CS6640 Computational Photography

16. Camera shake removal

Approaches to shake deblurring

- **Measure shake vs. fully blind approach**
- **Estimate kernel and deconvolve vs. full-image estimation**
- **In this lecture:**

  BenEzra & Nayar 2004: measured, direct deconvolve

  Fergus et al. 2006: blind kernel estimation

  Shan et al. 2008: blind, full-image estimation

  Joshi et al. 2010: measured, semi-blind kernel estimation

Ben-Ezra & Nayar

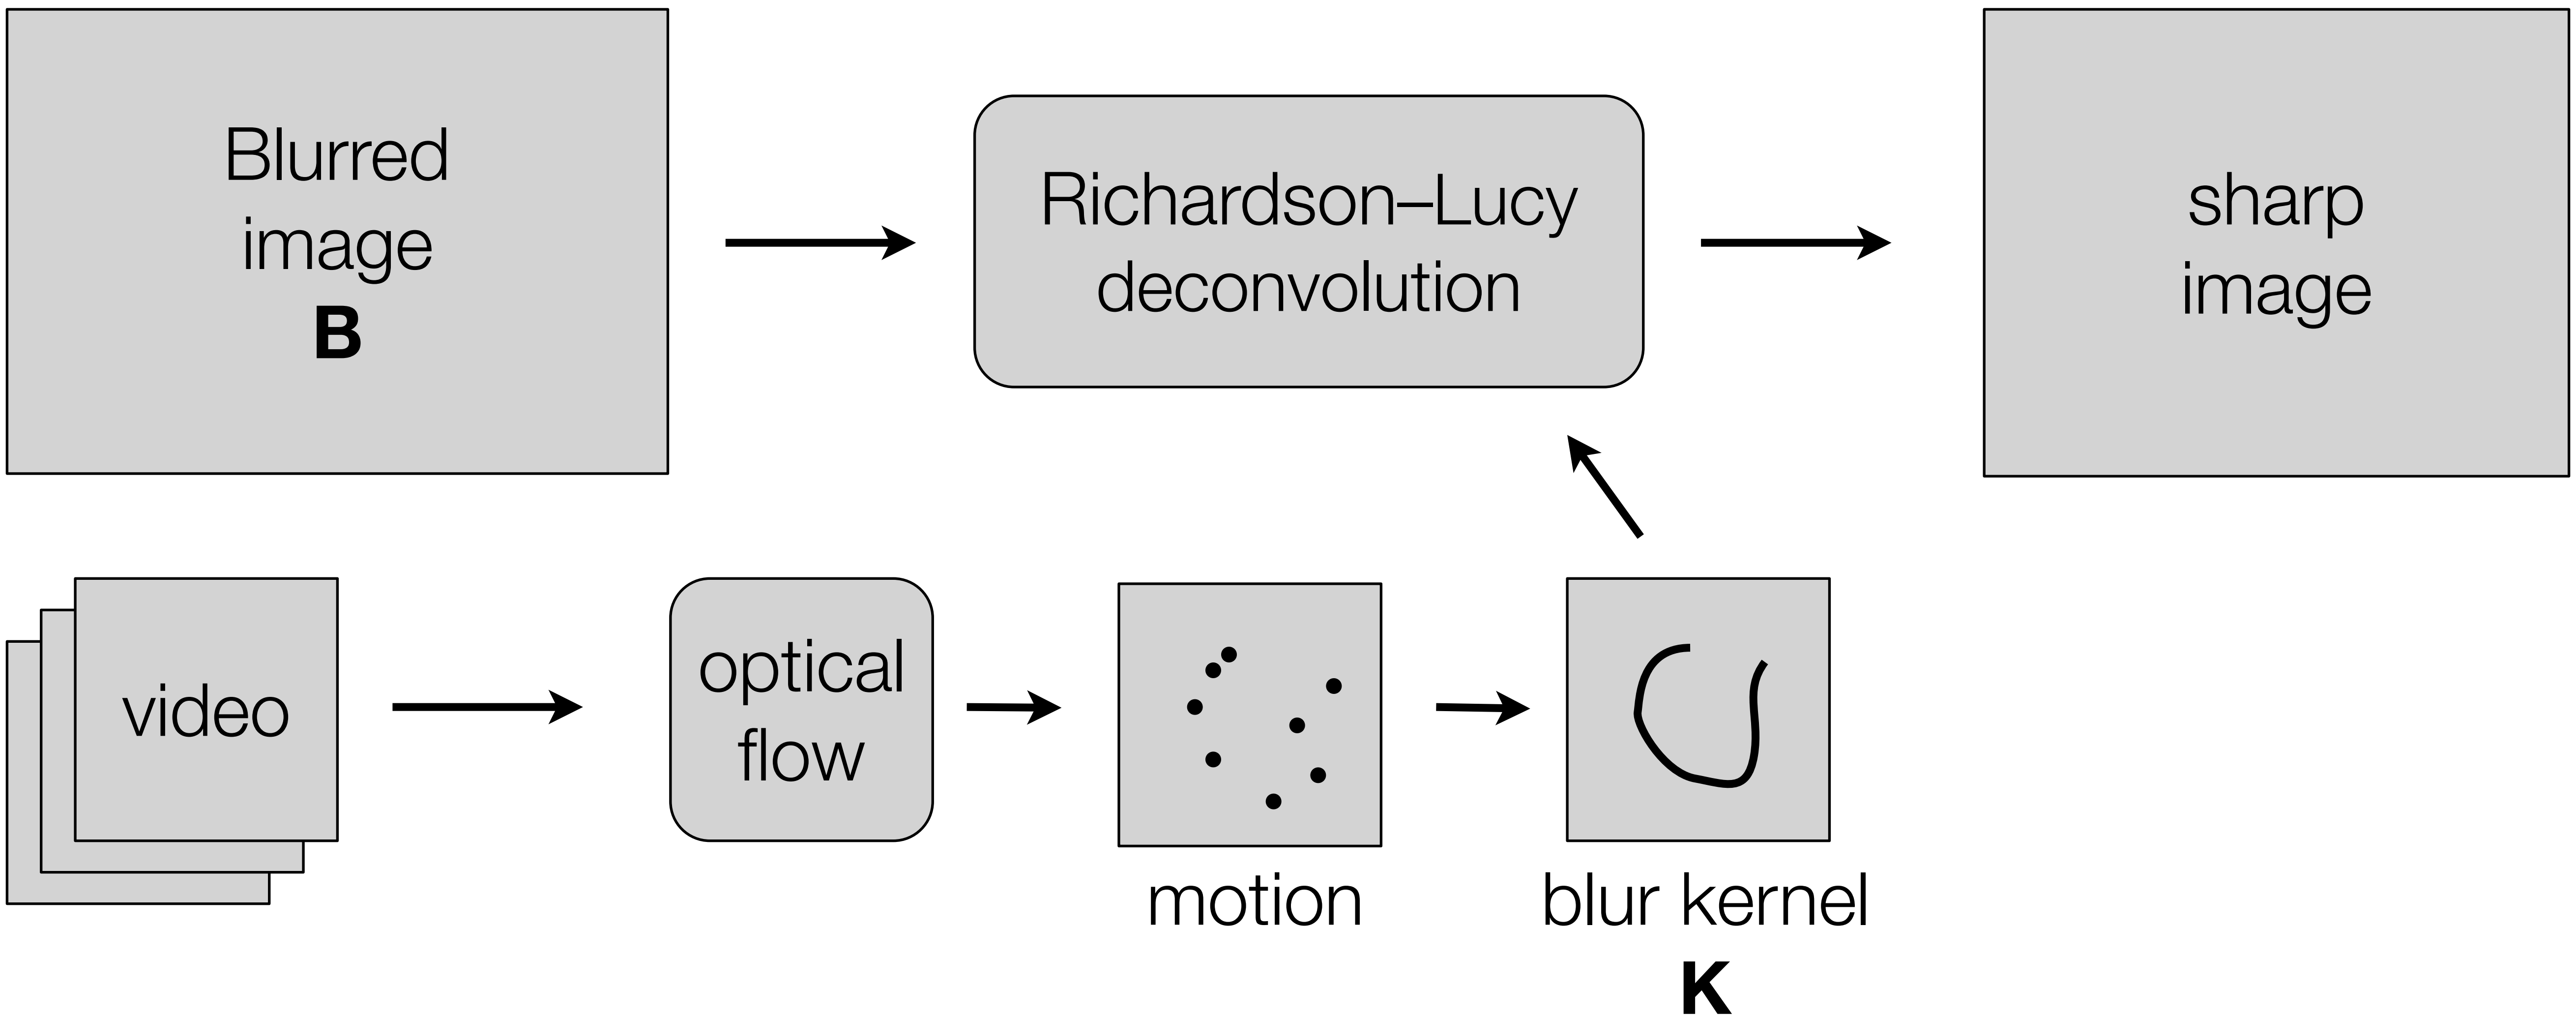

2-camera rig

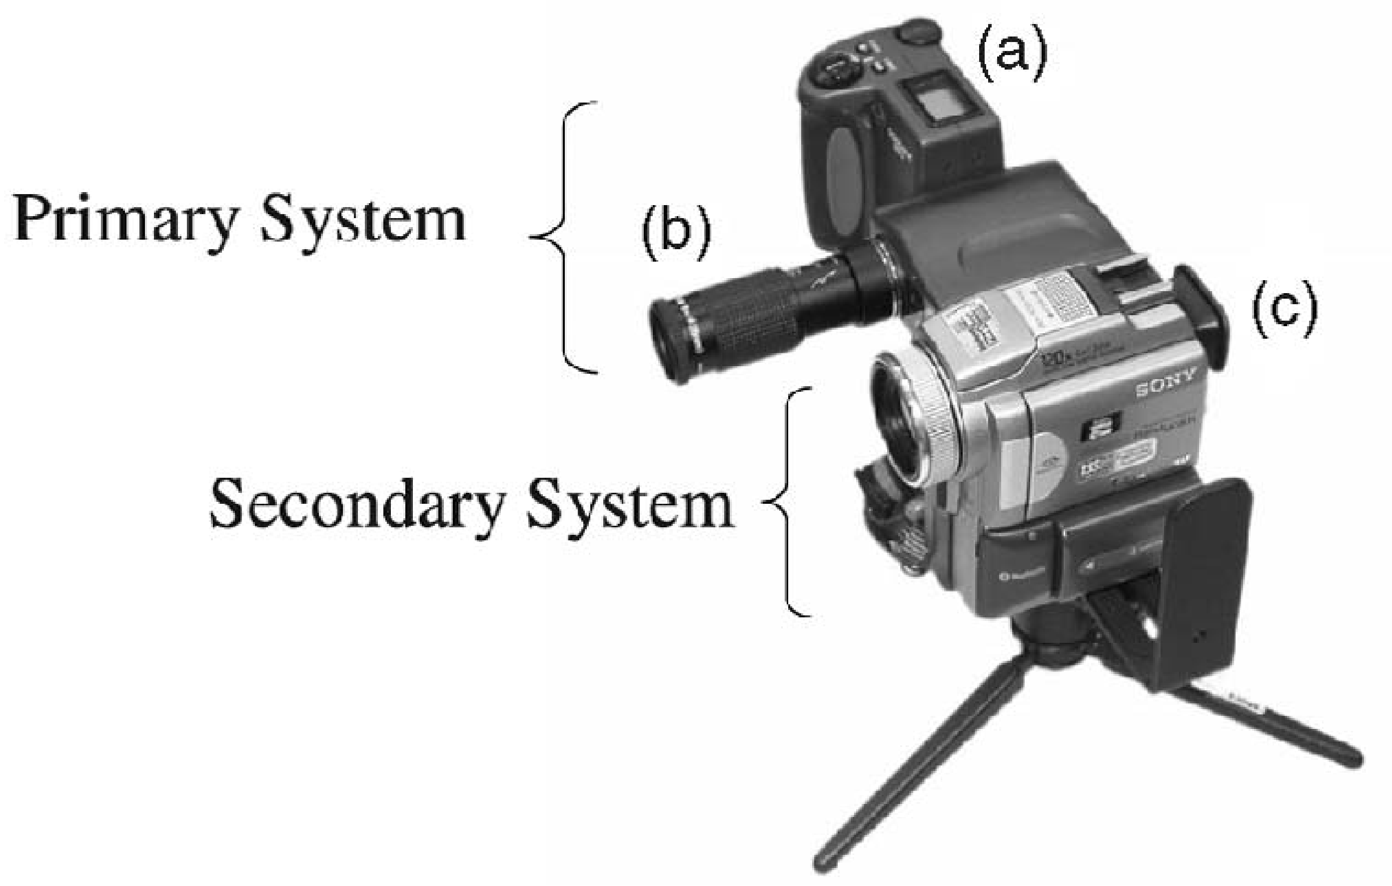

[Ben-Ezra & Nayar 2004]

