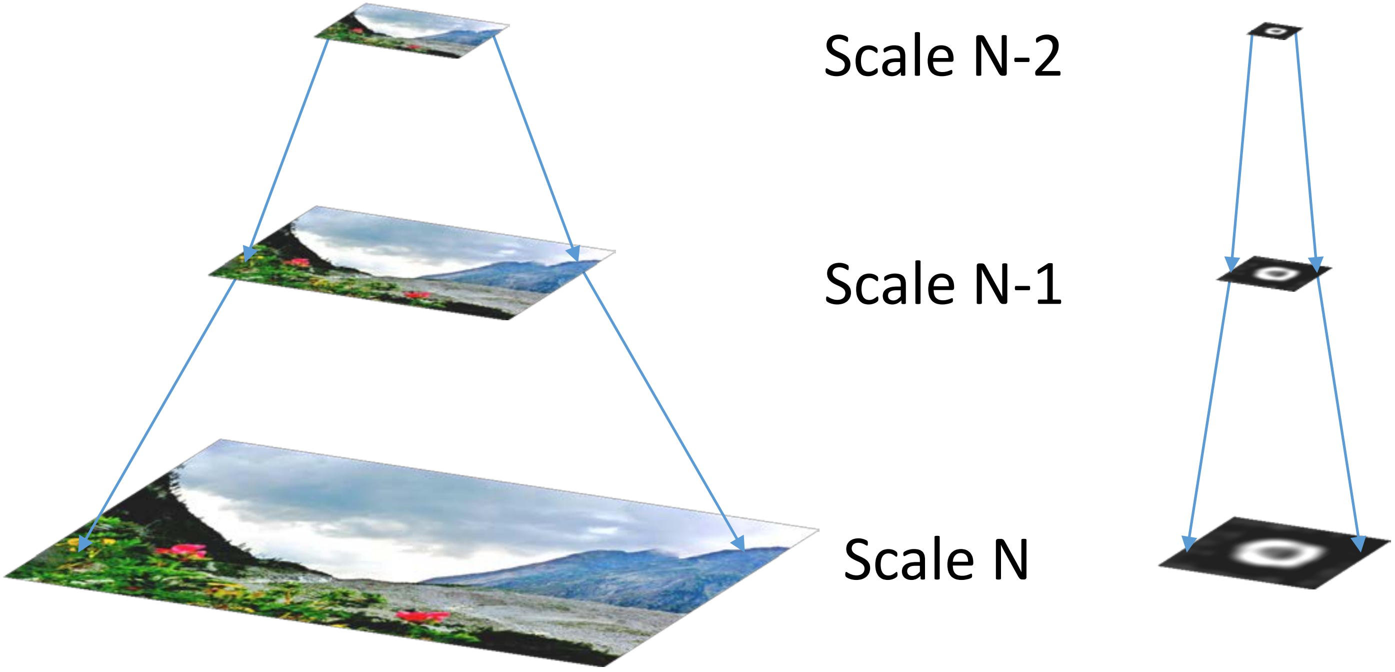

Scale N-2

Scale N-1

Scale N

image

blur

Bayesian Paradigm

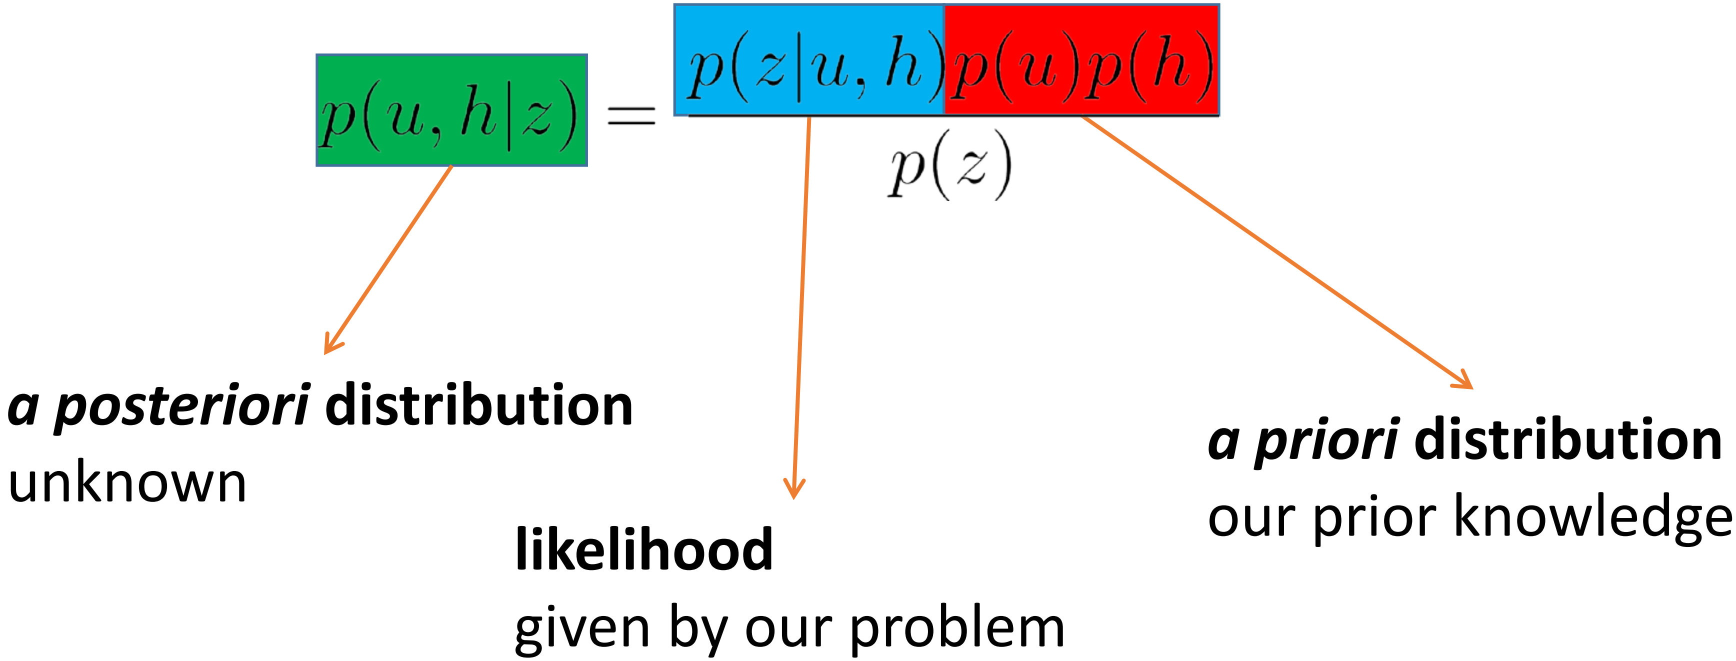

$$p(u,h|z) = \frac{p(z|u,h)p(u)p(h)}{p(z)}$$

***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.

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

- Then maximize the posterior

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

Delta function

$\tilde{h} = \delta(x)$

MAP

$p(u, h|z)$

no-blur solution

Marginalize

$p(h|z)$

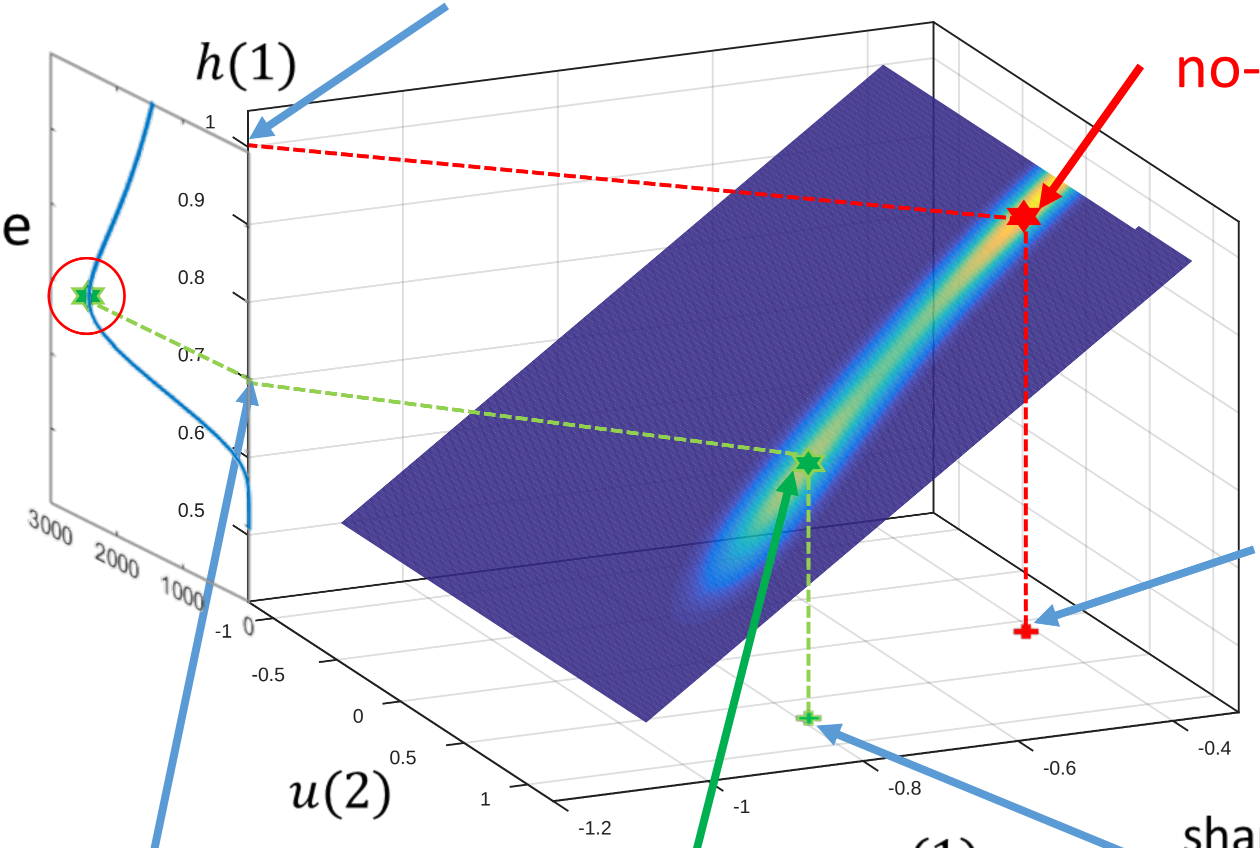

blurred image

$\tilde{u} = z(x)$

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.

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

Miskin AICA00
Fergus SIGGRAPH06
Tzikas TIP09
Whyte IJCV12
Levin PAMI11
Babacan ECCV12
Wipf JMLR14