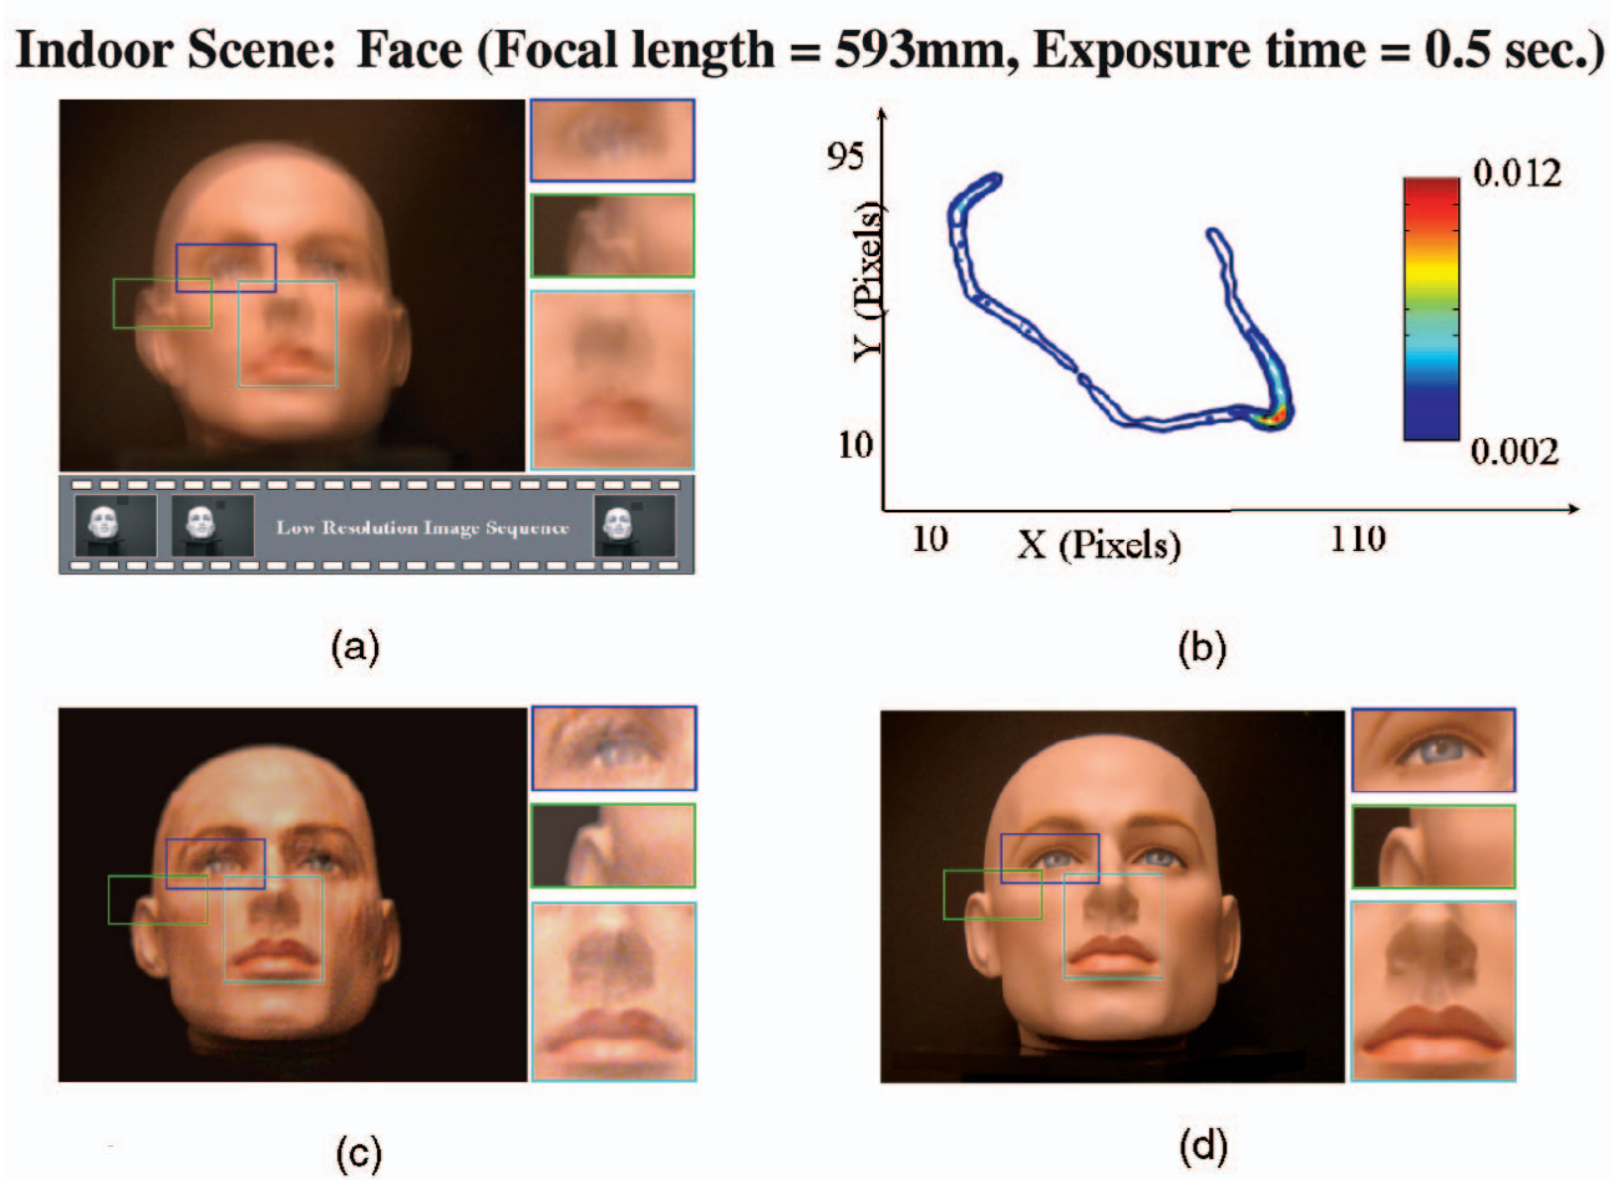

[Ben-Ezra & Nayar 2004]

Fig. 10. Experimental results for indoor face scene. (a) Input images, including the motion blurred image from the primary detector and a sequence of low-resolution frames from the secondary detector. (b) The computed PSF. Notice the complexity of its path and its energy distribution. (c) The deblurring result. The magnified windows show details. (d) Ground truth image that was captured without motion blur using a tripod.

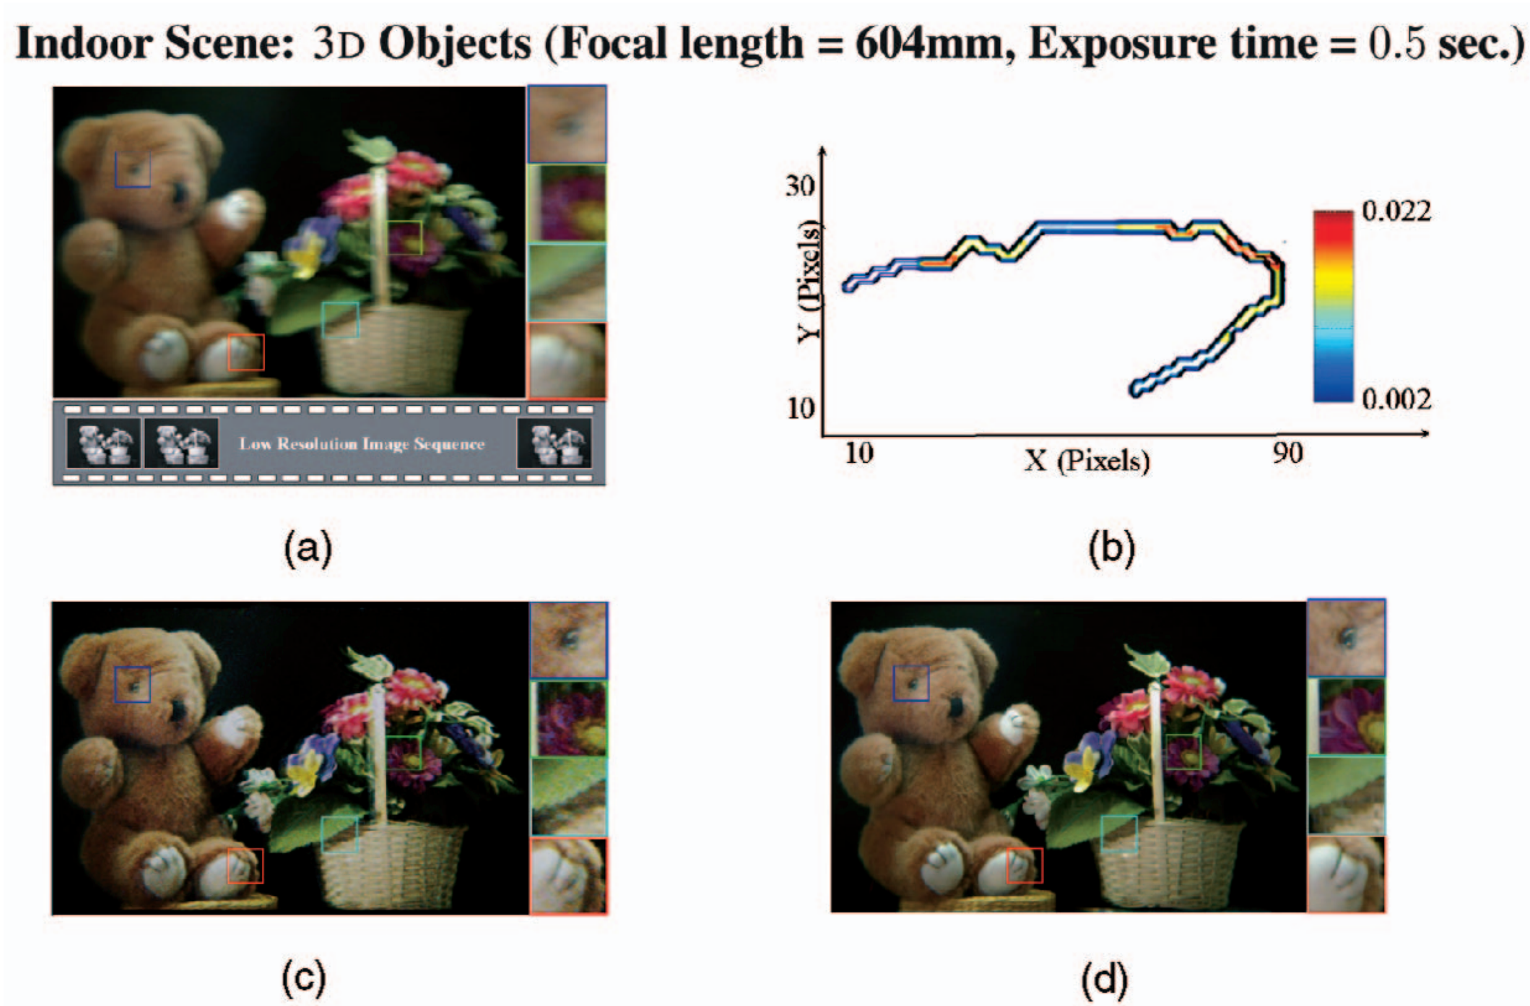

[Ben-Ezra & Nayar 2004]

Fig. 9. Experimental results for indoor 3D objects scene. (a) Input images, including the motion blurred image from the primary detector and a sequence of low-resolution frames from the secondary detector. (b) The computed PSF. Notice the complexity of its path and its energy distribution. (c) The deblurring result. The magnified windows show details. (d) Ground truth image that was captured without motion blur using a tripod.

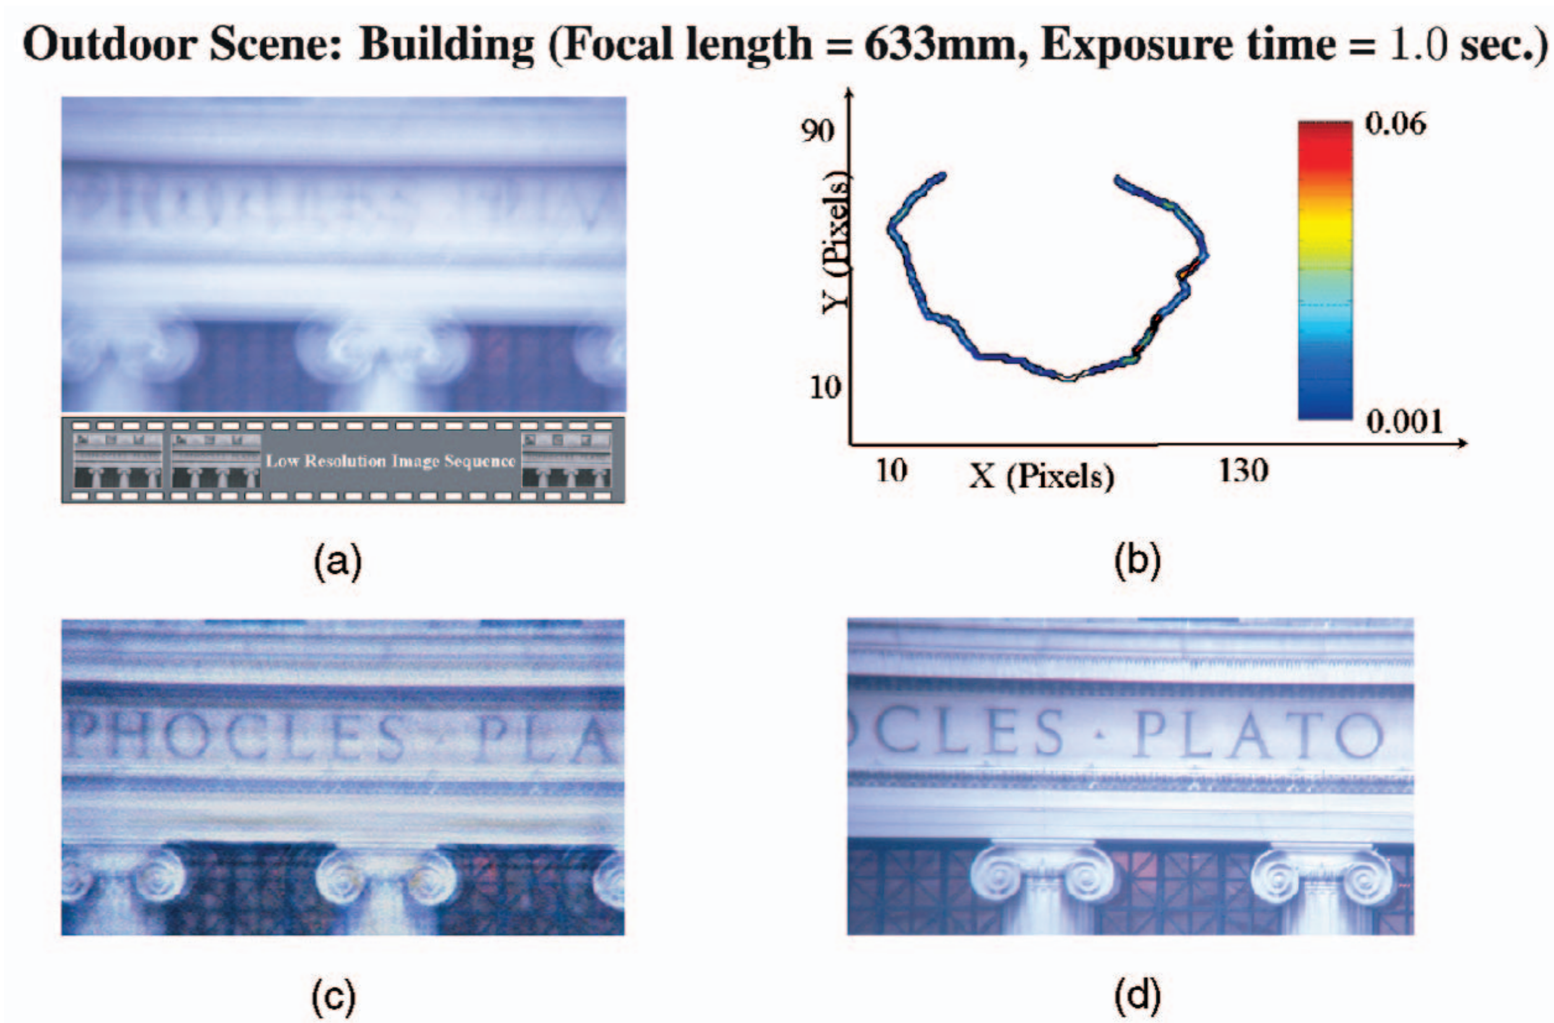

[Ben-Ezra & Nayar 2004]

Fig. 11. Experimental results for outdoor building scene. (a) Input images, including the motion blurred image from the primary detector and a sequence of low-resolution frames from the secondary detector. (b) The computed PSF. Notice the complexity of its paths and its energy distribution. (c) The deblurring result. Notice the clarity of the text. (d) Ground truth image that were captured without motion blur using a tripod.

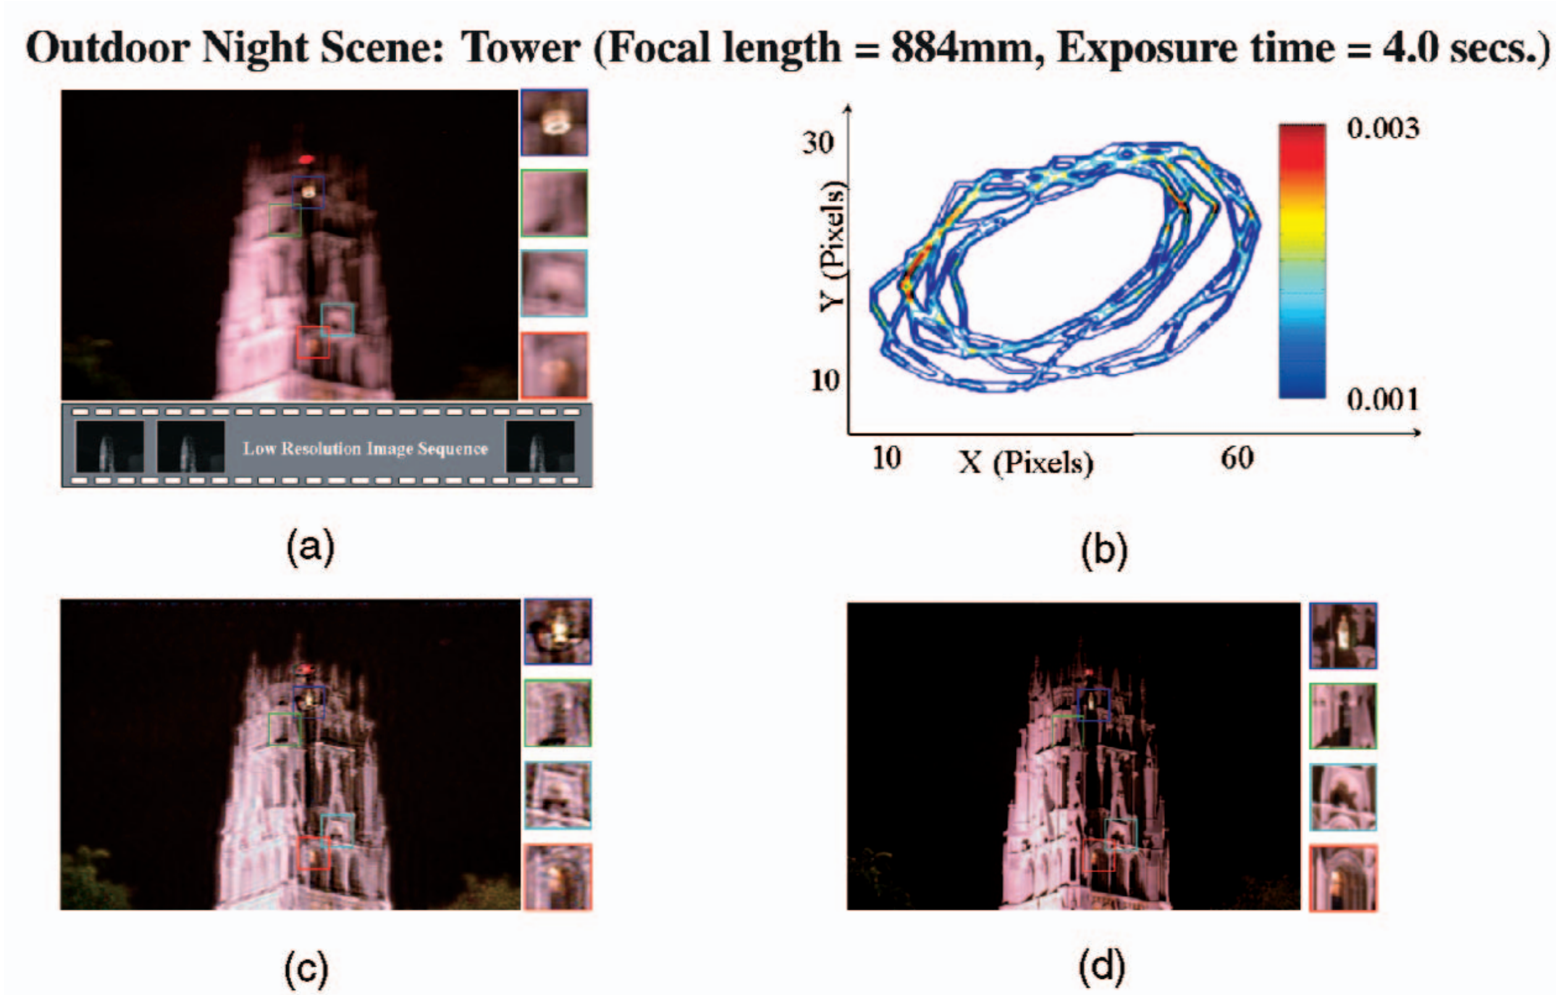

[Ben-Ezra & Nayar 2004]

Fig. 12. Experimental results for outdoor tower scene. (a) Input images, including the motion blurred image from the primary detector and a sequence of low-resolution frames from the secondary detector. (b) The computed PSF. Notice the complexity of its path and its energy distribution. (c) The deblurring result. (d) Ground truth image that was captured without motion blur using a tripod.

Fergus et al.

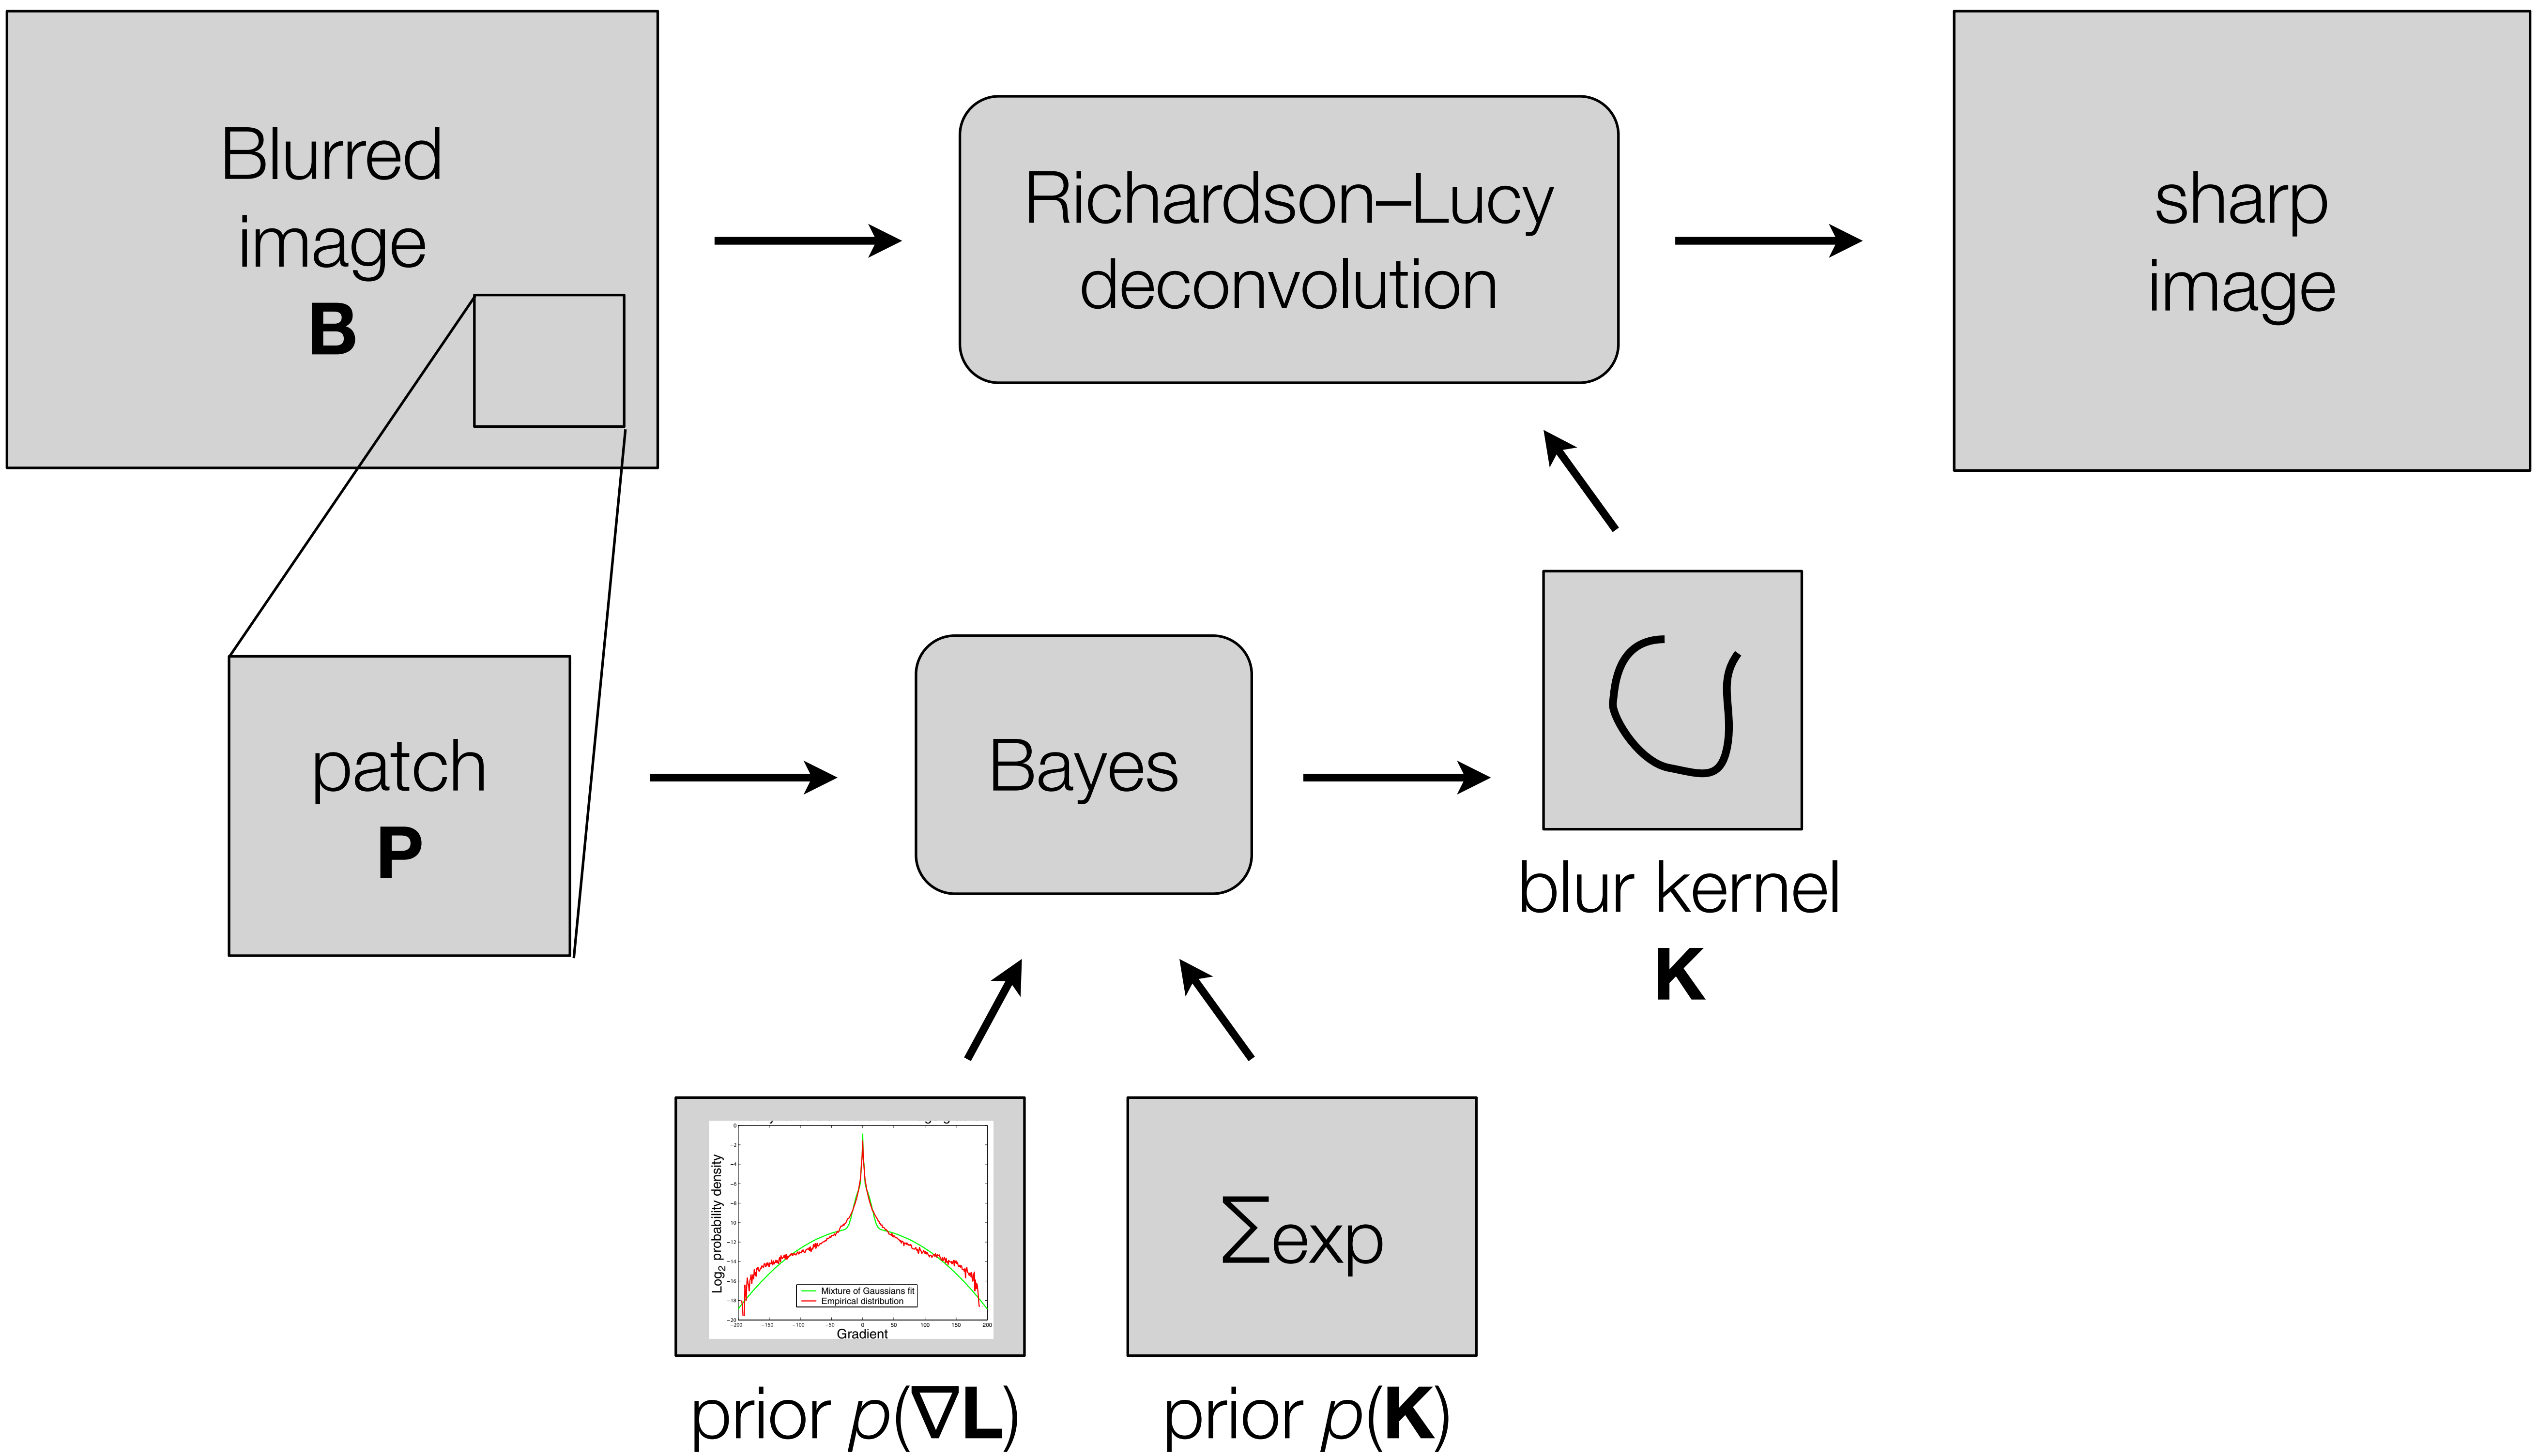

Blurred image
B
Richardson–Lucy deconvolution
sharp image
patch
P
Bayes
blur kernel
K
Log2 probability density
Mixture of Gaussians fit
Empirical distribution
Gradient
Σexp
prior p(∇L)
prior p(K)