Marginalization in blind deconvolution

Blurred image
$z(x)$

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

Limitations of BD

- Gaussian noise

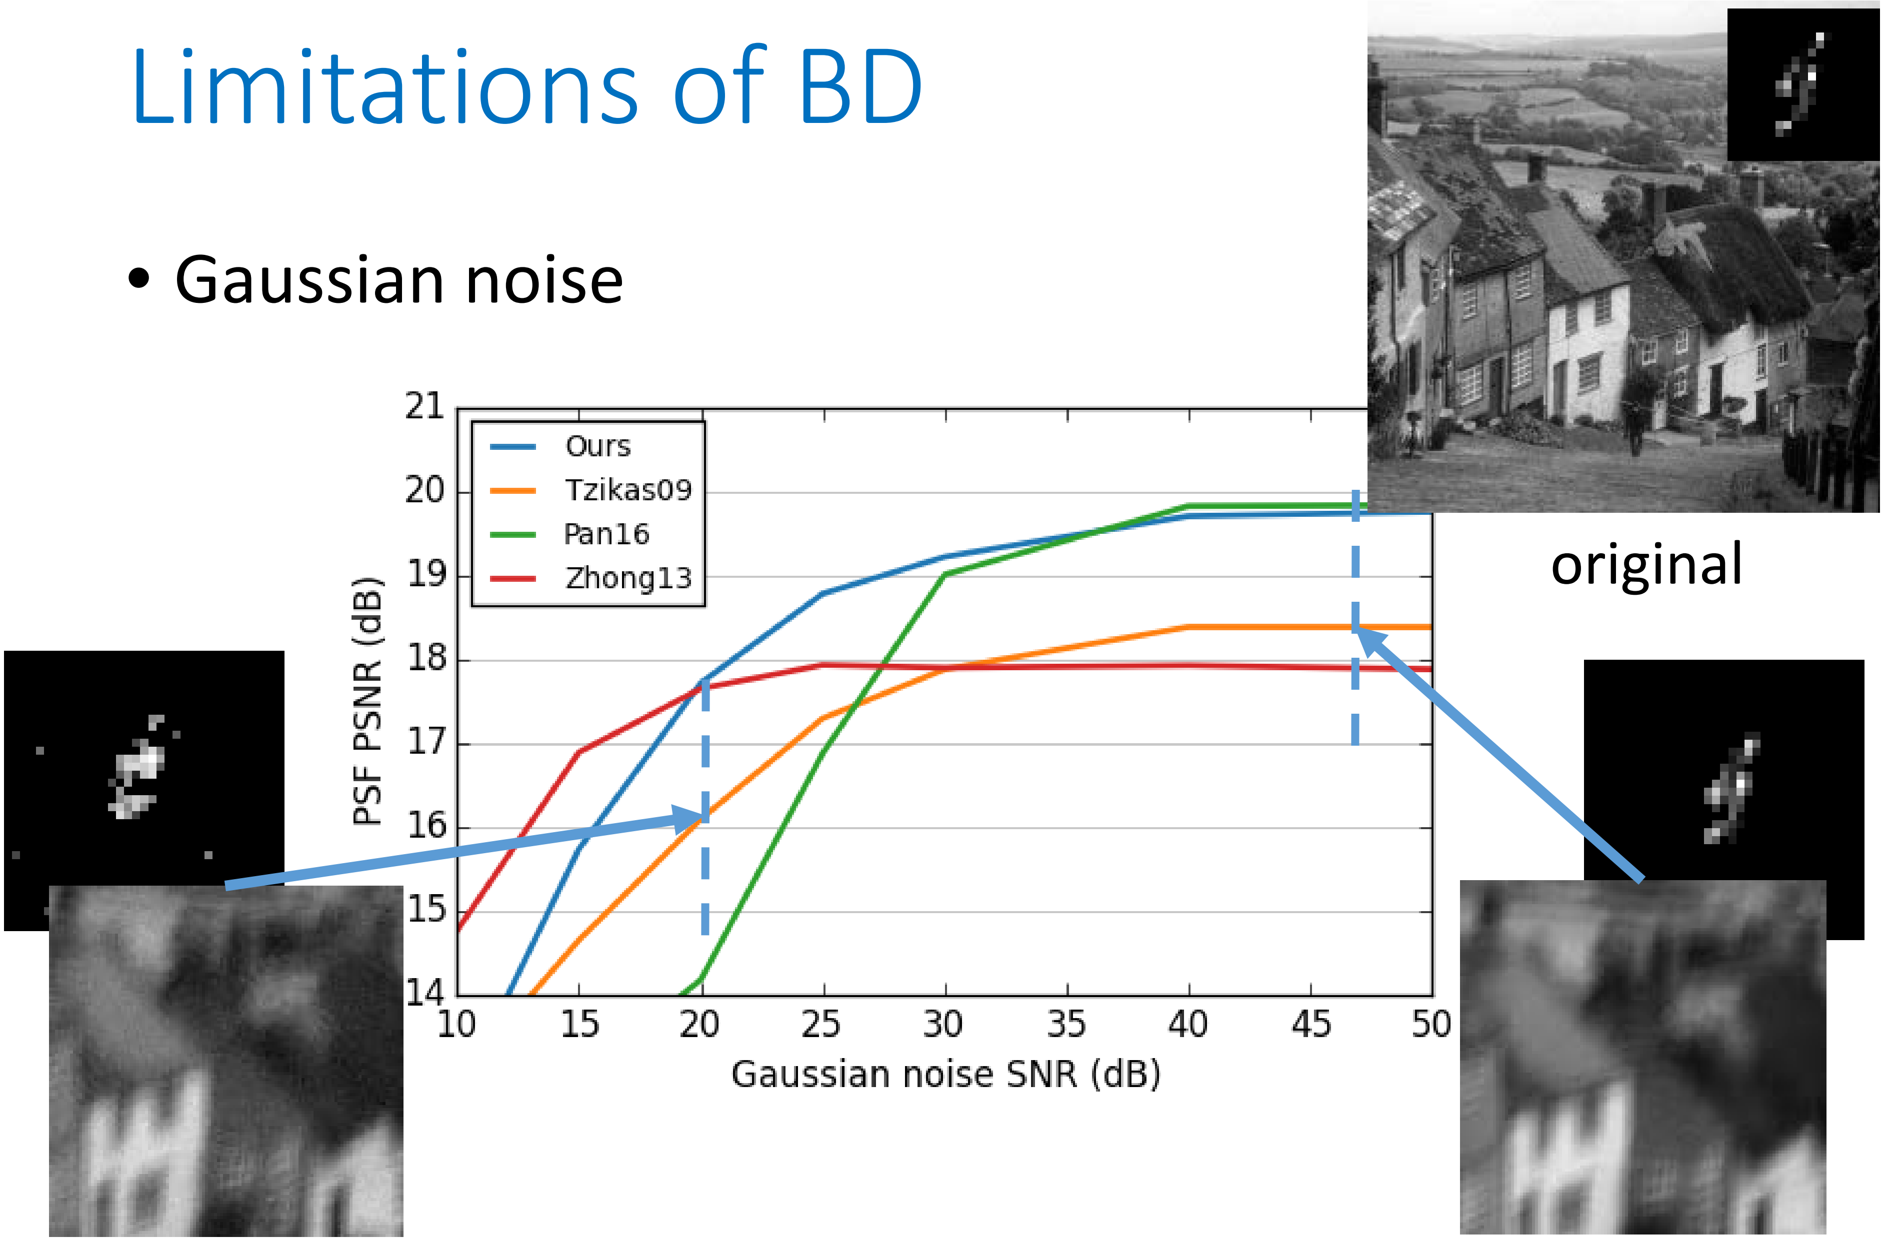

original

SNR=20dB

SNR=50dB

Limitations of BD

- Model discrepancies

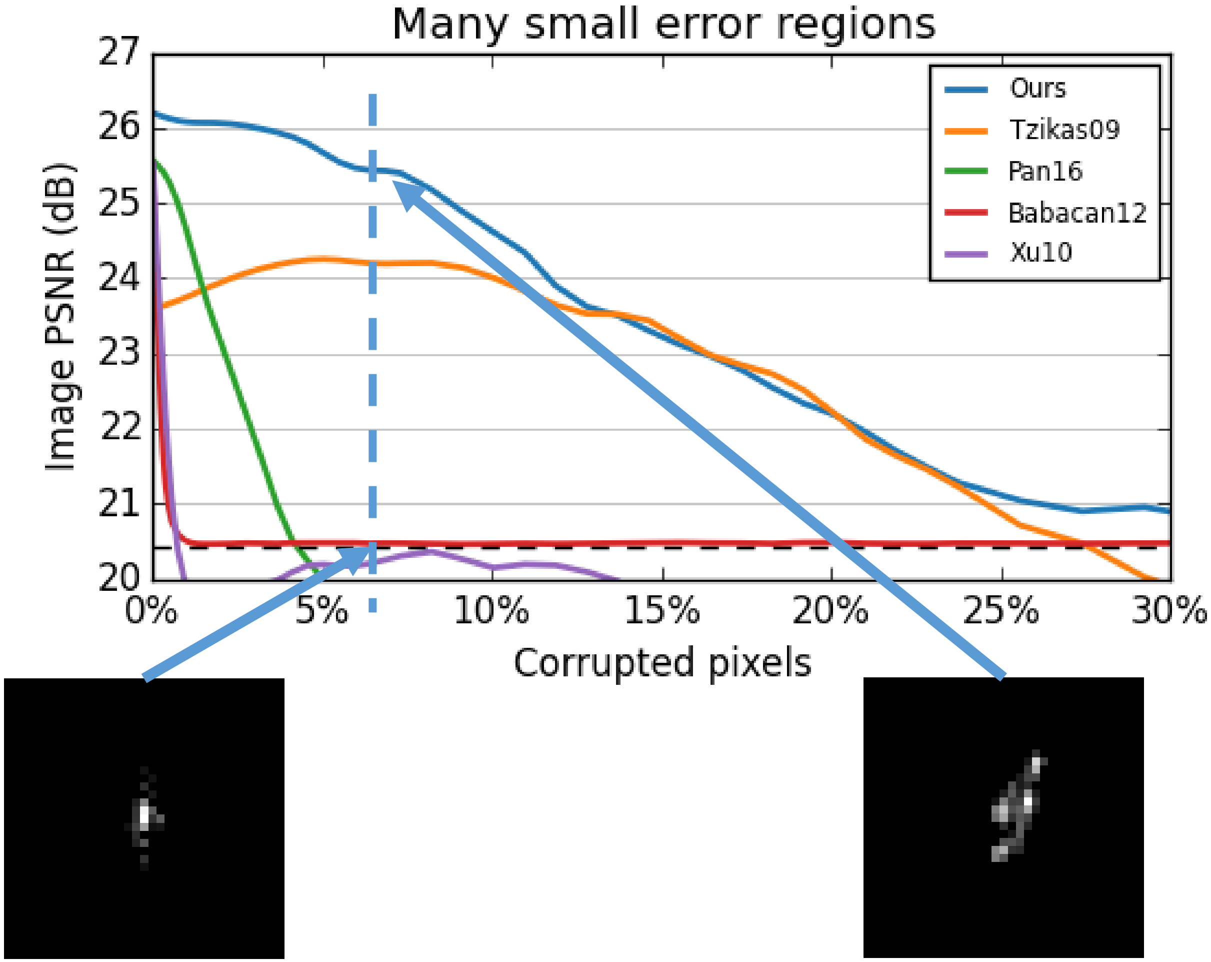

original

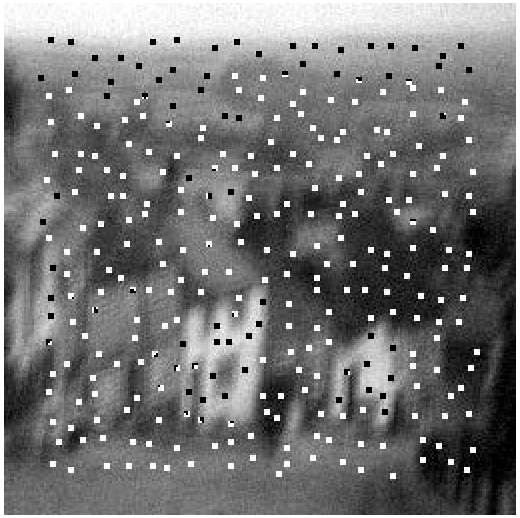

7% discrepancy

Automatic Relevance Determination (Students’ *t*-distribution)