Argument for spatial invariace

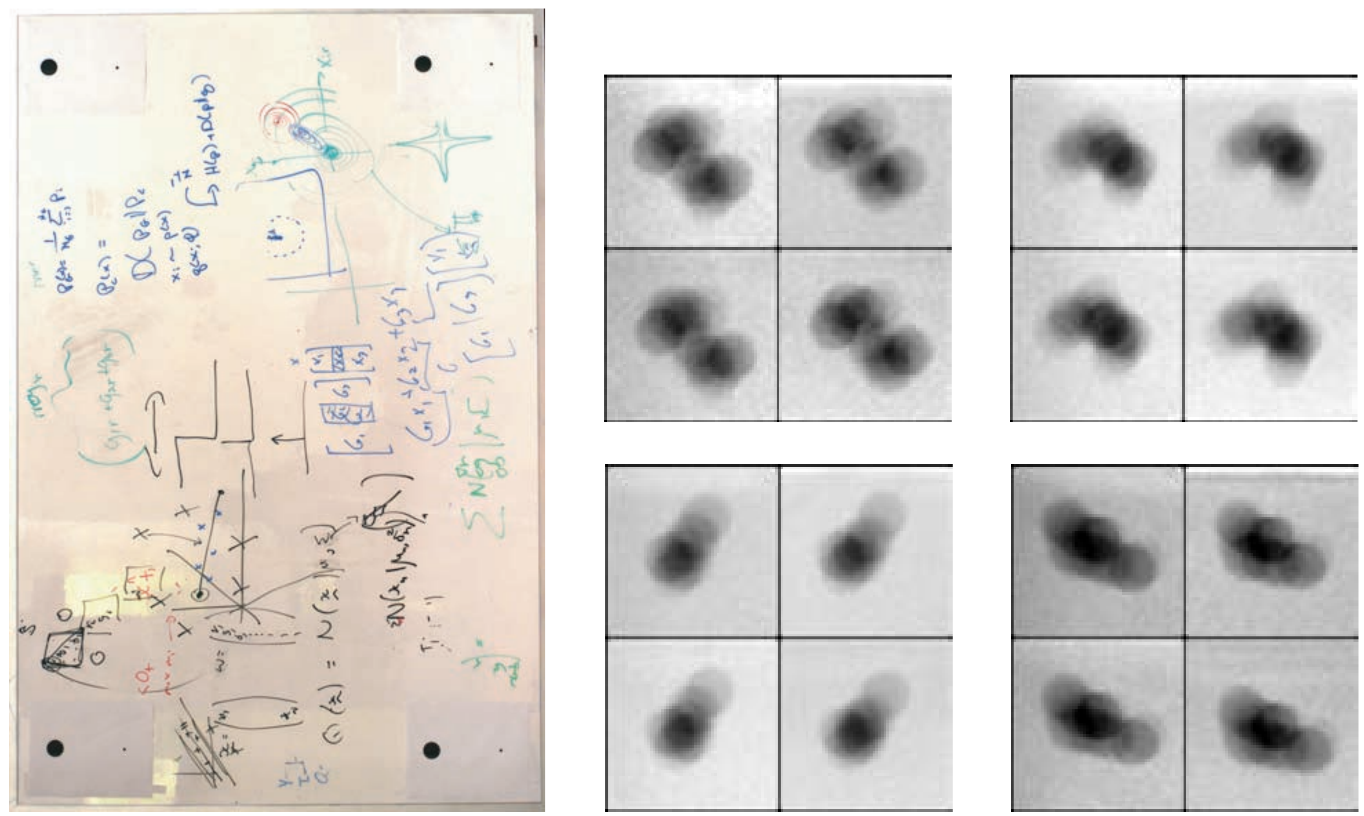

Figure 4: *Left*: The whiteboard test scene with dots in each corner. *Right*: Dots from the corners of images taken by different people. Within each image, the dot trajectories are very similar suggesting that image blur is well modeled as a spatially invariant convolution.

[Fergus et al. 2006]

Bayesian estimate of kernel

$$
\begin{aligned}
p(\mathbf{K}, \nabla \mathbf{L}_p | \nabla \mathbf{P}) & \propto p(\nabla \mathbf{P} | \mathbf{K}, \nabla \mathbf{L}_p) p(\nabla \mathbf{L}_p) p(\mathbf{K}) & (2) \\
& = \prod_i \mathtt{N}(\nabla \mathbf{P}(i) | (\mathbf{K} \otimes \nabla \mathbf{L}_p(i)), \sigma^2) & (3) \\
& \quad \prod_i \sum_{c=1}^{C} \pi_c \mathtt{N}(\nabla \mathbf{L}_p(i) | 0, v_c) \prod_j \sum_{d=1}^{D} \pi_d \mathtt{E}(\mathbf{K}_j | \lambda_d)
\end{aligned}
$$

[Fergus et al. 2006]

Prior on image gradients

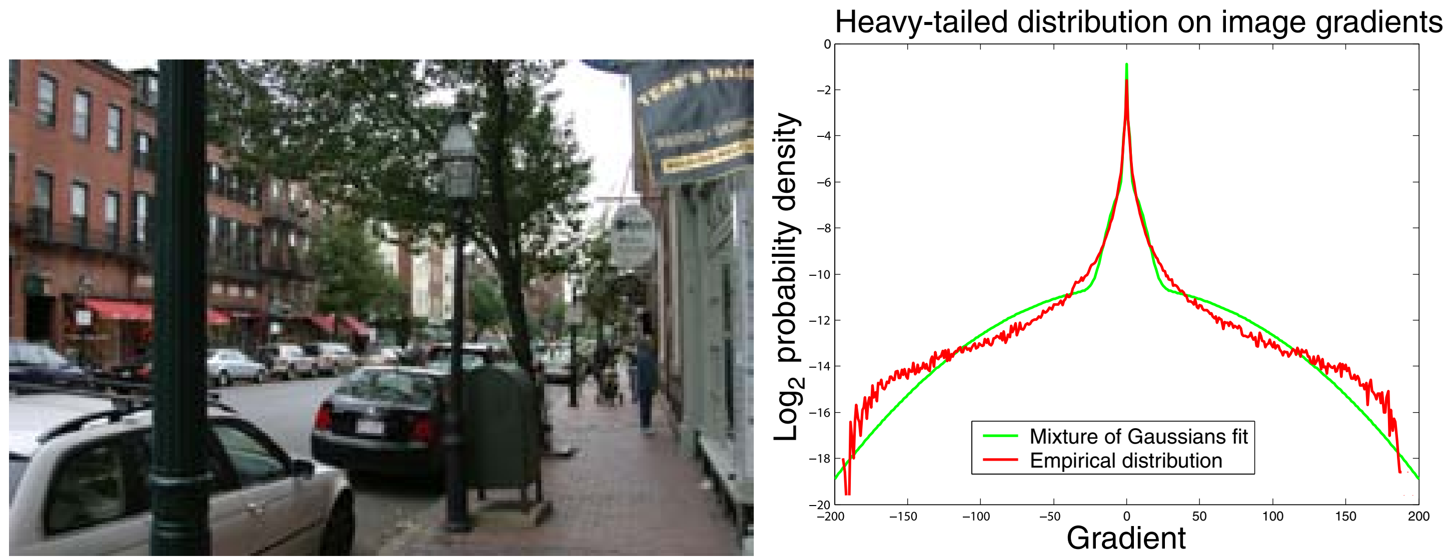

Figure 2: *Left:* A natural scene. *Right:* The distribution of gradient magnitudes within the scene are shown in red. The *y*-axis has a logarithmic scale to show the heavy tails of the distribution. The mixture of Gaussians approximation used in our experiments is shown in green.

[Fergus et al. 2006]

Estimation results

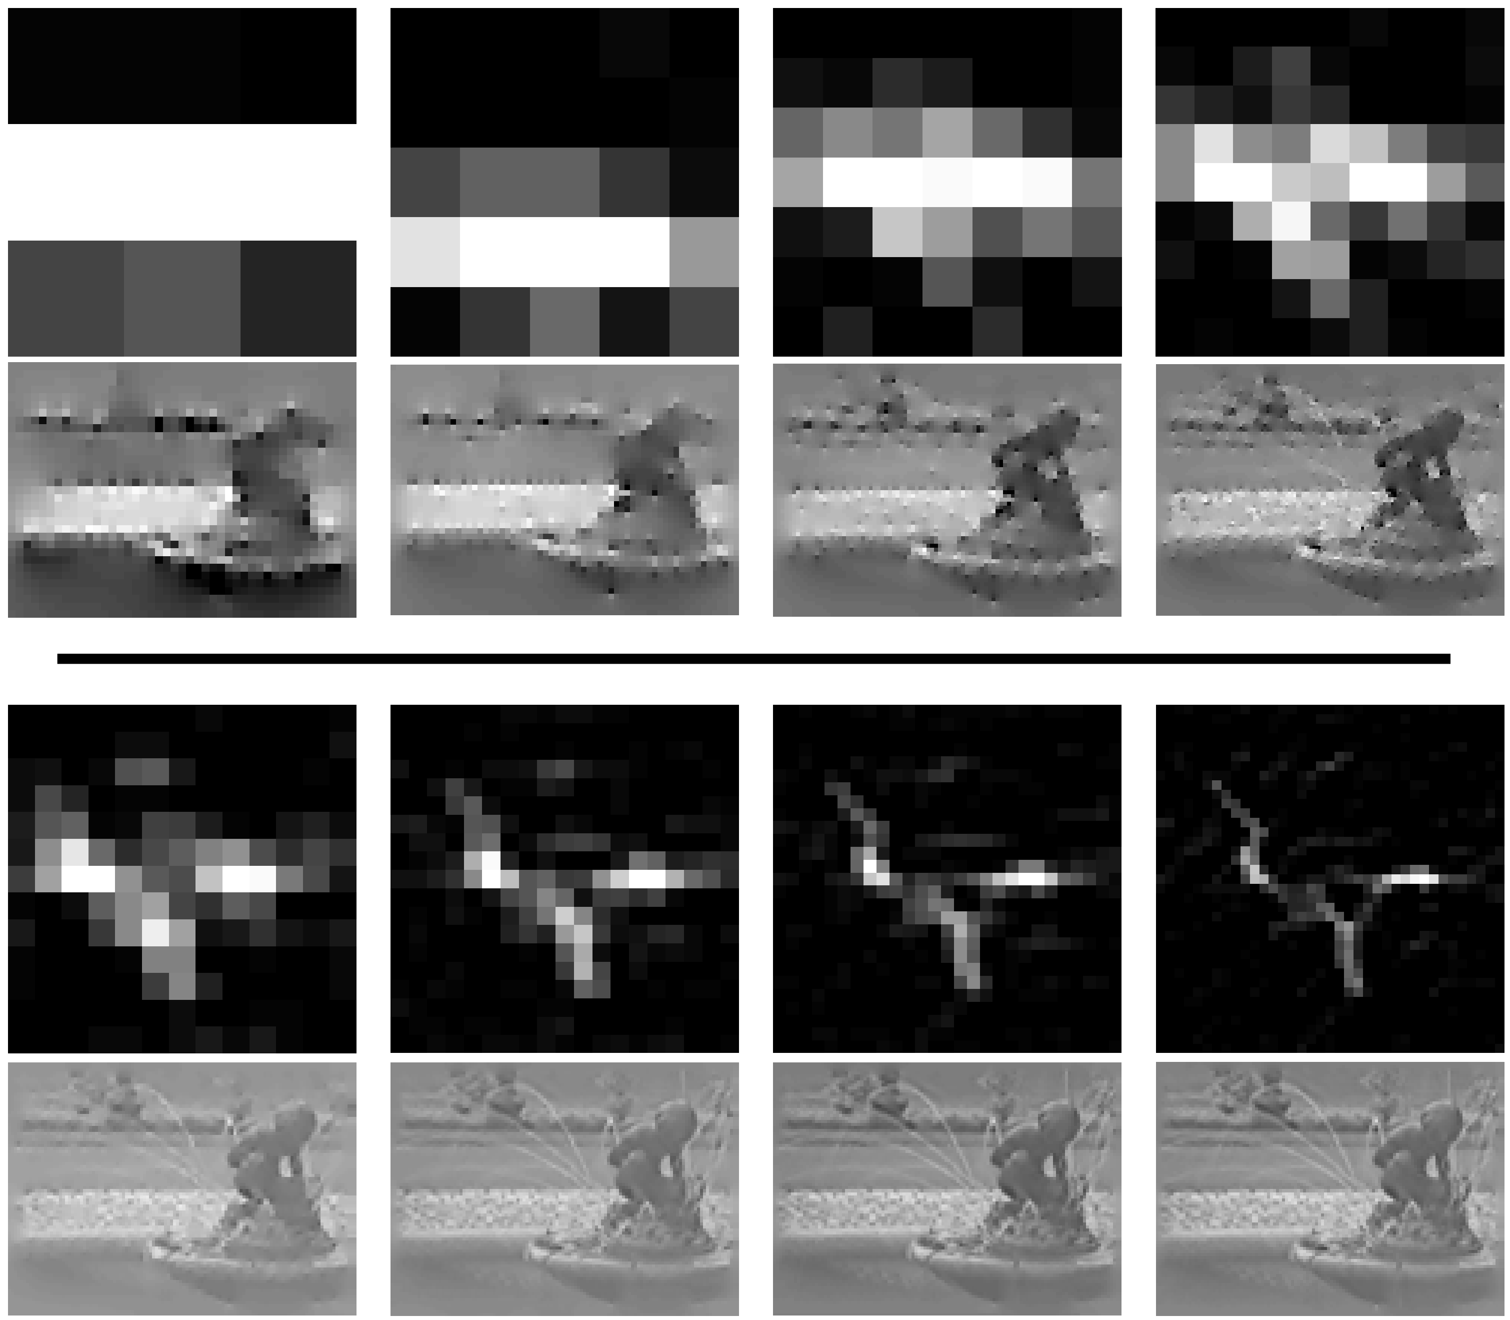

[Fergus et al. 2006]

[Fergus et al. 2006]

[Fergus et al. 2006]

[Fergus et al. 2006]

Shan et al.