- Standard data term

$$\|h * u - z\|^2 = \int \big([h * u](x) - z(x)\big)^2$$

replaced by

$$\|h * u - z\|_\gamma^2 = \int \gamma(x)\big([h * u](x) - z(x)\big)^2$$

z

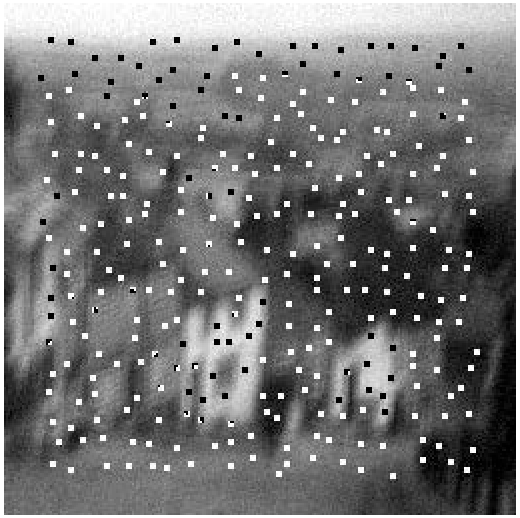

γ

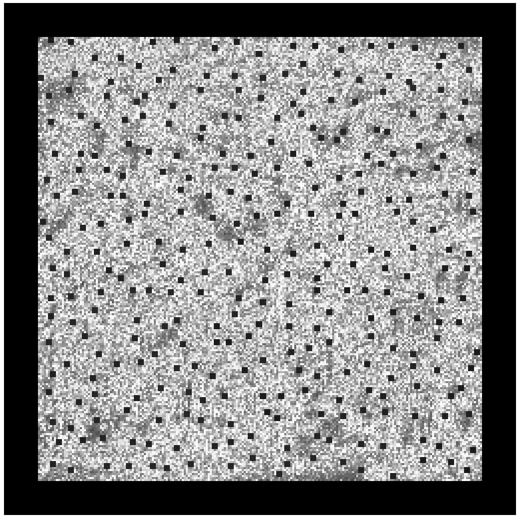

with γ

without γ

Space-variant Blur

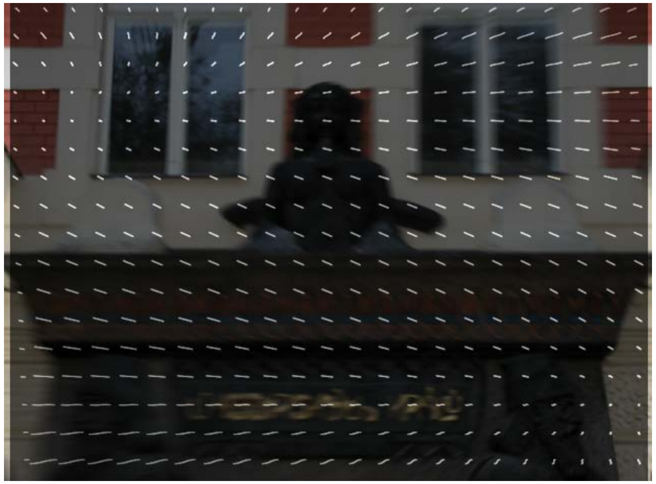

Camera motion

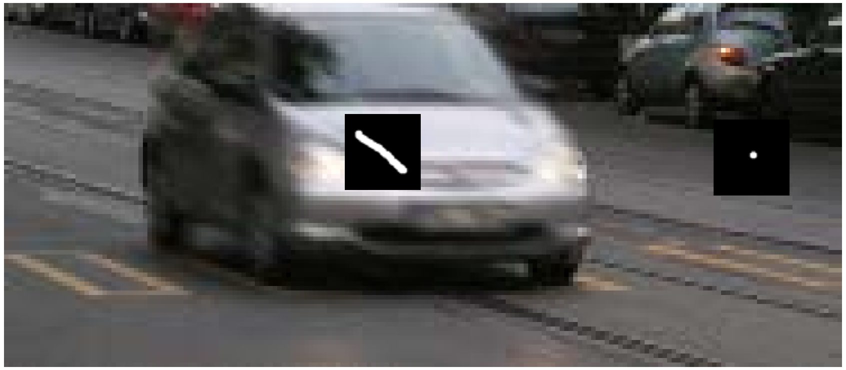

Object motion

Optical aberrations

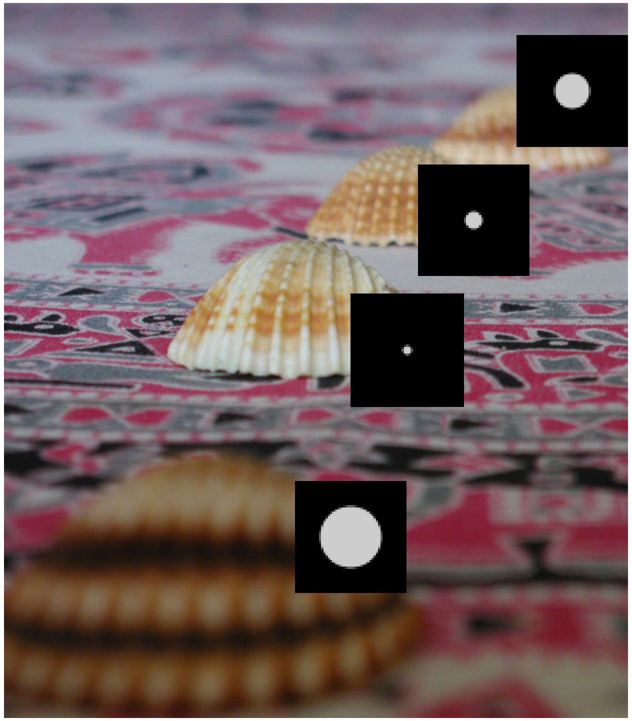

Scene depth

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

- Conversion to space-invariant
  - HW design

Levin SIGGRAPH08, Ben-Ezra PAMI04,
Tai PAMI10, Wavefront coding

- Local object motion
  - Segmentation
  - Line blur

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

Fast Moving Objects

D.R.,J.K.,F.S.,L.N.,J.M.
arXiv:1611.07889

- FMO moves over a distance exceeding its size within exposure time

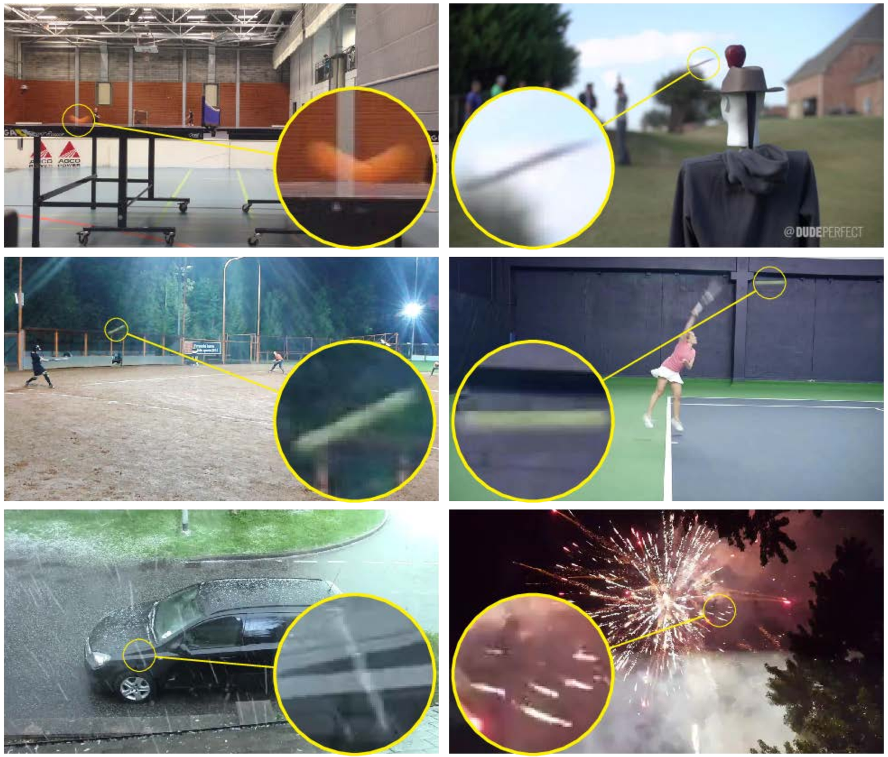

Formation model

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

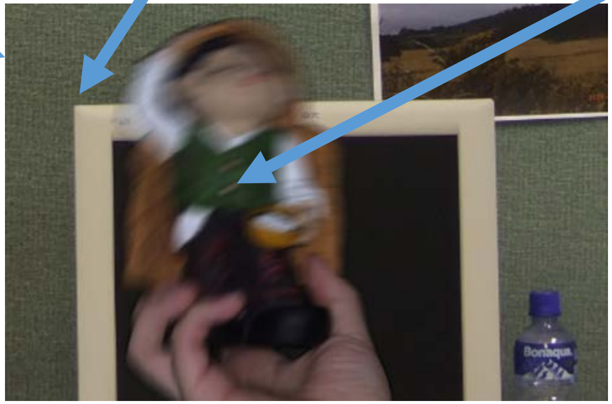

FMO Appearance Estimation

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

- Assumptions:
  - Background b is known → previous frame
  - Blur h is a FMO trajectory → estimated by a tracker
- Optimization problem:

$$E(f, m, h) = \|(1 - h * m)b + h * f - z\|_1 + \lambda Q(f)$$

- Alternating minimization w.r.t. f, m, h
- Initial estimation of h is critical → novel tracker designed for FMOs

FMO Localization

- The pipeline consists of three stages:

1) Explorer

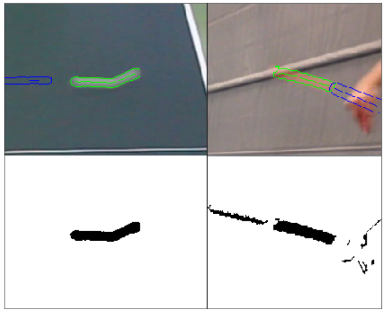

2) Detector

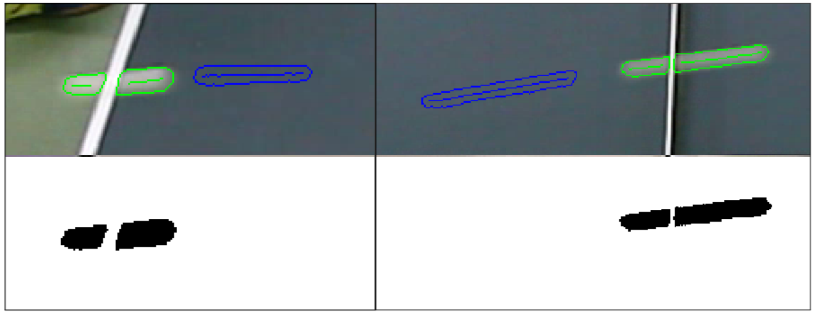

3) Tracker

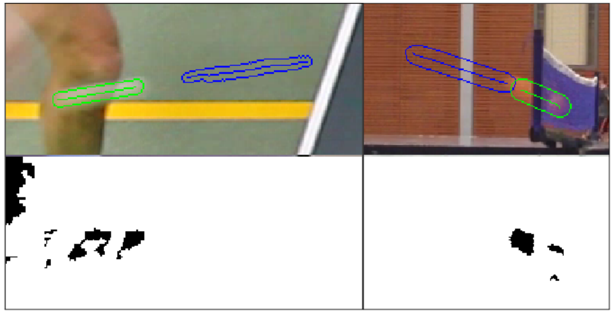

Input FMO Video

FMO Tracking + Reconstruction

Temporal SR - Table tennis

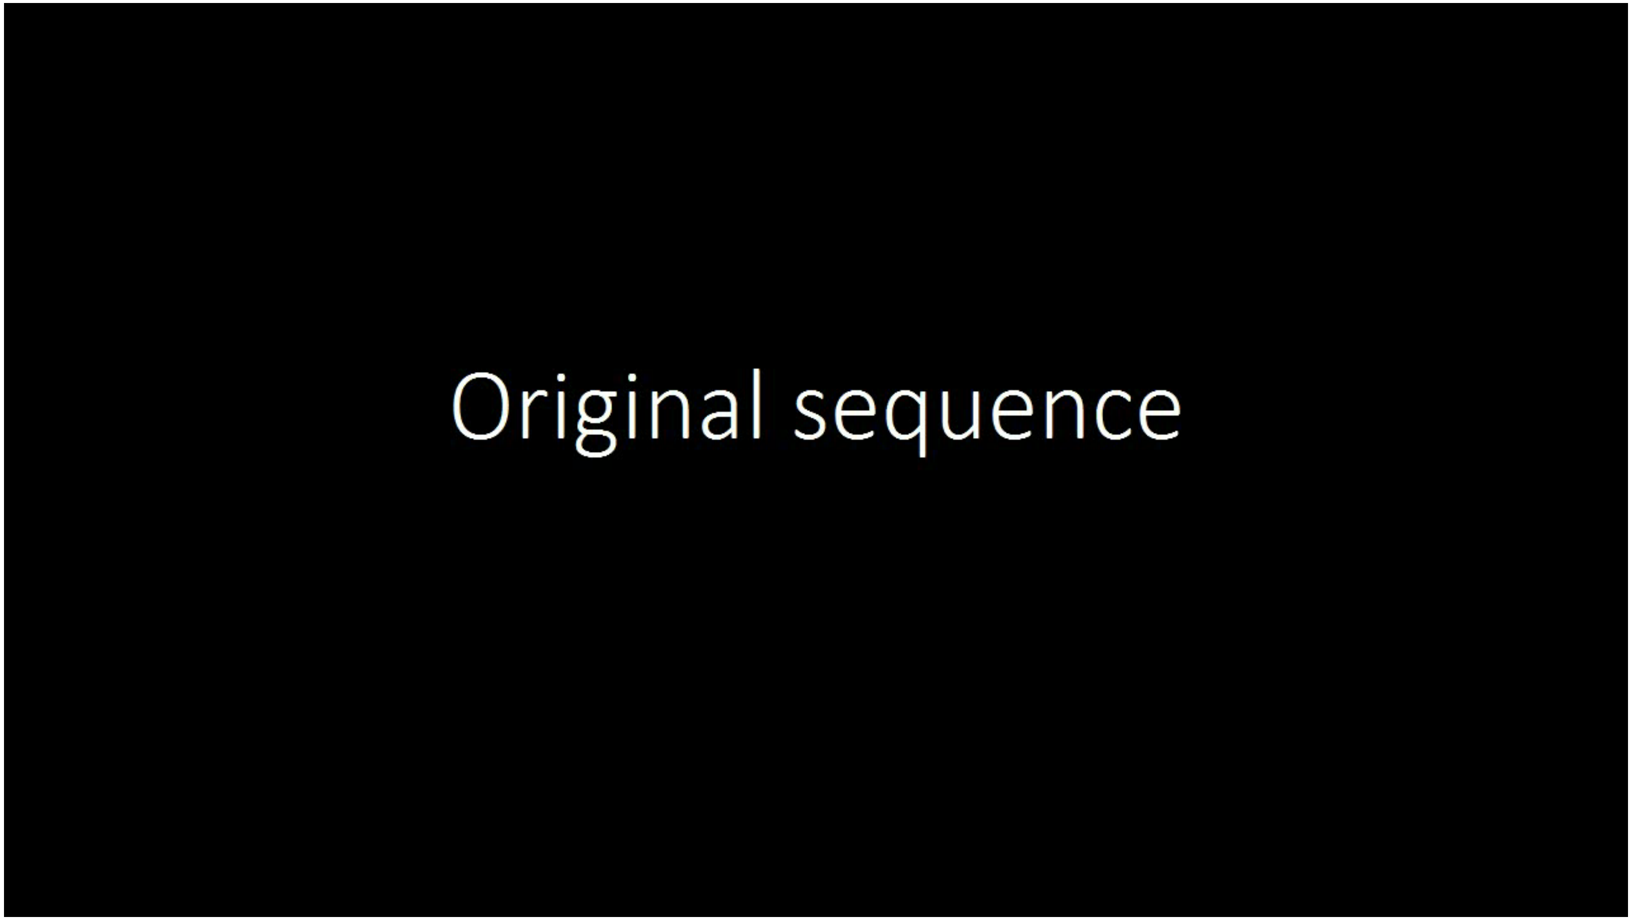

Original sequence

Temporal SR - Darts

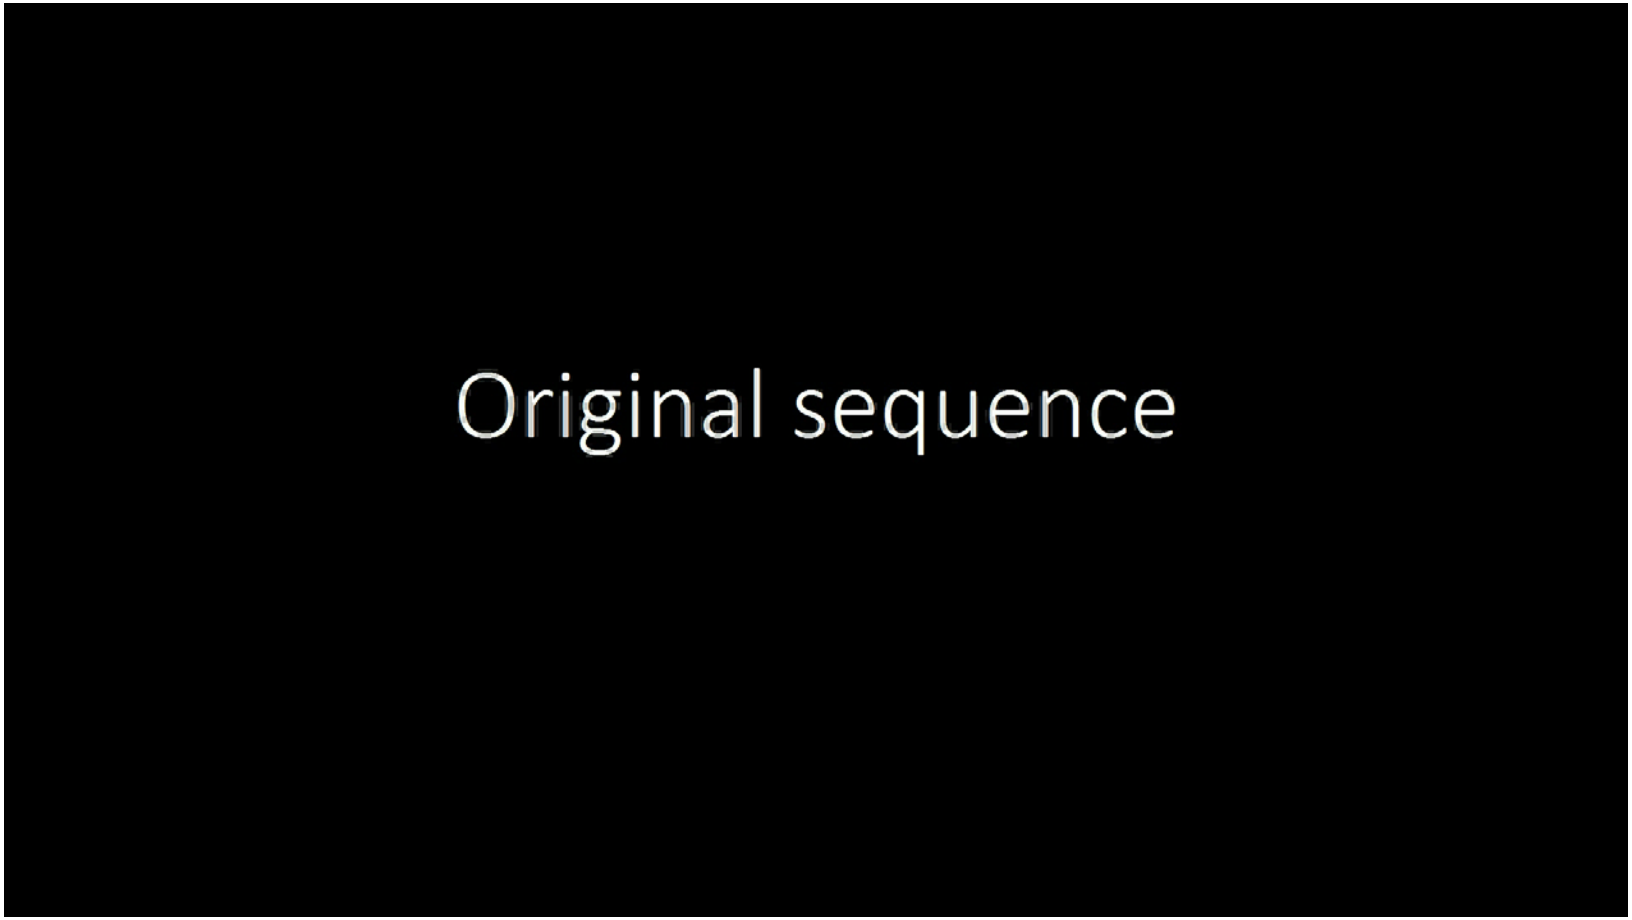

Original sequence

Rotating FMOs

- Simple convolution

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

replaced by space-variant convolution

$$z = (1 - \mathcal{H}m)b + \mathcal{H}f$$

$\mathcal{H}$ is a function of trajectory and angular velocity

f is a 3D object (sphere)

- Gradient descent w.r.t f
- Exhaustive search w.r.t angular velocity

Temporal SR with Rotation

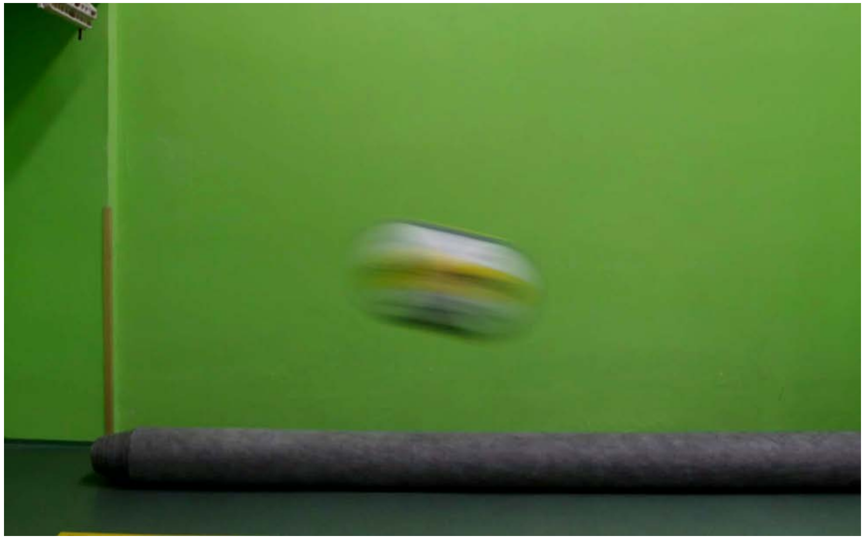

Input video frame

Estimated
high-speed video

Limitations

- Poor collaboration between tracking and restoration
  - Estimated blur and FMO appearance improve tracking
- Unstable FMO appearance estimation
  - Combine multiple video frames
- Track multiple FMOs

Thank You

For Your Attention

Multichannel Model

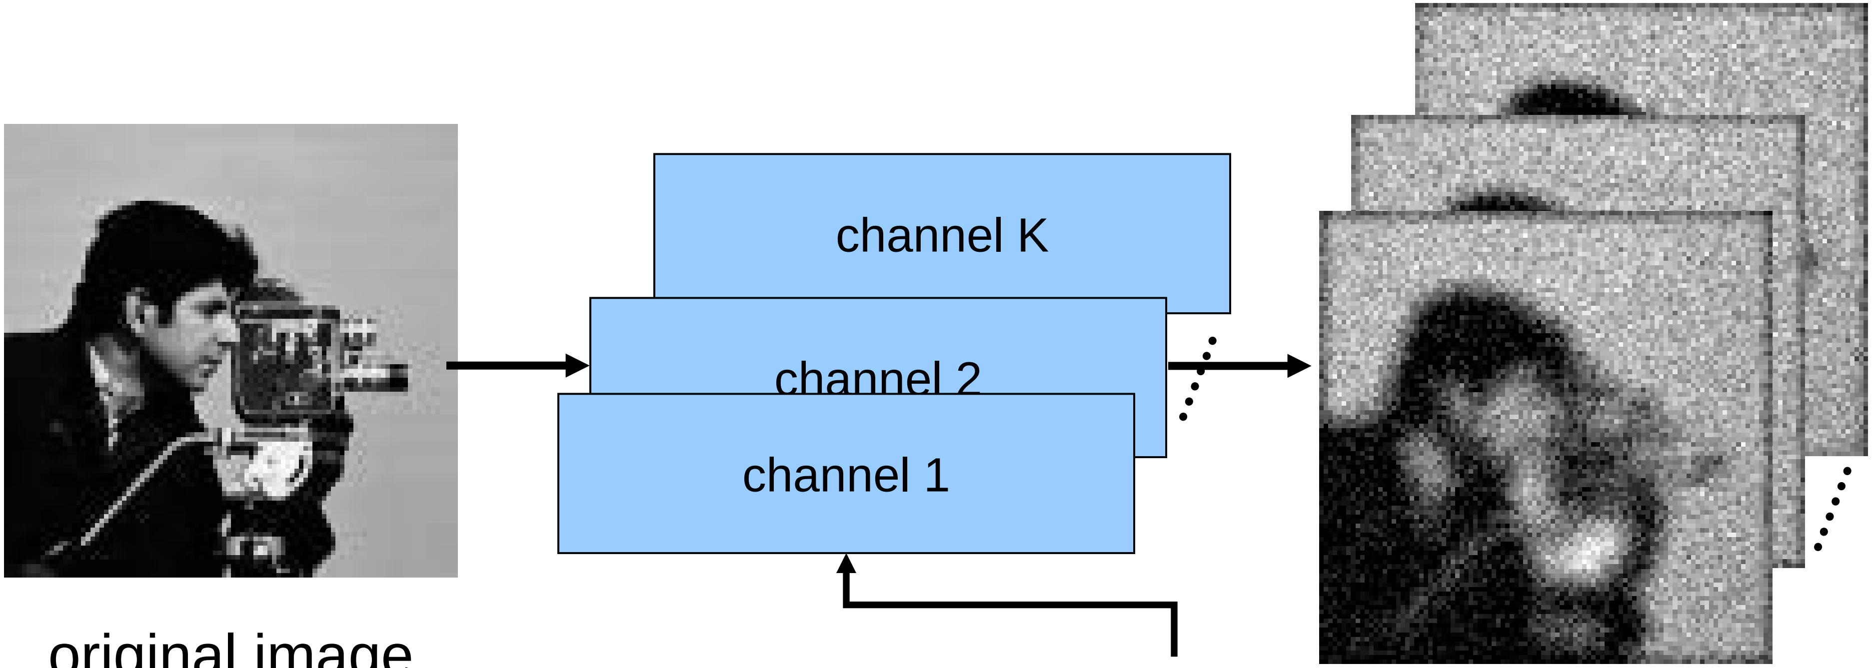

original image

noise

acquired images

$$[u \quad * \quad h_k] \quad + \quad n_k \quad = \; z_k$$

Multichannel Model

- Acquisition model:

$$\begin{aligned} z_1 &= (h_1 * u) + n_1 \\ &\vdots \\ z_K &= (h_K * u) + n_K \end{aligned}$$

- Optimization problem

$$E(u, \{h_k\}) = \underbrace{\frac{1}{2}\sum_k \|(h_k * u) - z_k\|^2}_{\text{Data term}} + \underbrace{\lambda Q(u)}_{\text{Image regularization term}} + \underbrace{\gamma R(\{h_k\})}_{\text{Blur Regularization term}}$$

Data term

Image regularization term

Blur Regularization term

Gaussian noise
L2 norm

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

Multichannel Blur Regularization

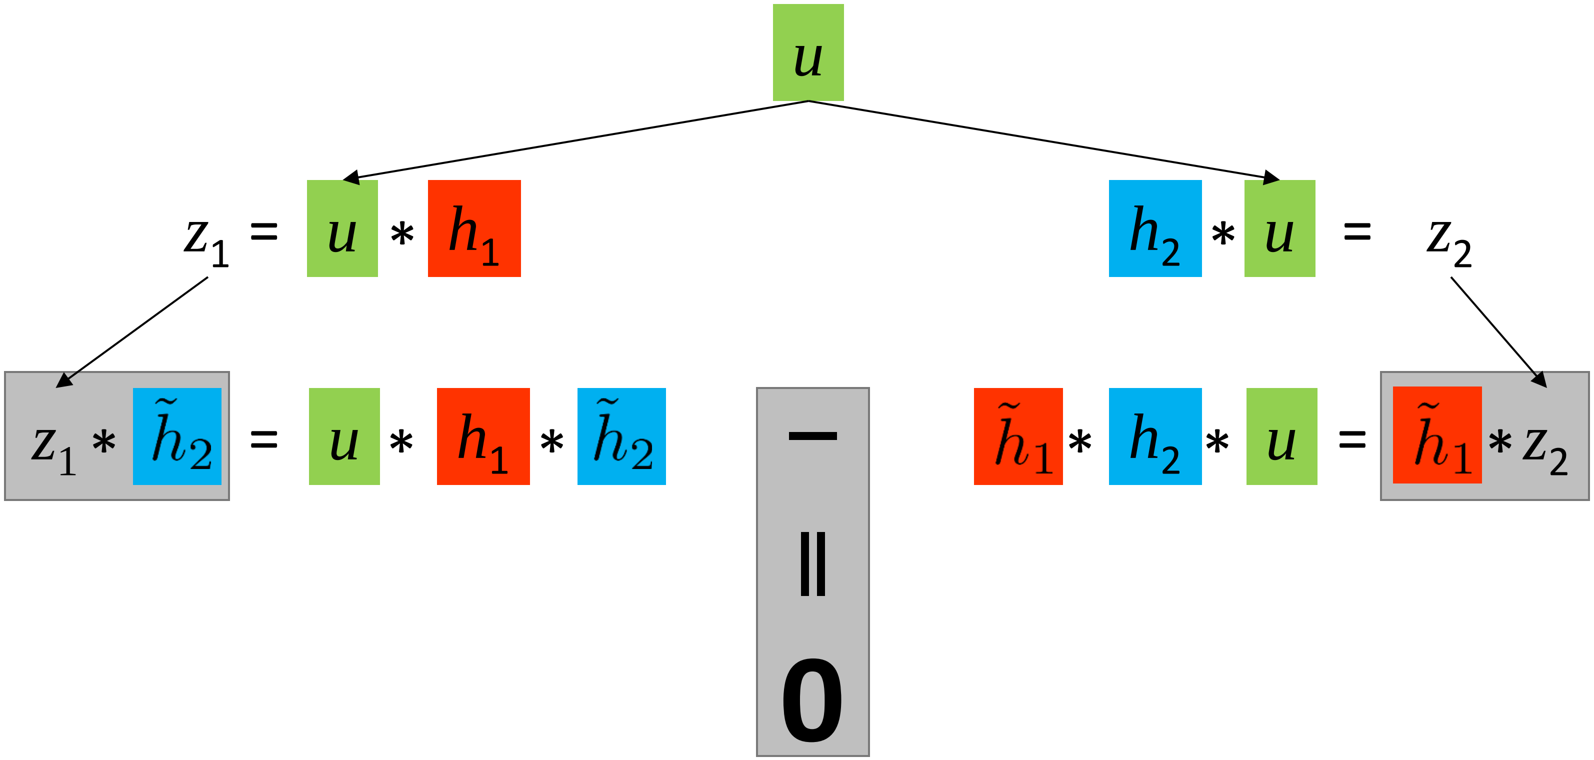

u

$z_1 = u * h_1$

$h_2 * u = z_2$

$z_1 * \tilde{h}_2 = u * h_1 * \tilde{h}_2$

$\tilde{h}_1 * h_2 * u = \tilde{h}_1 * z_2$

$-$

$\parallel$

$\mathbf{0}$

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

Camera motion