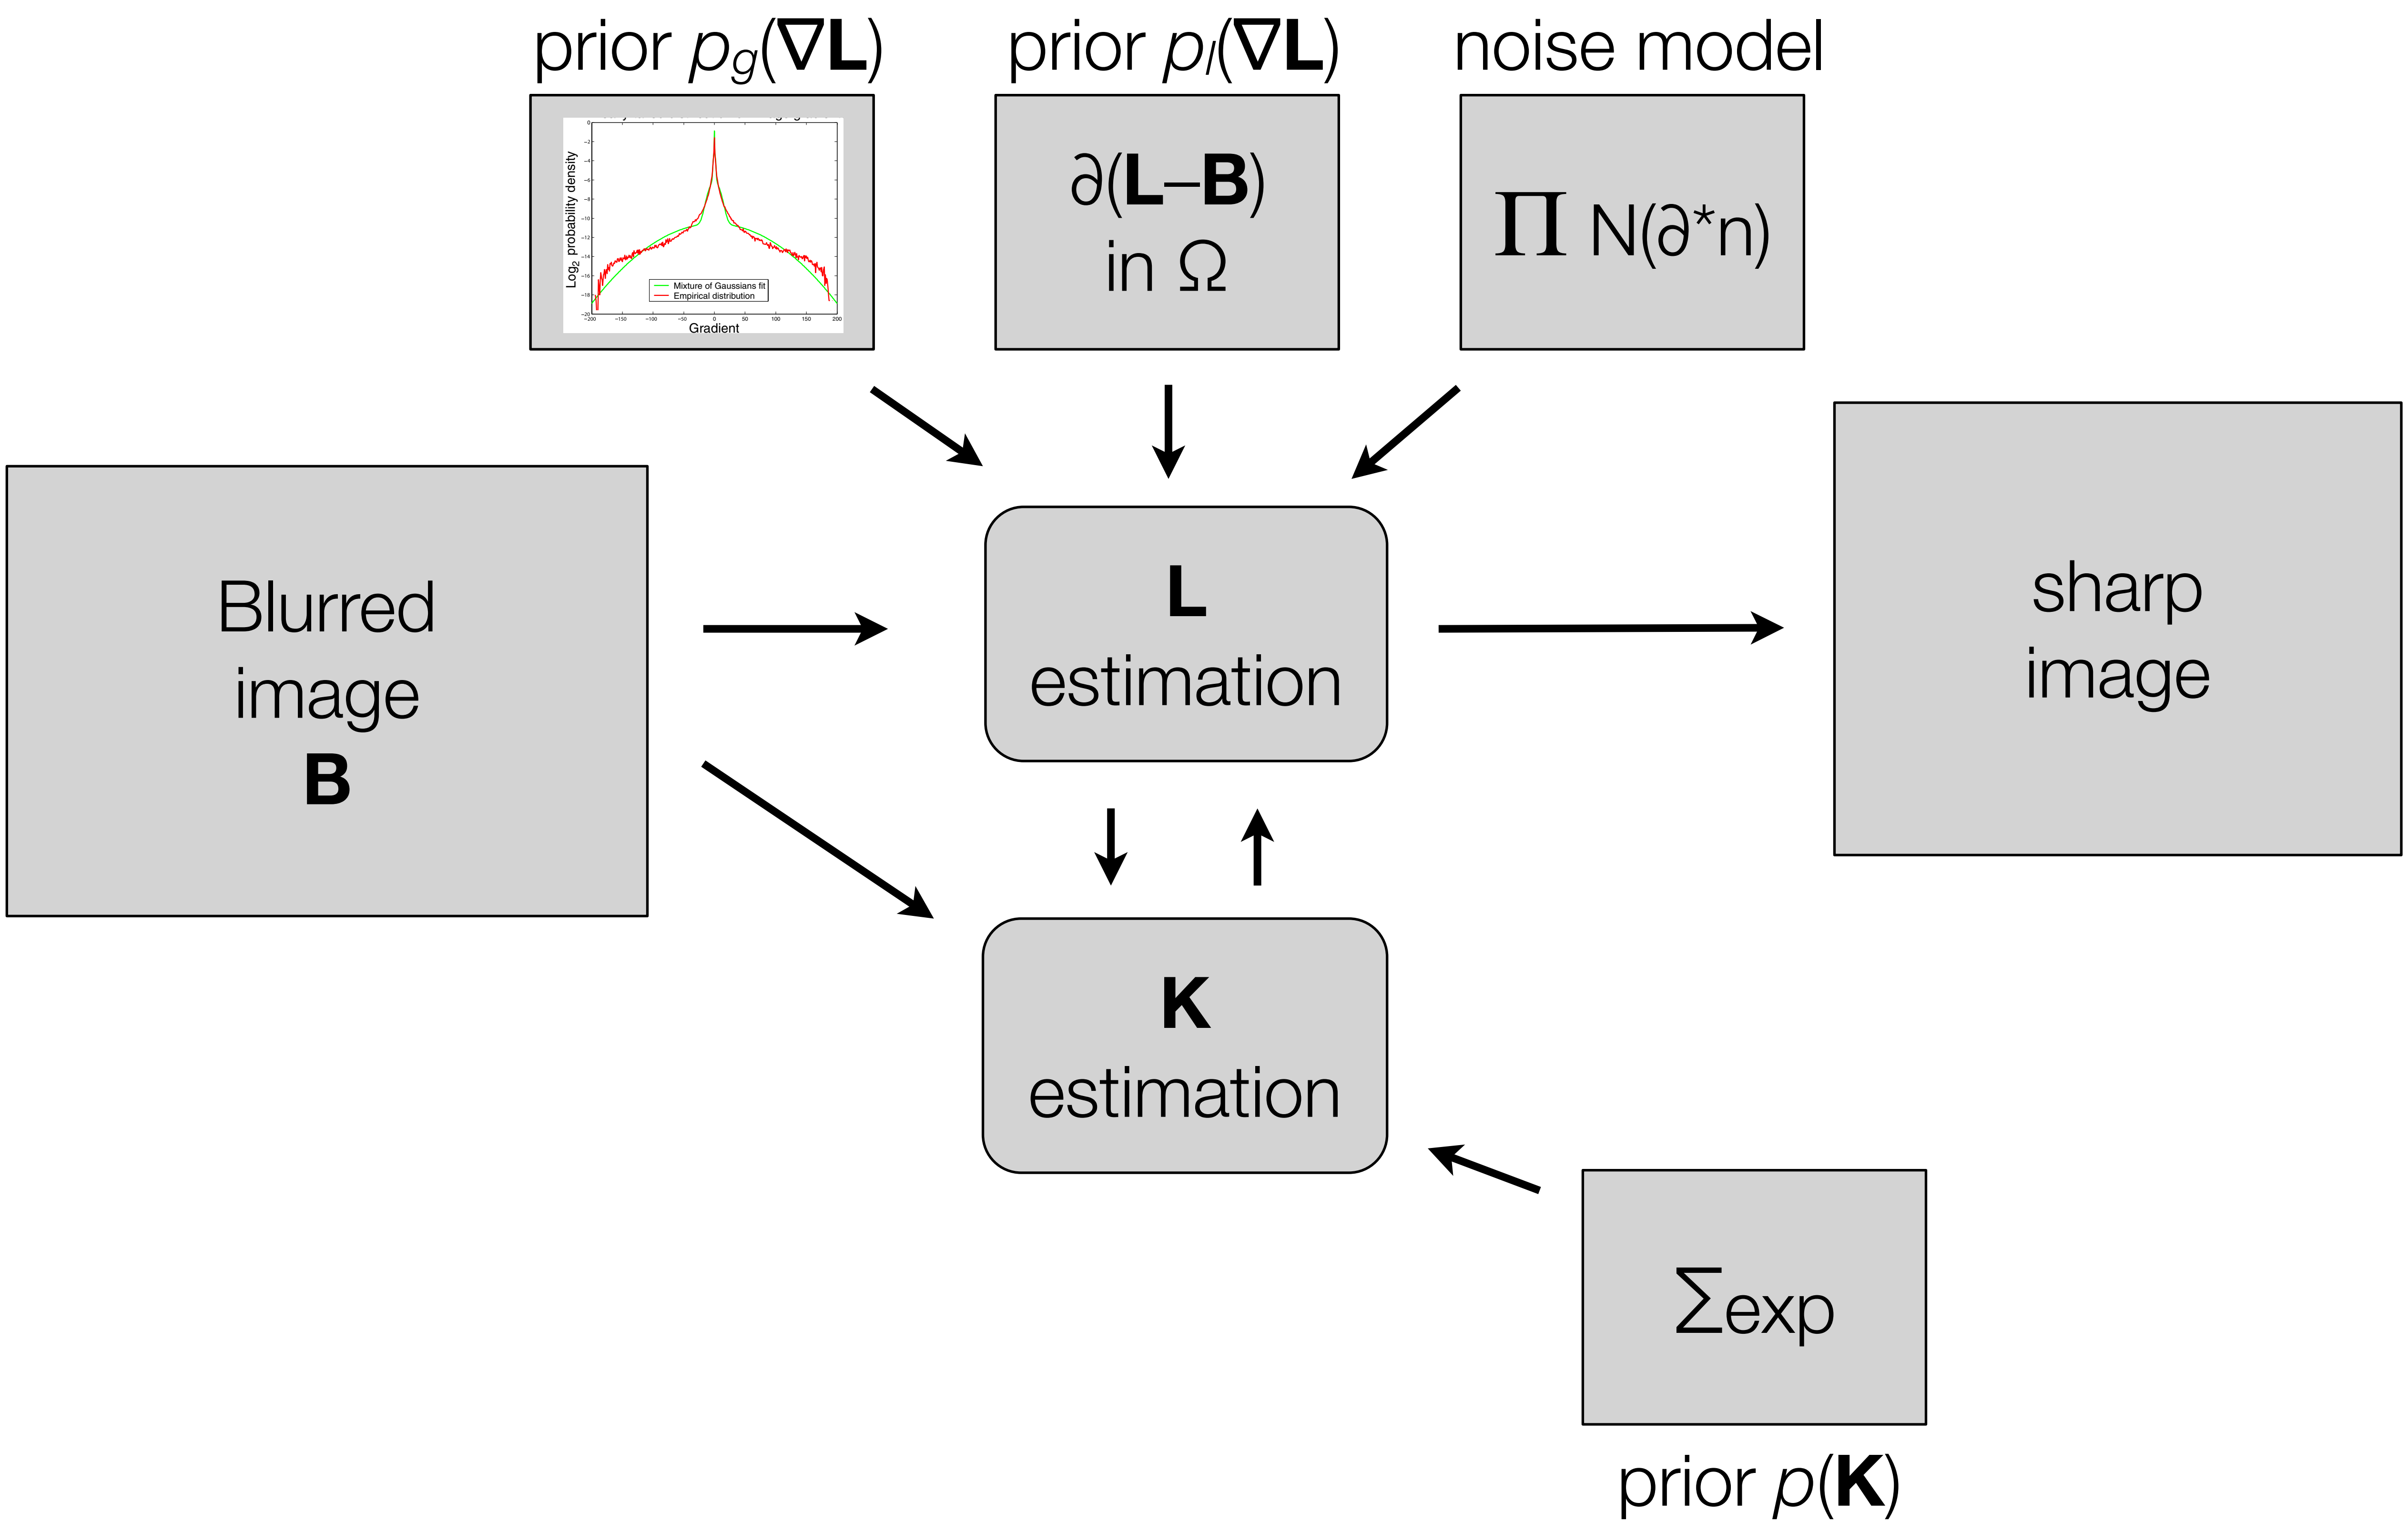

prior $p_g(\nabla \mathbf{L})$
Log2 probability density
Mixture of Gaussians fit
Empirical distribution
Gradient
prior $p_l(\nabla \mathbf{L})$
$\partial(\mathbf{L}-\mathbf{B})$ in Ω
noise model
$\prod N(\partial^* n)$
Blurred image **B**
L estimation
sharp image
K estimation
$\Sigma \exp$
prior $p(\mathbf{K})$

Challenges (1)

Assuming no noise

[Slides by Qi Shan]

Challenges (2)