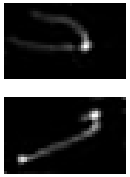

Astronomical images

Degraded images

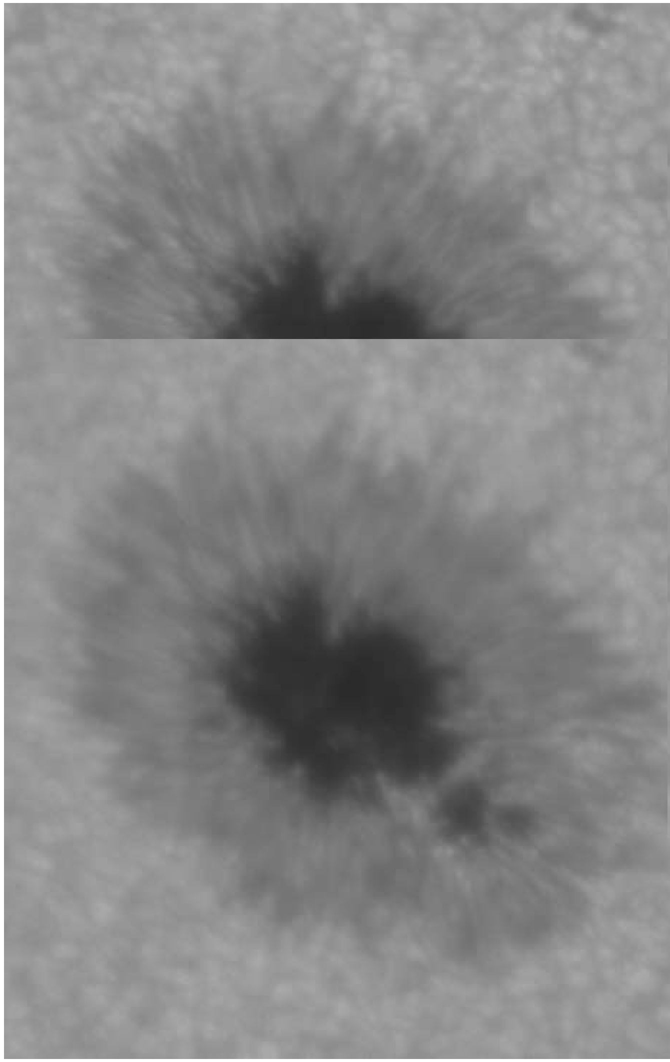

Reconstructed image

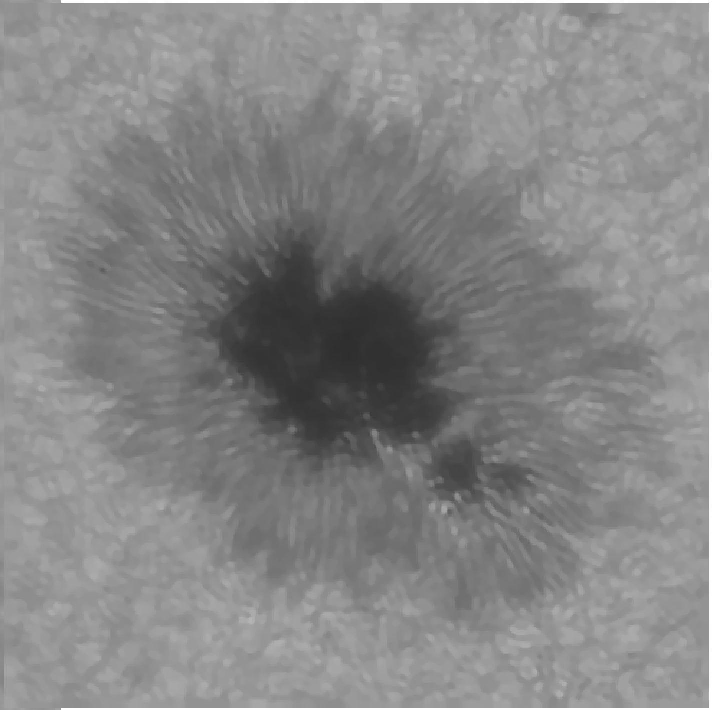

PSF estimation

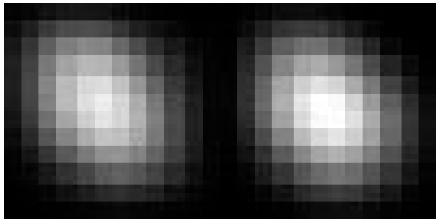

MC Model with Decimation

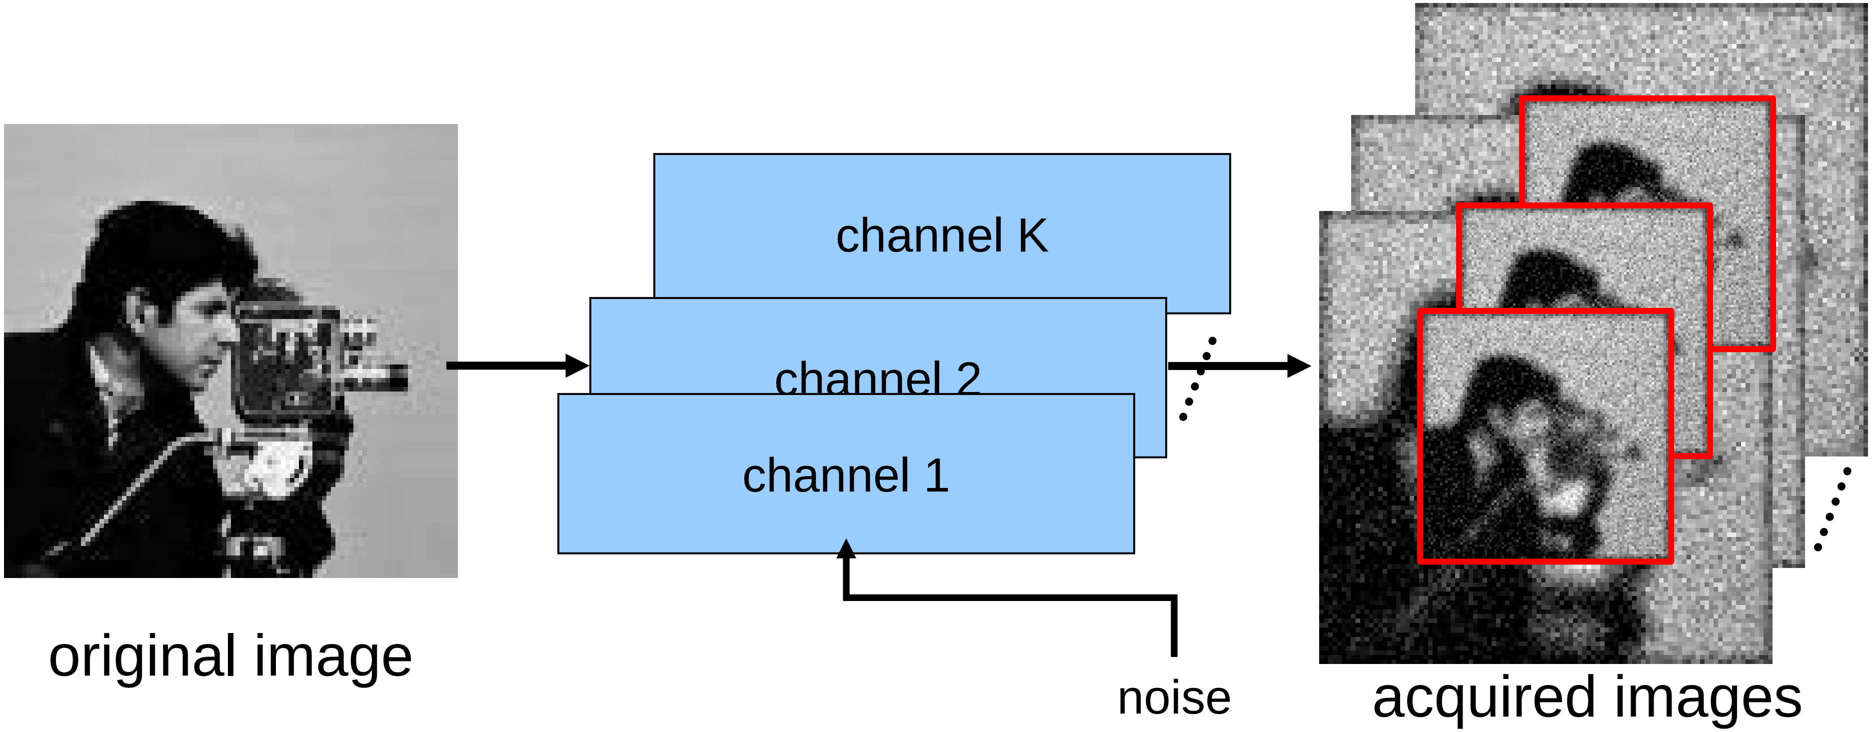

$$D\,[u \quad * \quad h_k] \quad + \quad n_k \quad = \quad z_k$$

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

Robustness to misalignment

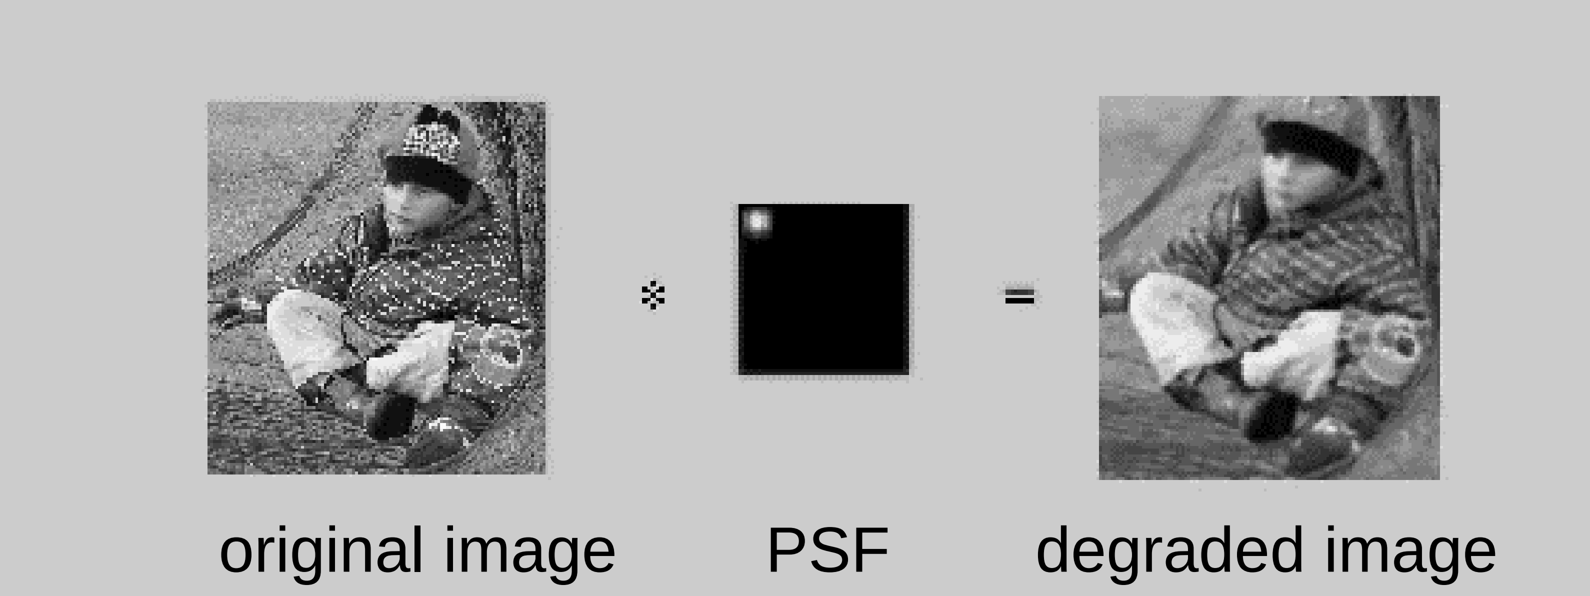

$$D(u * h_k)[\tau_k(x)] + n_k(x) = z_k(x)$$

$$D(u * g_k)(x) + n_k(x) = z_k(x)$$

Super-resolution

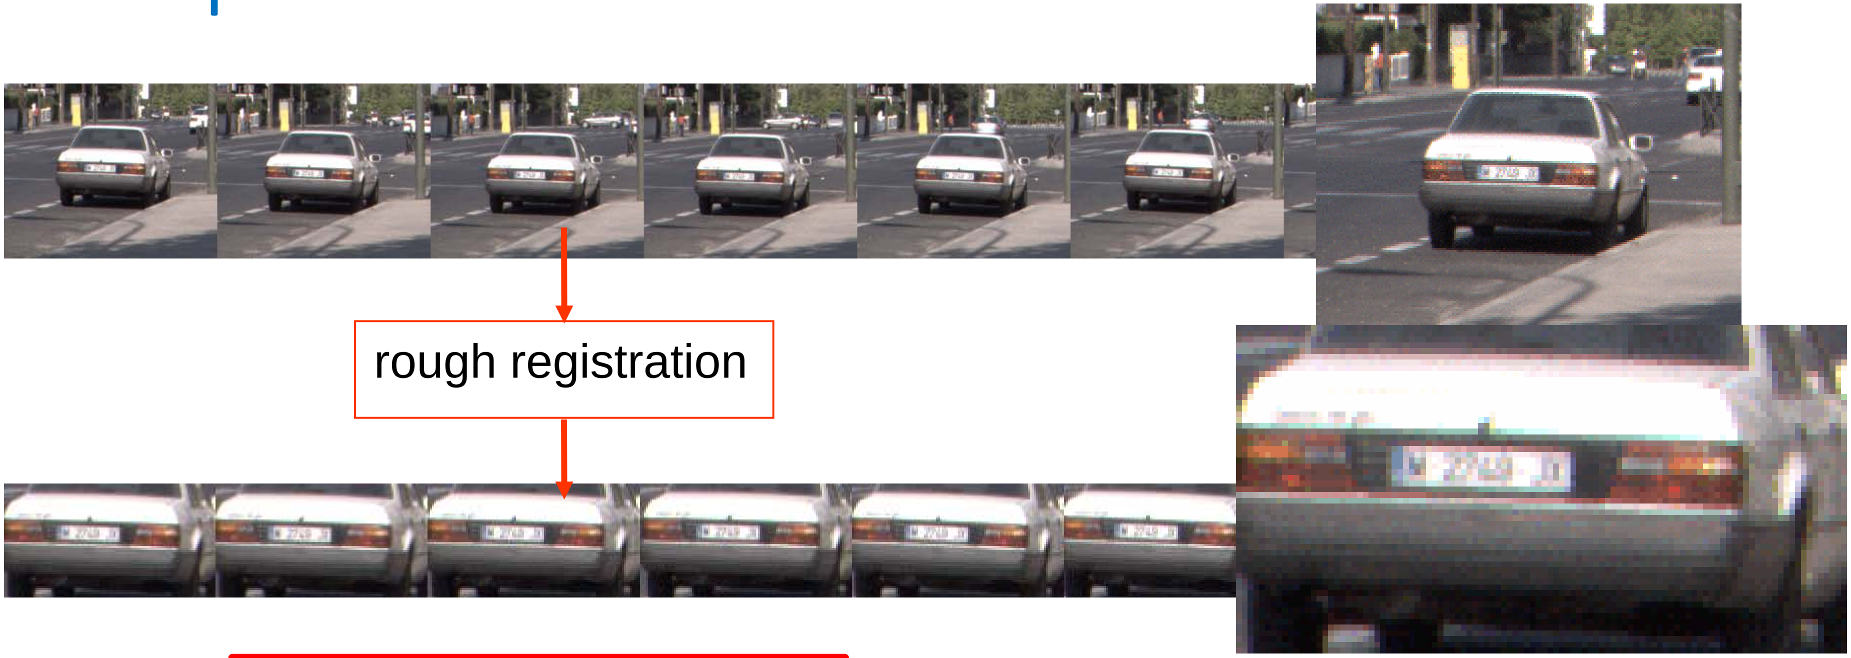

rough registration

Superresolved image (2x)

Optical zoom (ground truth)

original

8 images

SR 2x

SR 3x

Video Super-resolution

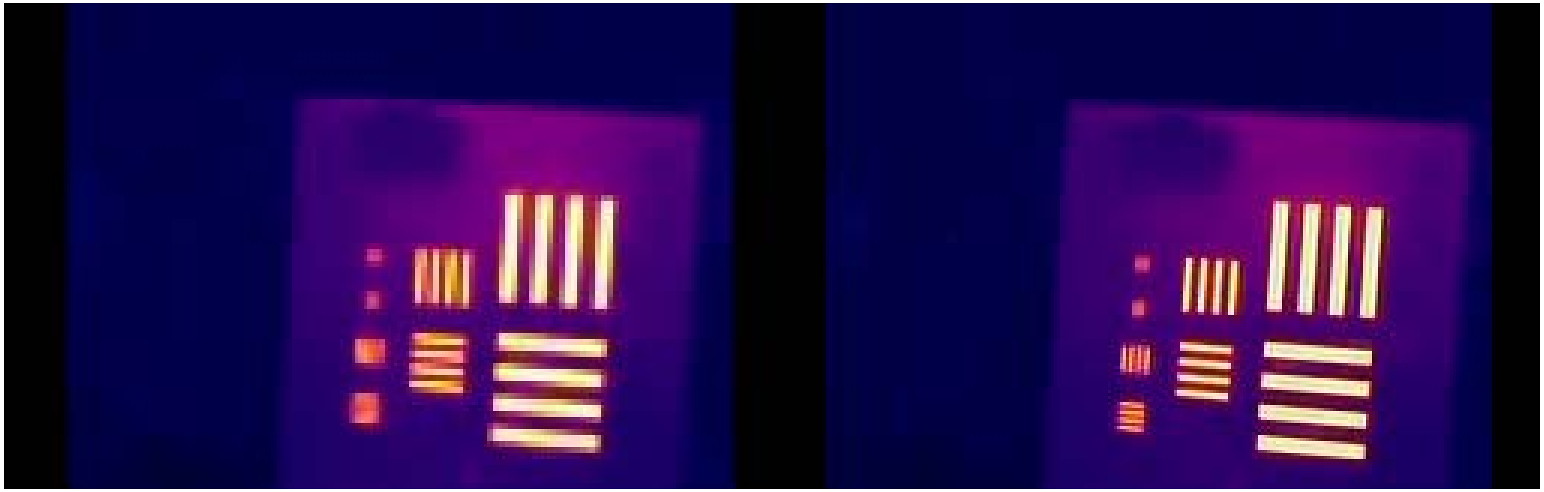

Input video

Super-resolution

Embedded Systems

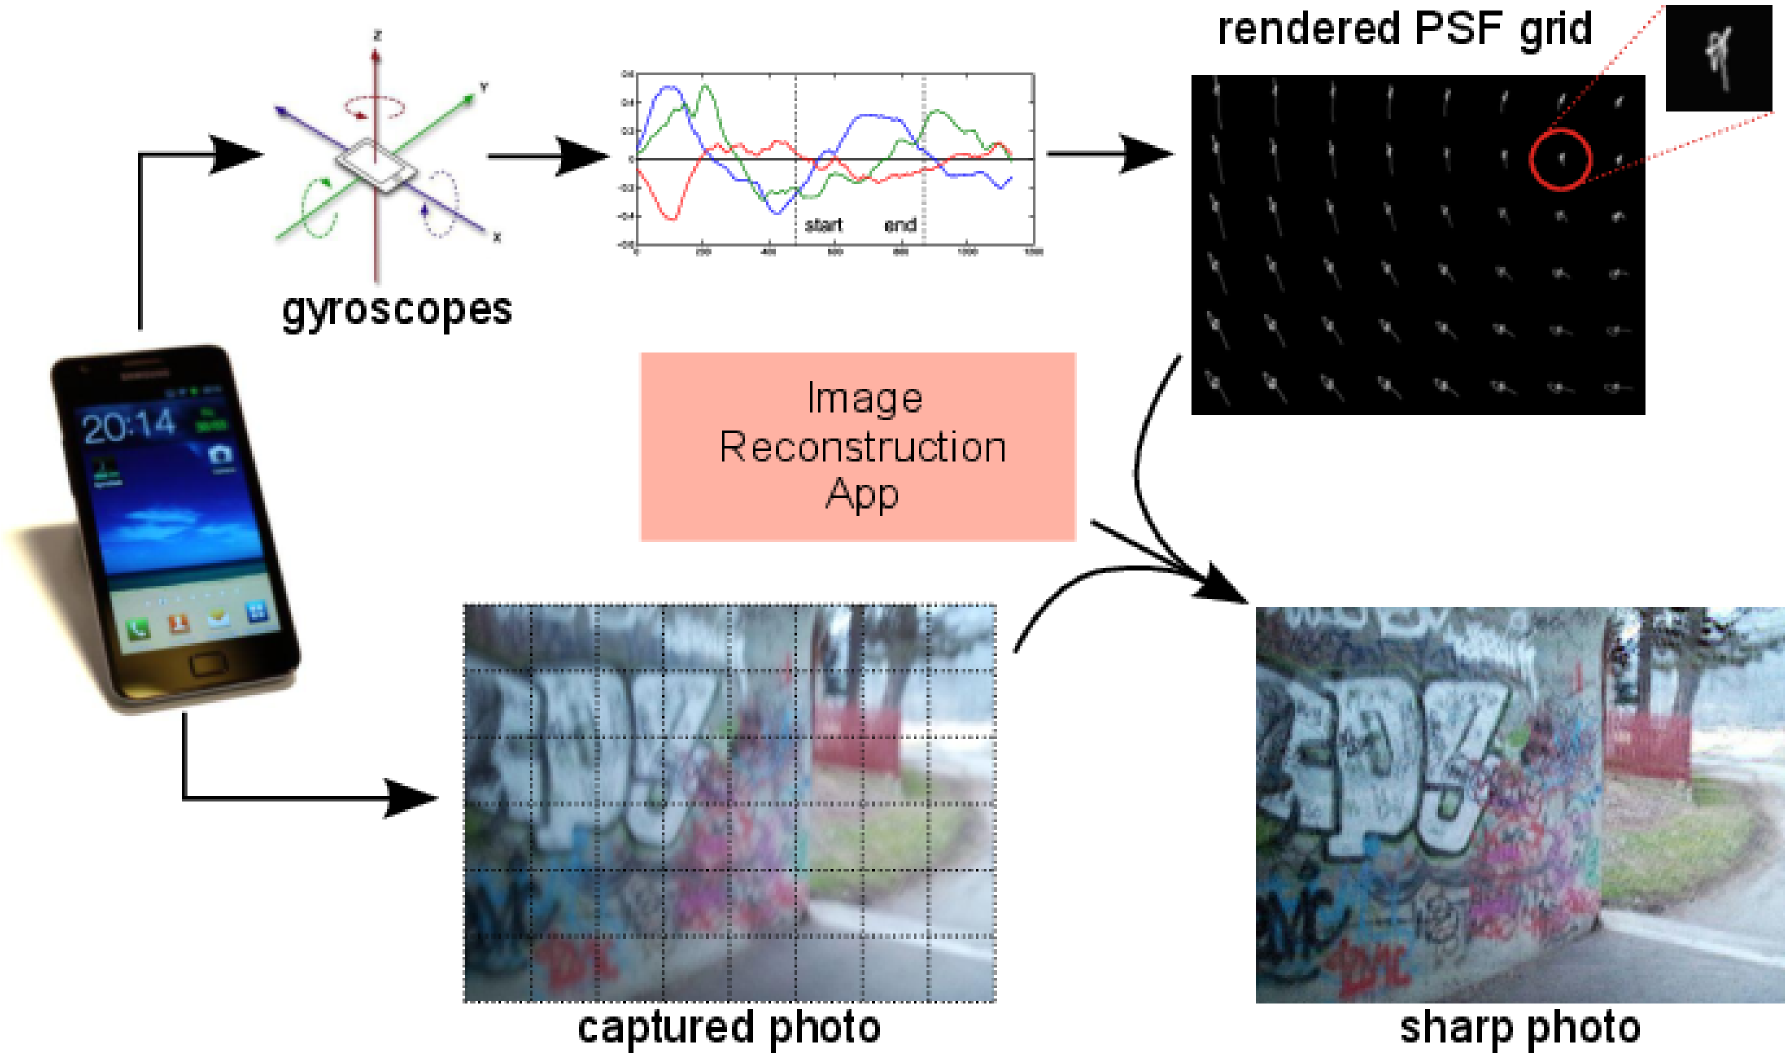

rendered PSF grid

gyroscopes

Image
Reconstruction
App

captured photo

sharp photo

Acquired blurred image

Blur estimation

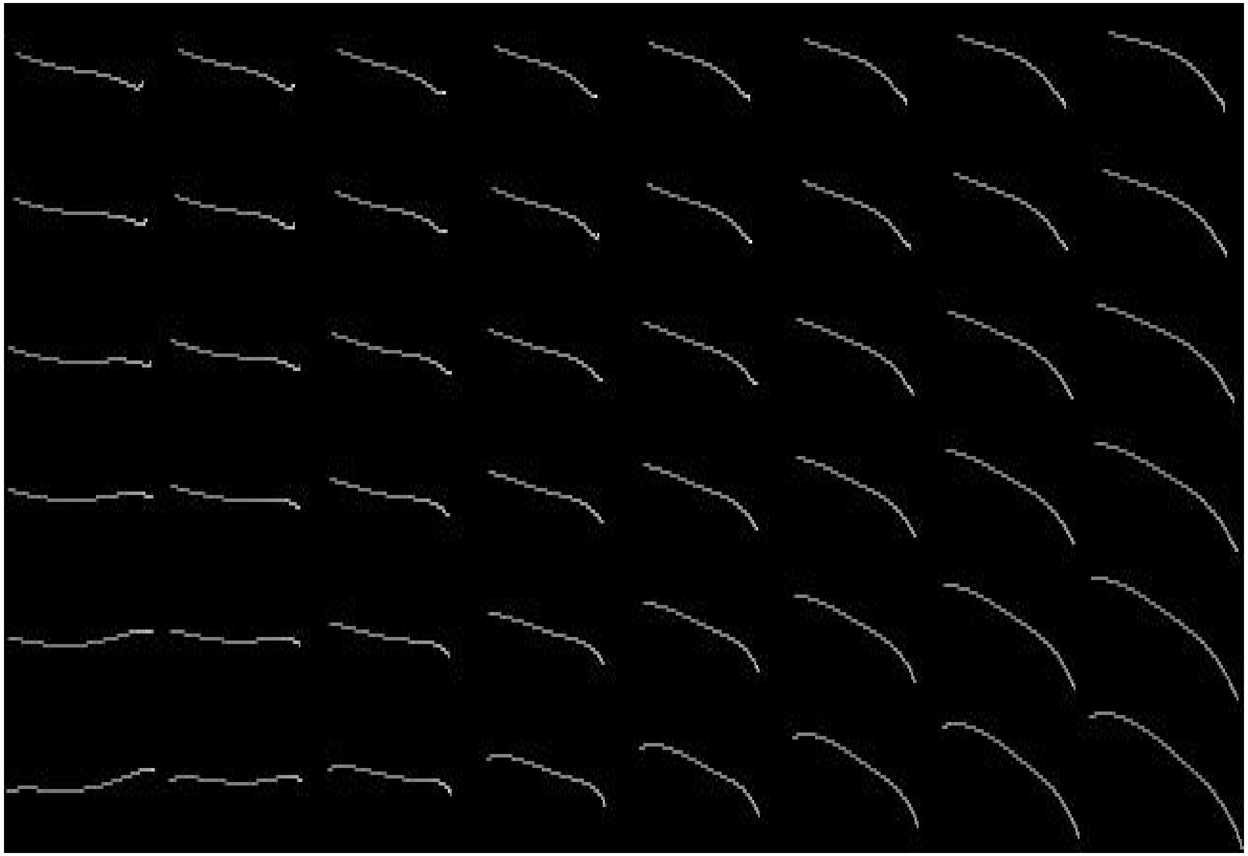

Patch-wise Deconvolution

Super-resolution

JPEG

SR

Embedded Super-Resolution

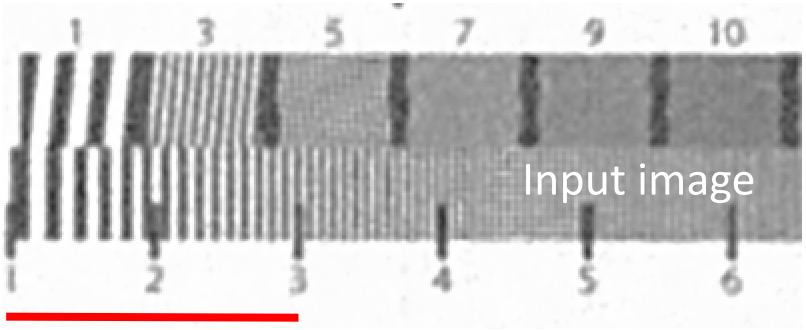

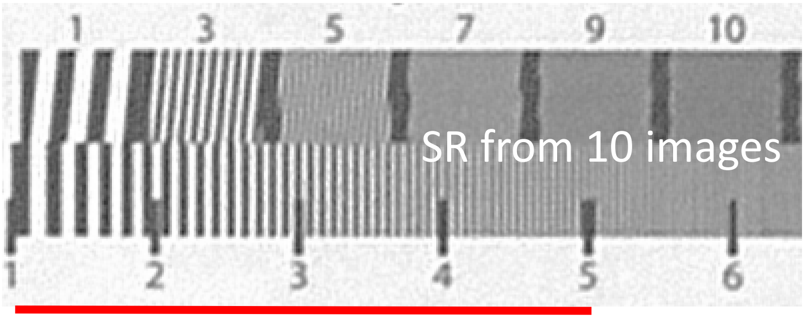

Patch-wise restoration

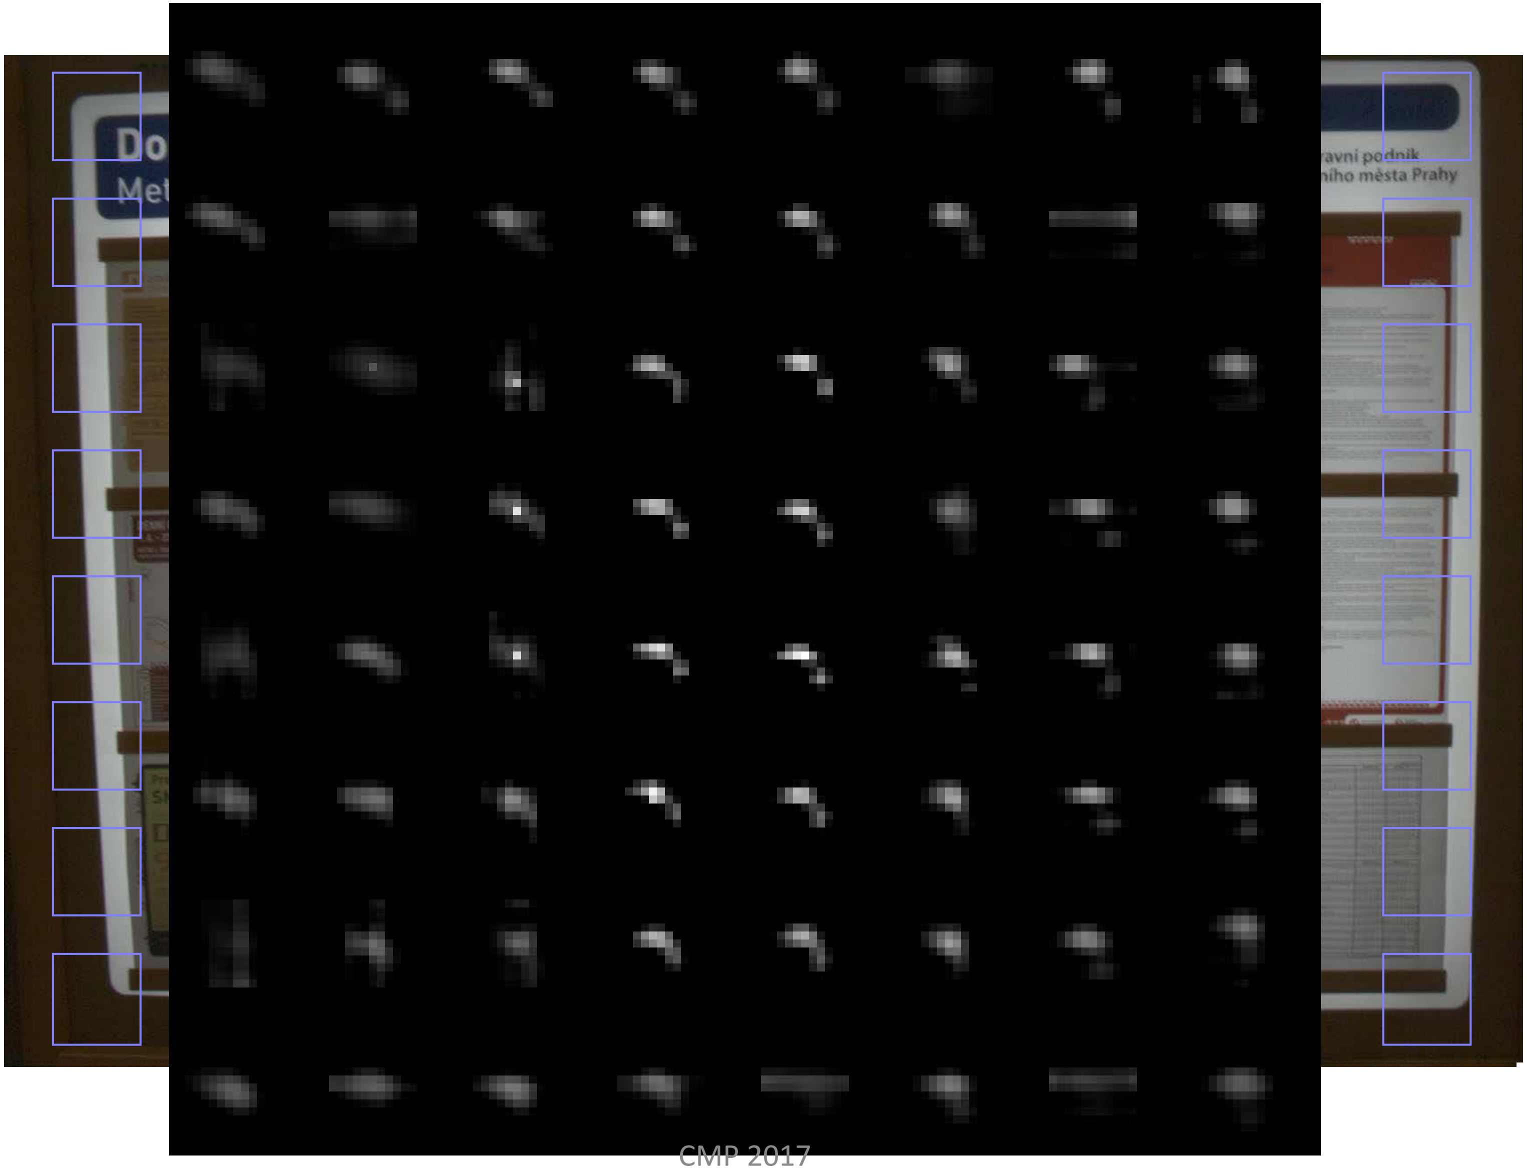

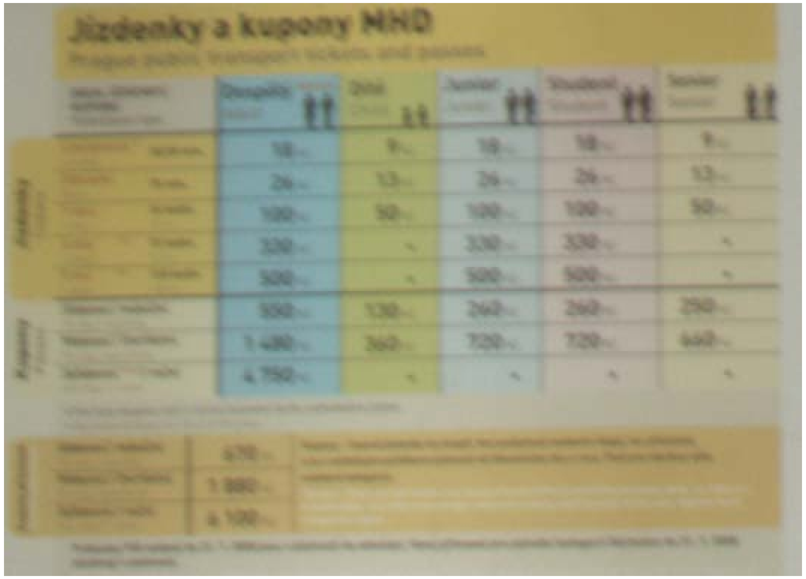

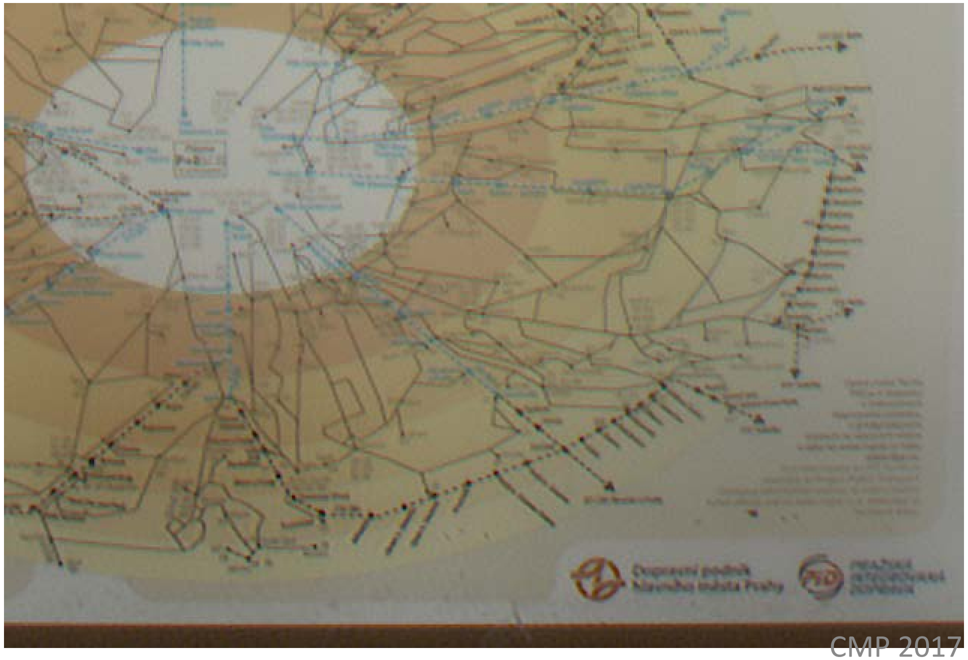

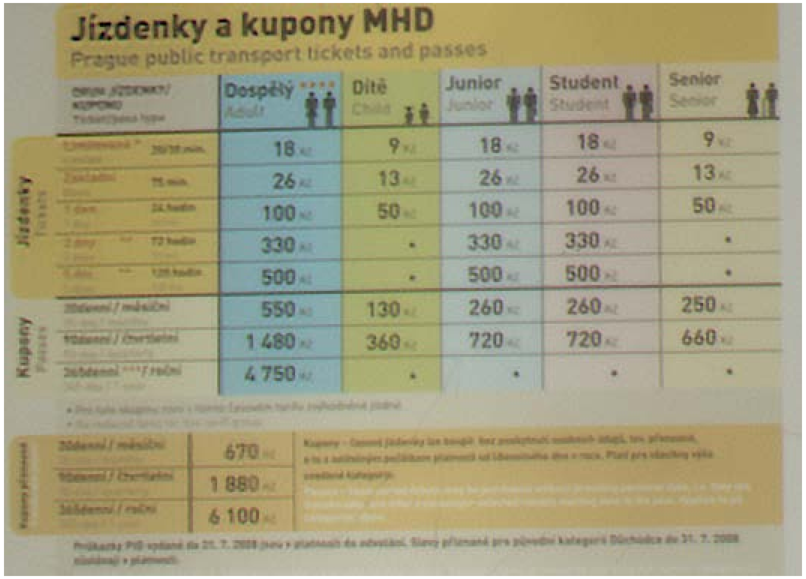

Jízdenky a kupony MHD

Prague public transport tickets and passes

		Dospělý Adult	Dítě Child	Junior Junior	Student Student	Senior Senior