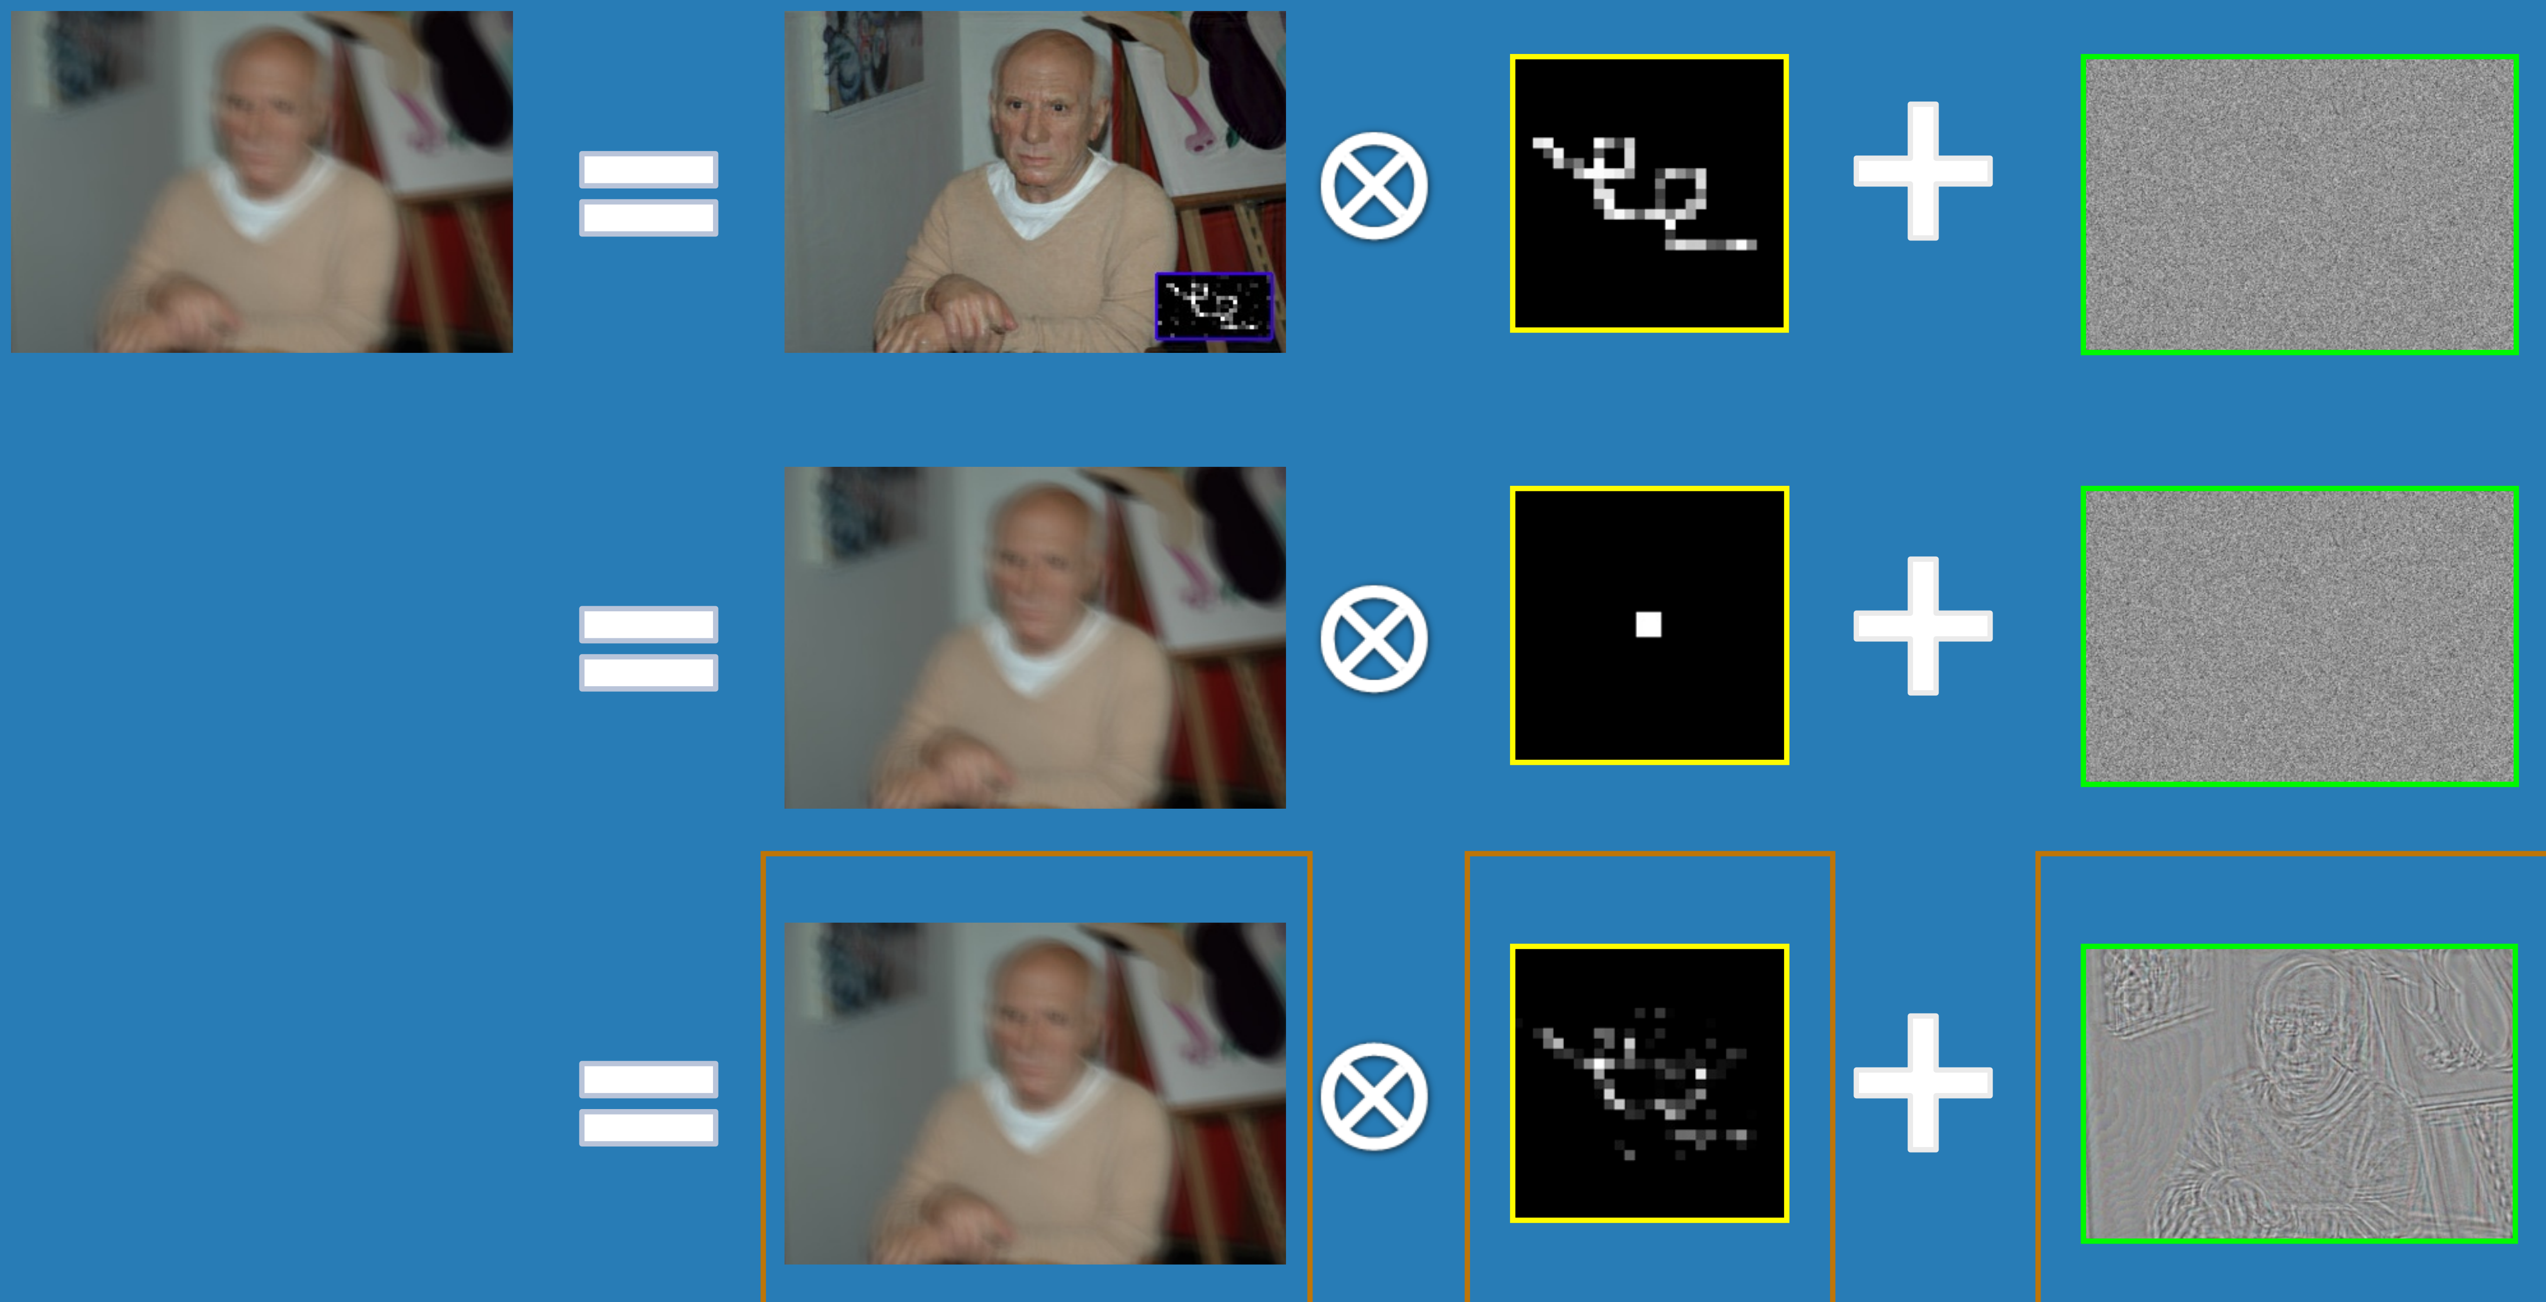

With noise

[Slides by Qi Shan]

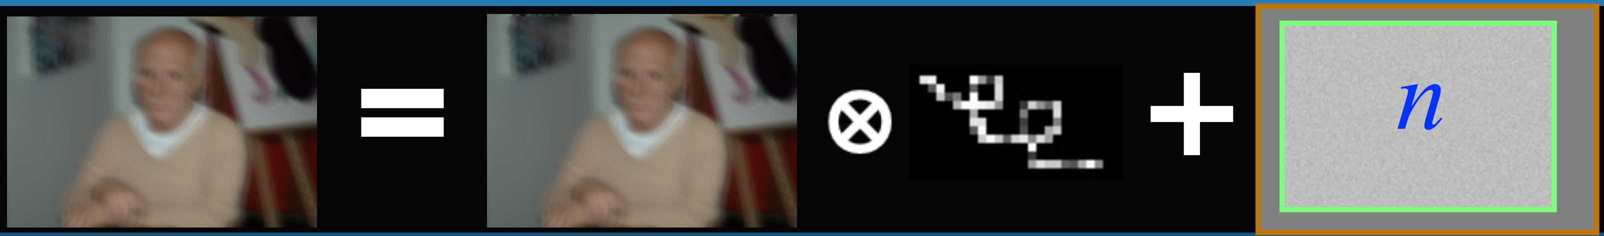

Noise constraint

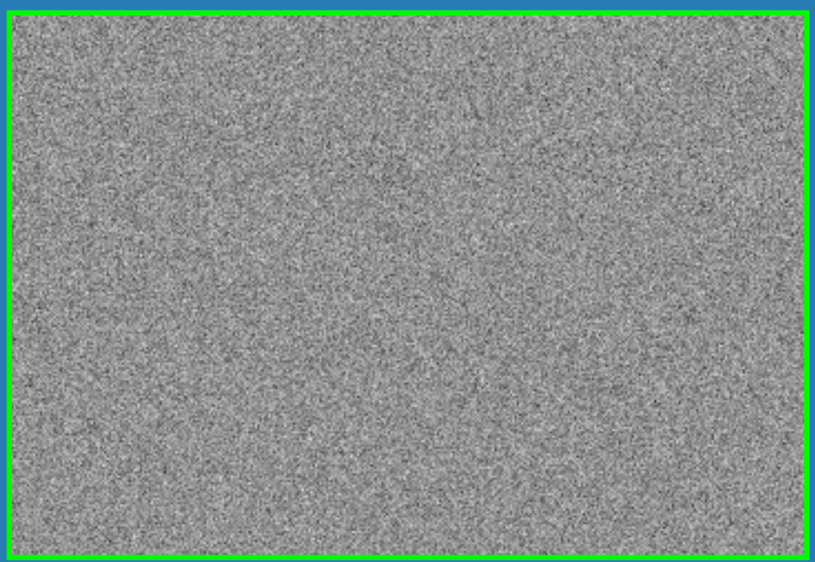

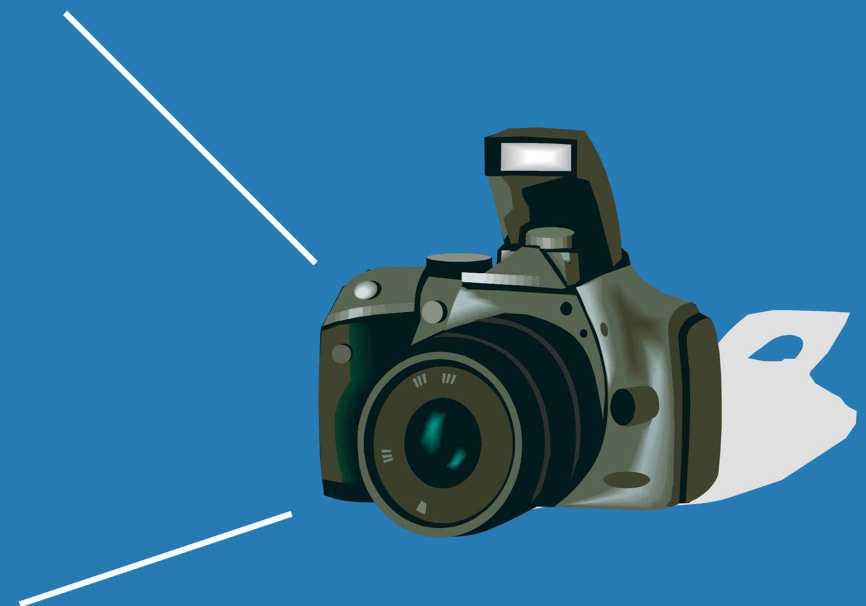

[Slides by Qi Shan]

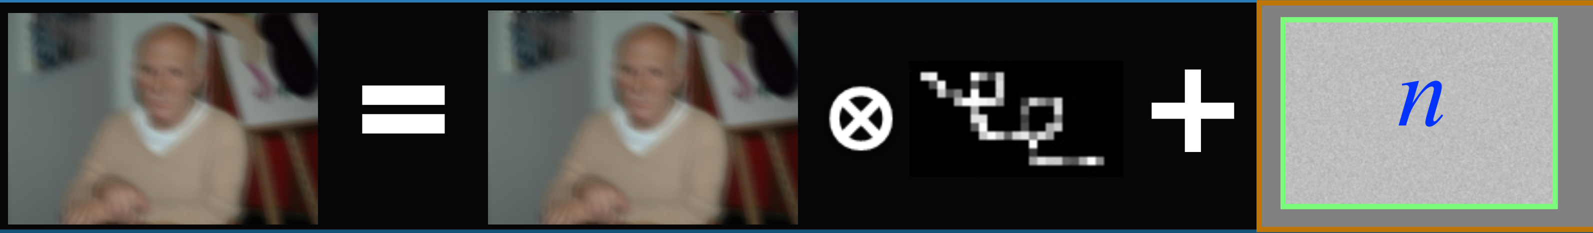

Noise constraint

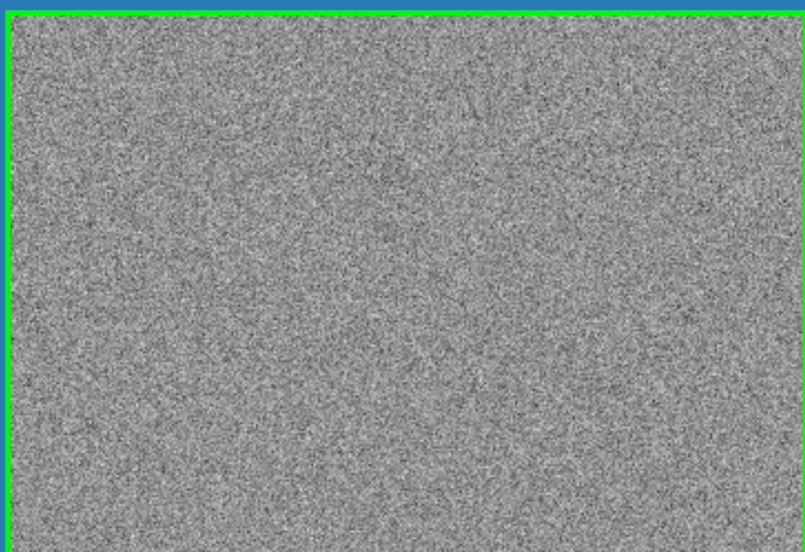

Possible noise models:

(1) $$\prod_i N(n_i \mid 0, \zeta_0)$$

(2) $$\prod_i N(\nabla n_i \mid 0, \zeta_1)$$

$$\prod_i N(n_i \mid 0, \zeta_0) \prod_i N(\nabla n_i \mid 0, \zeta_1)$$

[Slides by Qi Shan]

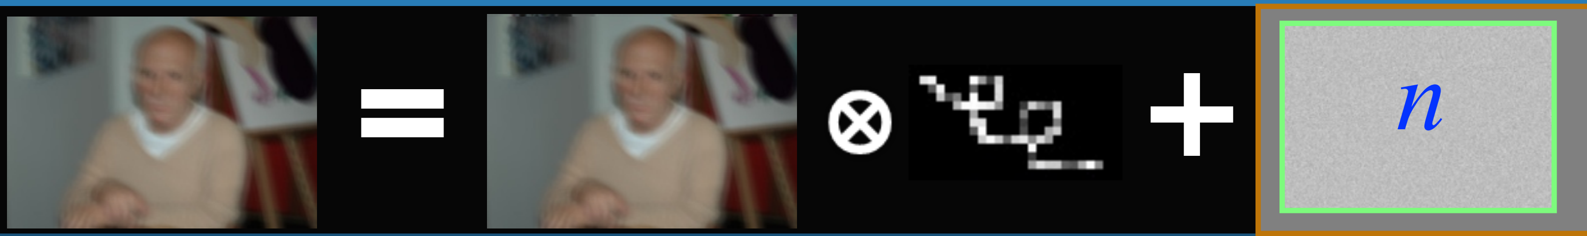

SIGGRAPH2008

Noise constraint

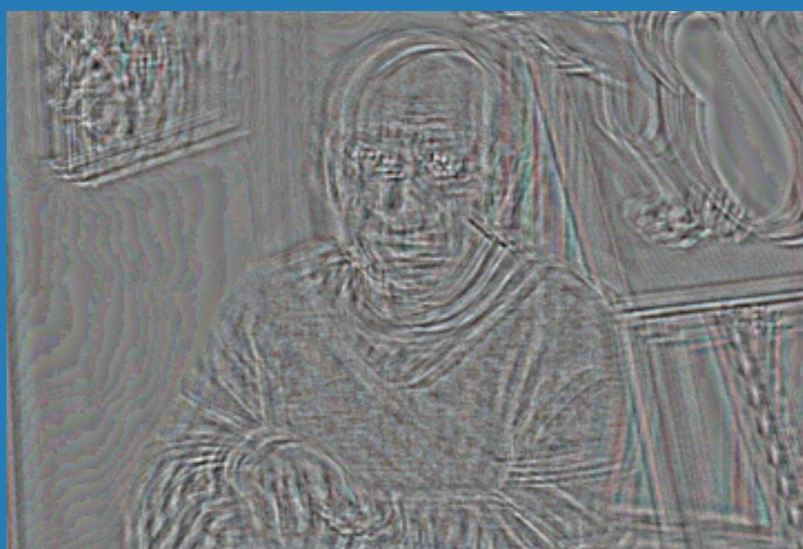

Possible noise models:

(1) $\prod_i N(n_i \mid 0, \zeta_0)$

(2) $\prod_i N(\nabla n_i \mid 0, \zeta_1)$

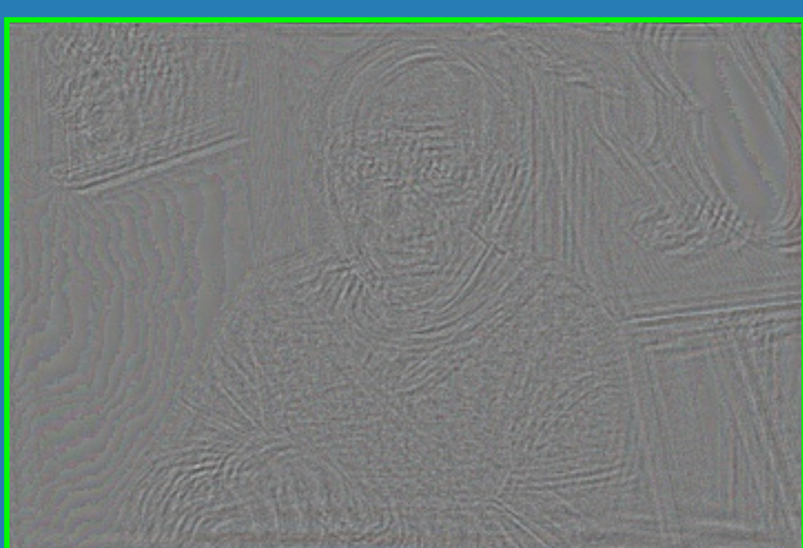

$$\prod_i N(n_i \mid 0, \zeta_0) \prod_i N(\nabla n_i \mid 0, \zeta_1) \prod_i N(\nabla\nabla n_i \mid 0, \zeta_2)$$

[Slides by Qi Shan]

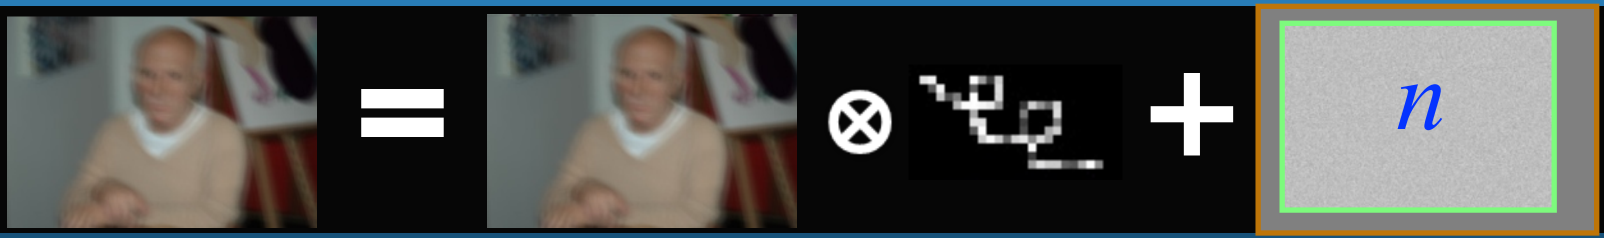

A random variable following an independent Gaussian distribution also has its any order derivative following it. *[Simon 2002]*

$$P(n) = \prod_i N(n_i \mid 0, \zeta_0) \prod_i N(\nabla n_i \mid 0, \zeta_1) \prod_i N(\nabla(\nabla n_i) \mid 0, \zeta_2)$$

[Slides by Qi Shan]

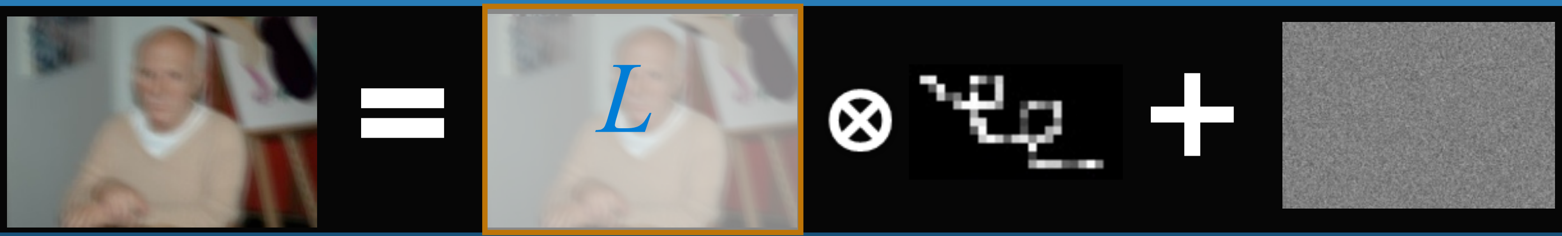

Image Global Statistics

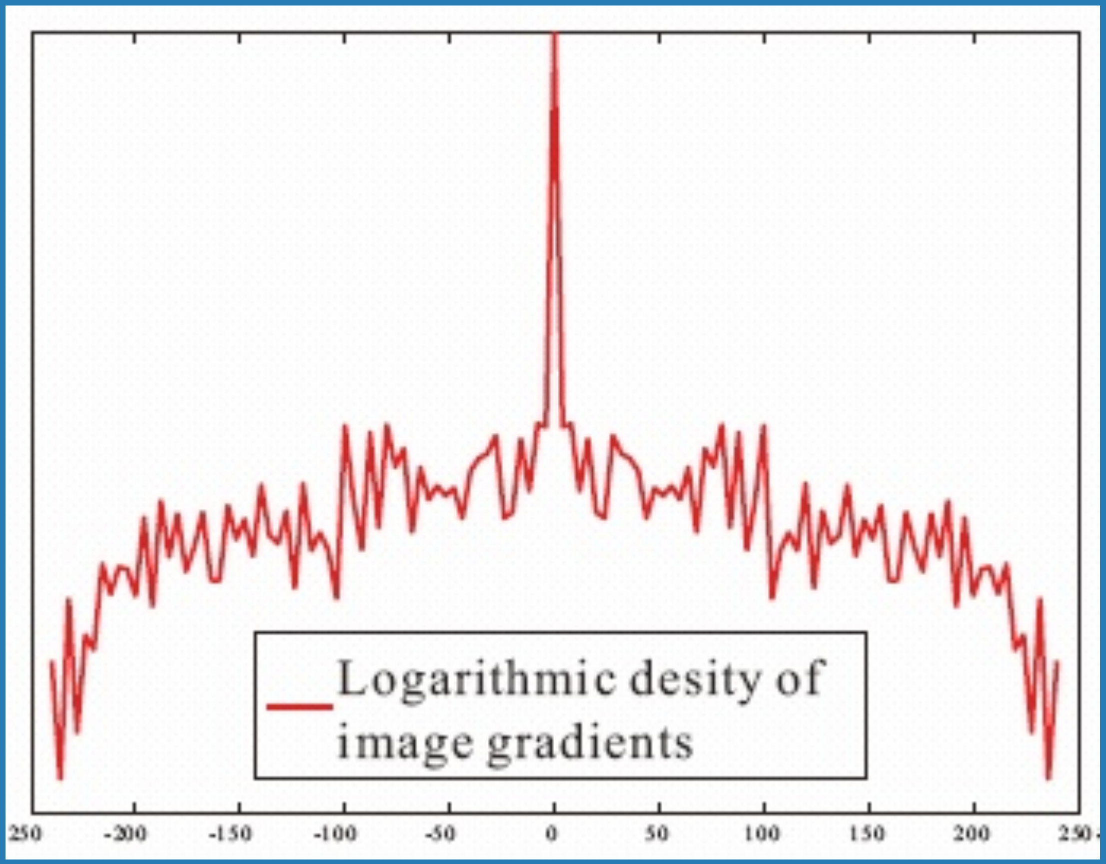

. . .

[Slides by Qi Shan]

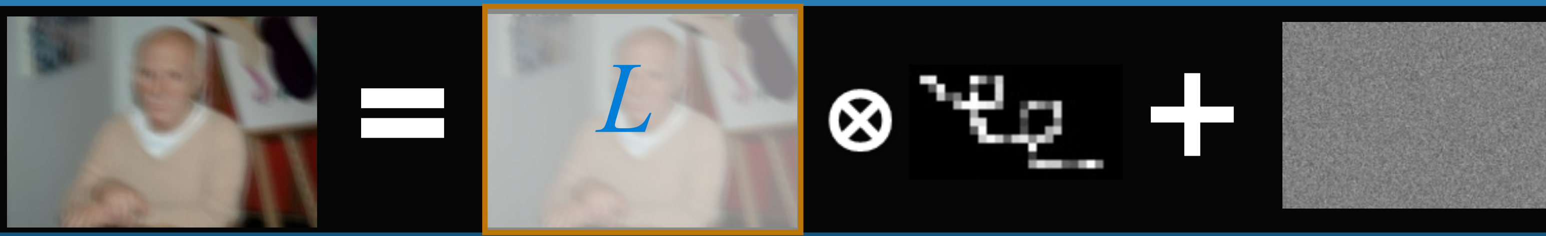

Image Global Statistics

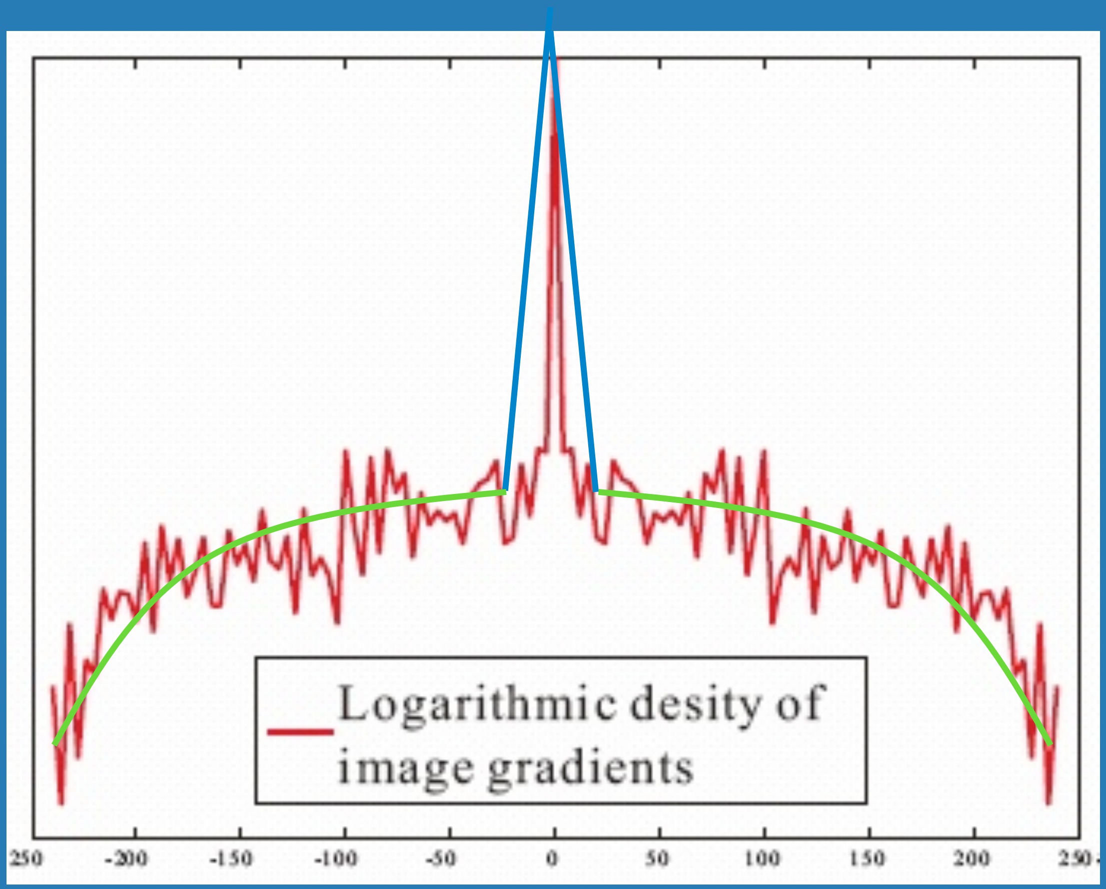

■ ■ ■

[Slides by Qi Shan]

Image Global Statistics

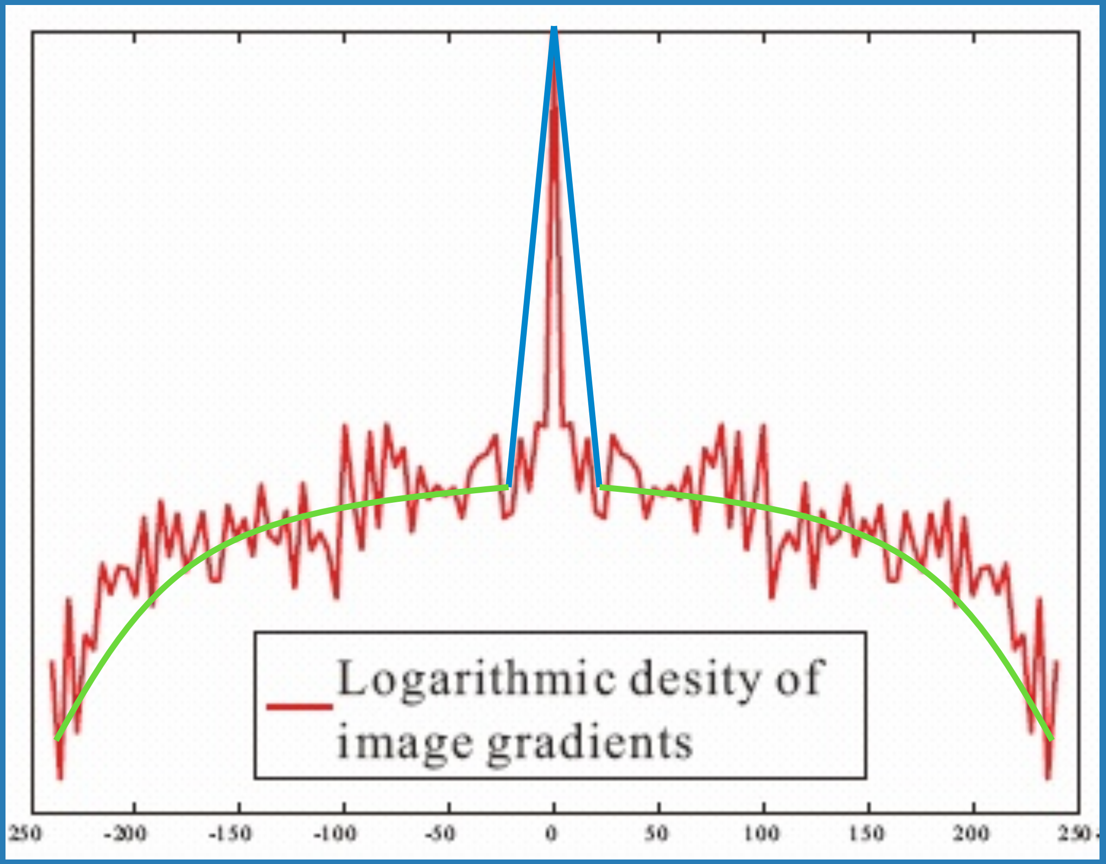

$$\log(P_1(\nabla L)) = \begin{cases} -k\,|\nabla L| & x \le c \\ -(a(\nabla L)^2 + b) & x > c \end{cases}$$