Jízdenky	Jednotlivé * 30/35 min.	18 Kč	9 Kč	18 Kč	18 Kč	9 Kč
	Základní 75 min.	26 Kč	13 Kč	26 Kč	26 Kč	13 Kč
	1 den 24 hodin	100 Kč	50 Kč	100 Kč	100 Kč	50 Kč
	3 dny ** 72 hodin	330 Kč	•	330 Kč	330 Kč	•
	5 dní ** 120 hodin	500 Kč	•	500 Kč	500 Kč	•
Kupony	30denní / měsíční	550 Kč	130 Kč	260 Kč	260 Kč	250 Kč
	90denní / čtvrtletní	1 480 Kč	360 Kč	720 Kč	720 Kč	660 Kč
	365denní ***/ roční	4 750 Kč	•	•	•	•

Kupony přenosné	30denní / měsíční	670 Kč
	90denní / čtvrtletní	1 880 Kč
	365denní / roční	6 100 Kč

Regularization

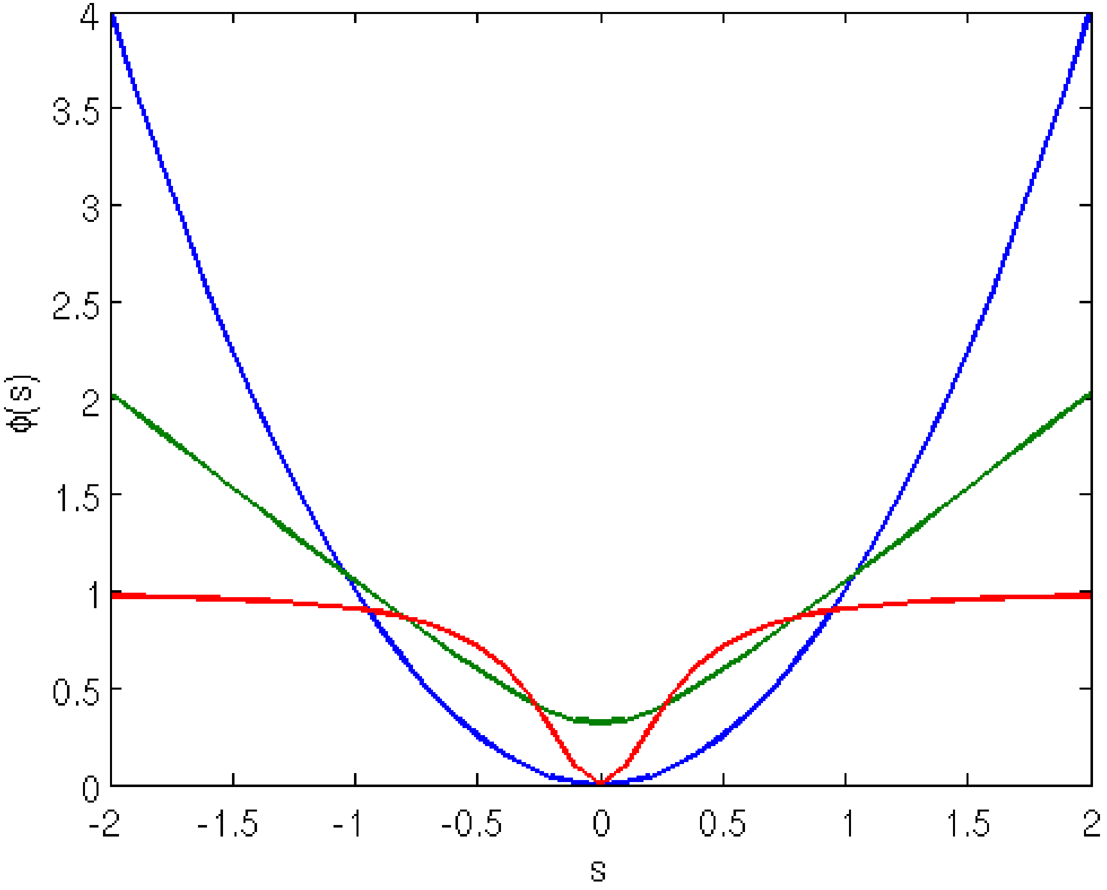

Tikhonov . . . $|s|^2$

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

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

Adjusting priors

u

$$Q(u) = \sum_i |\nabla u_i|^2$$

$$\sum_i |\nabla u_i|^1$$

$$\sum_i |\nabla u_i|^{0.5}$$

- 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.