[Slides by Qi Shan]

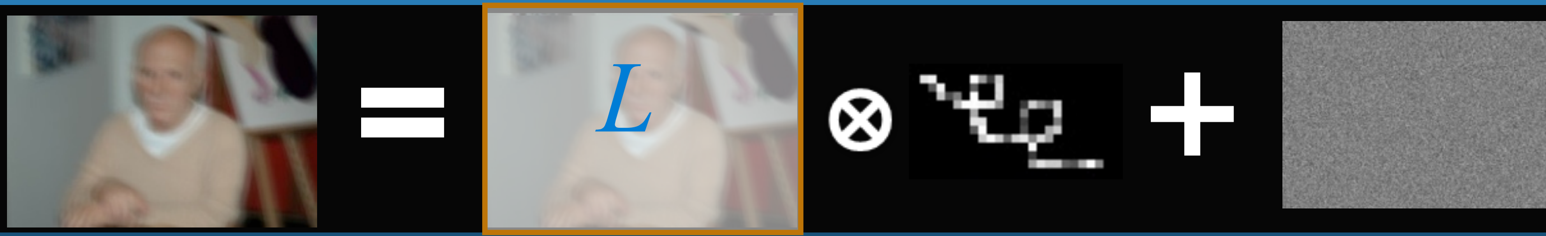

Image Local Constraint

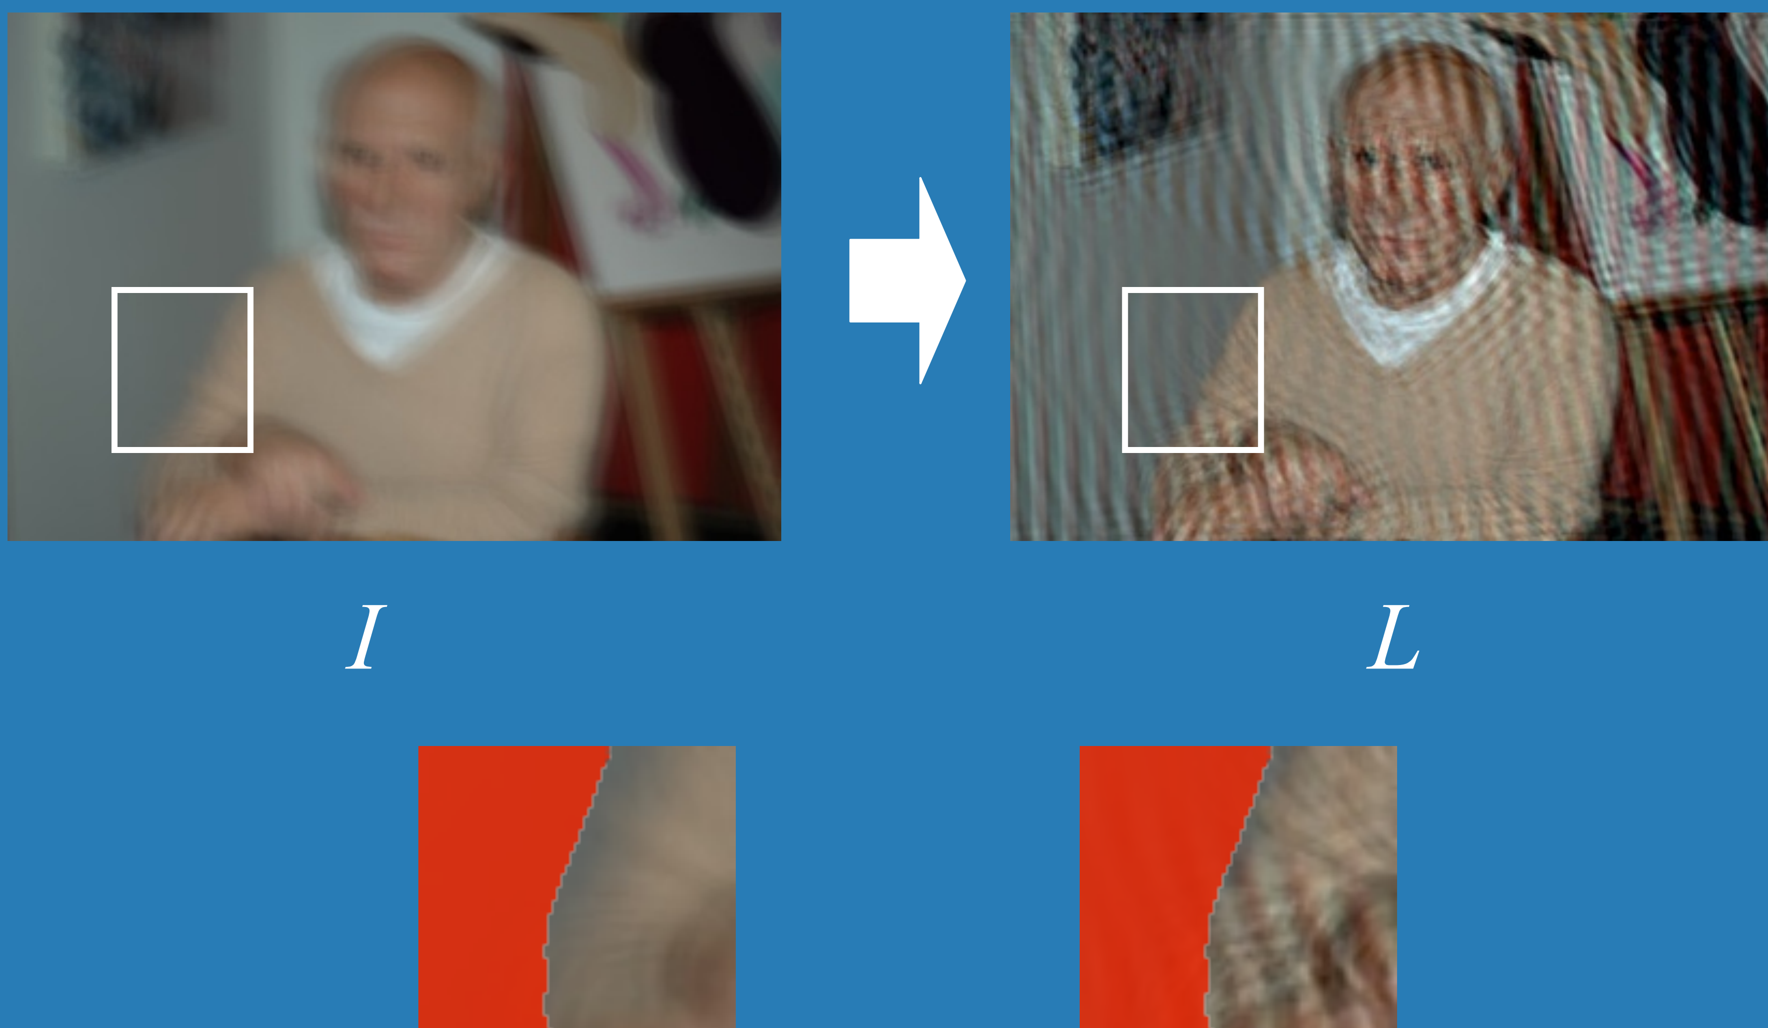

[Slides by Qi Shan]

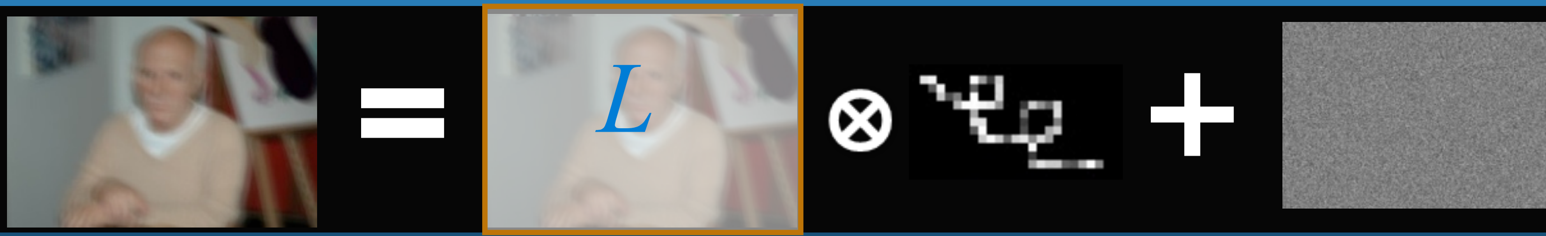

Image Local Constraint

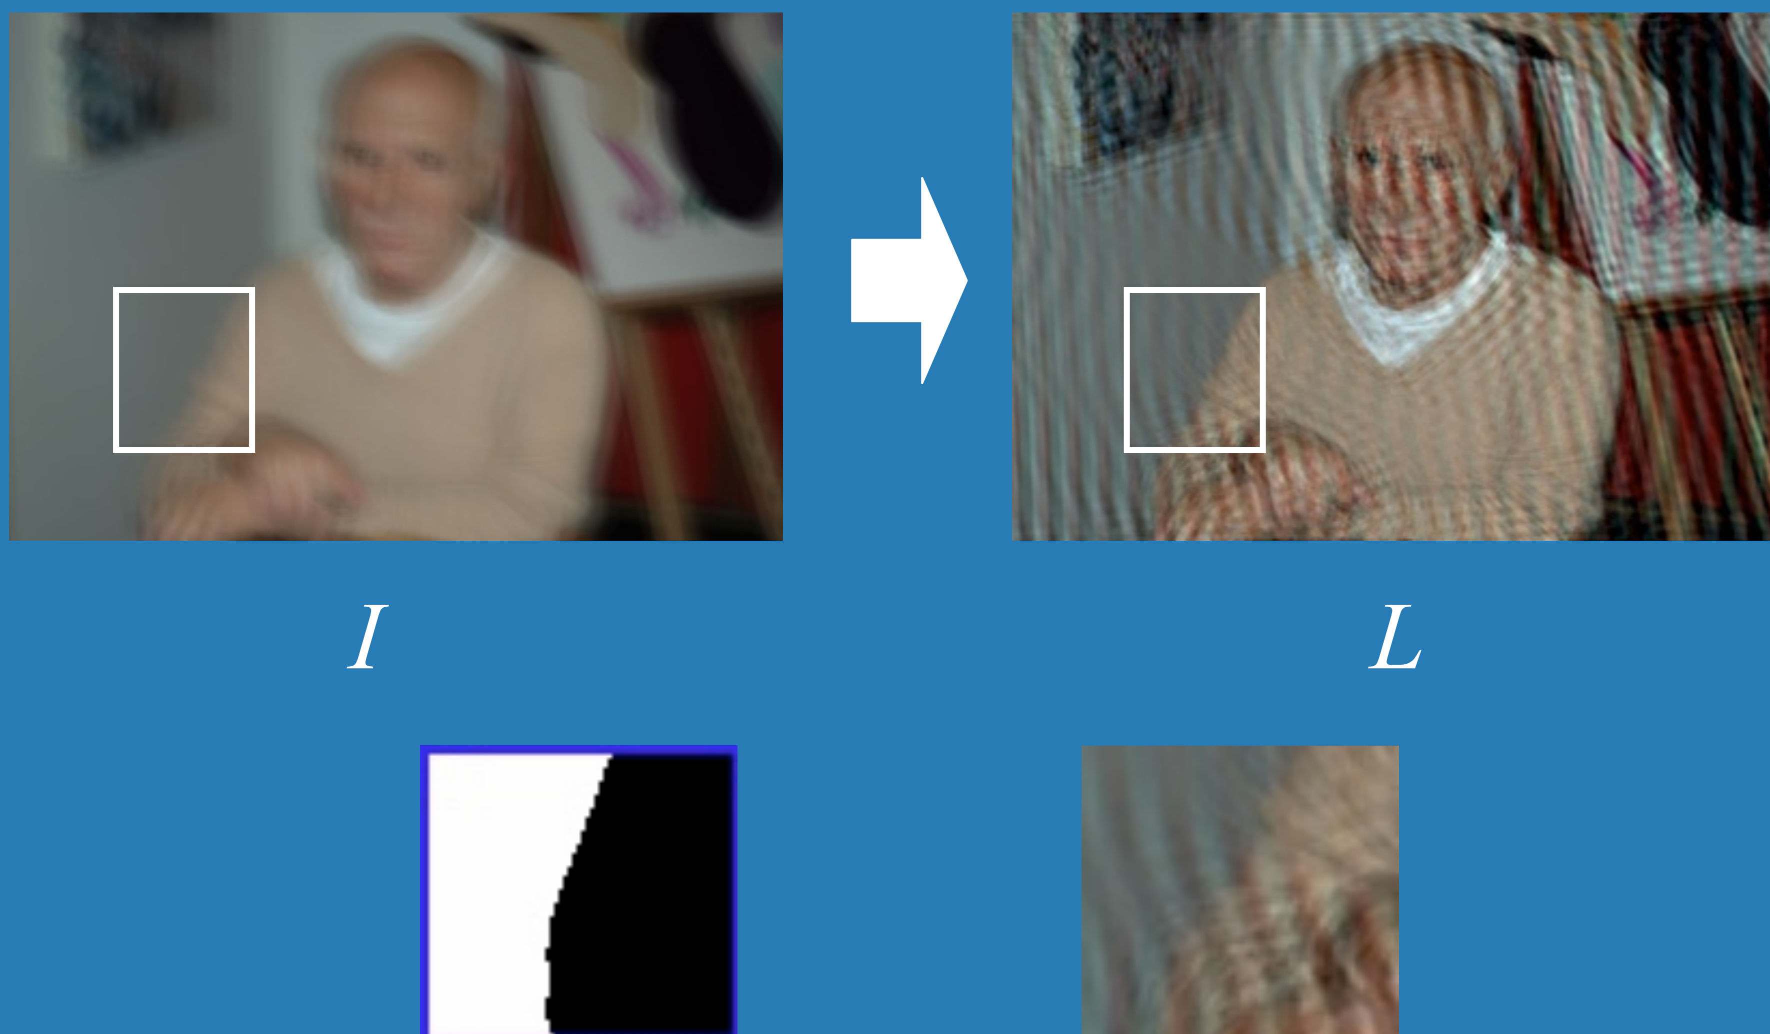

[Slides by Qi Shan]

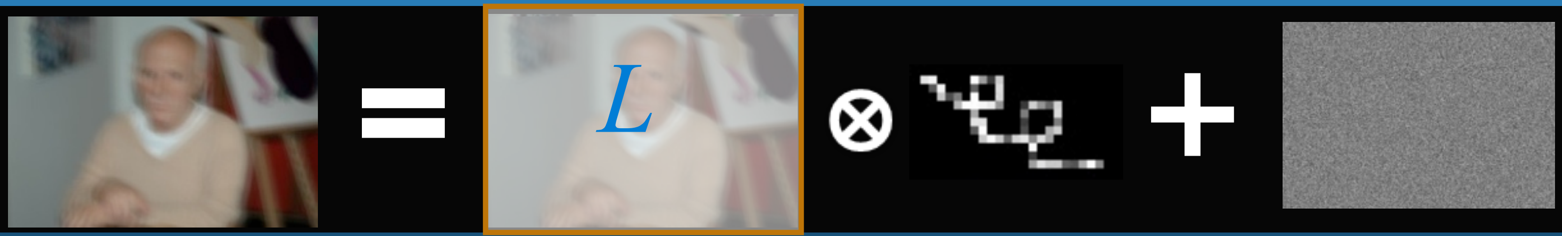

Image Local Constraint

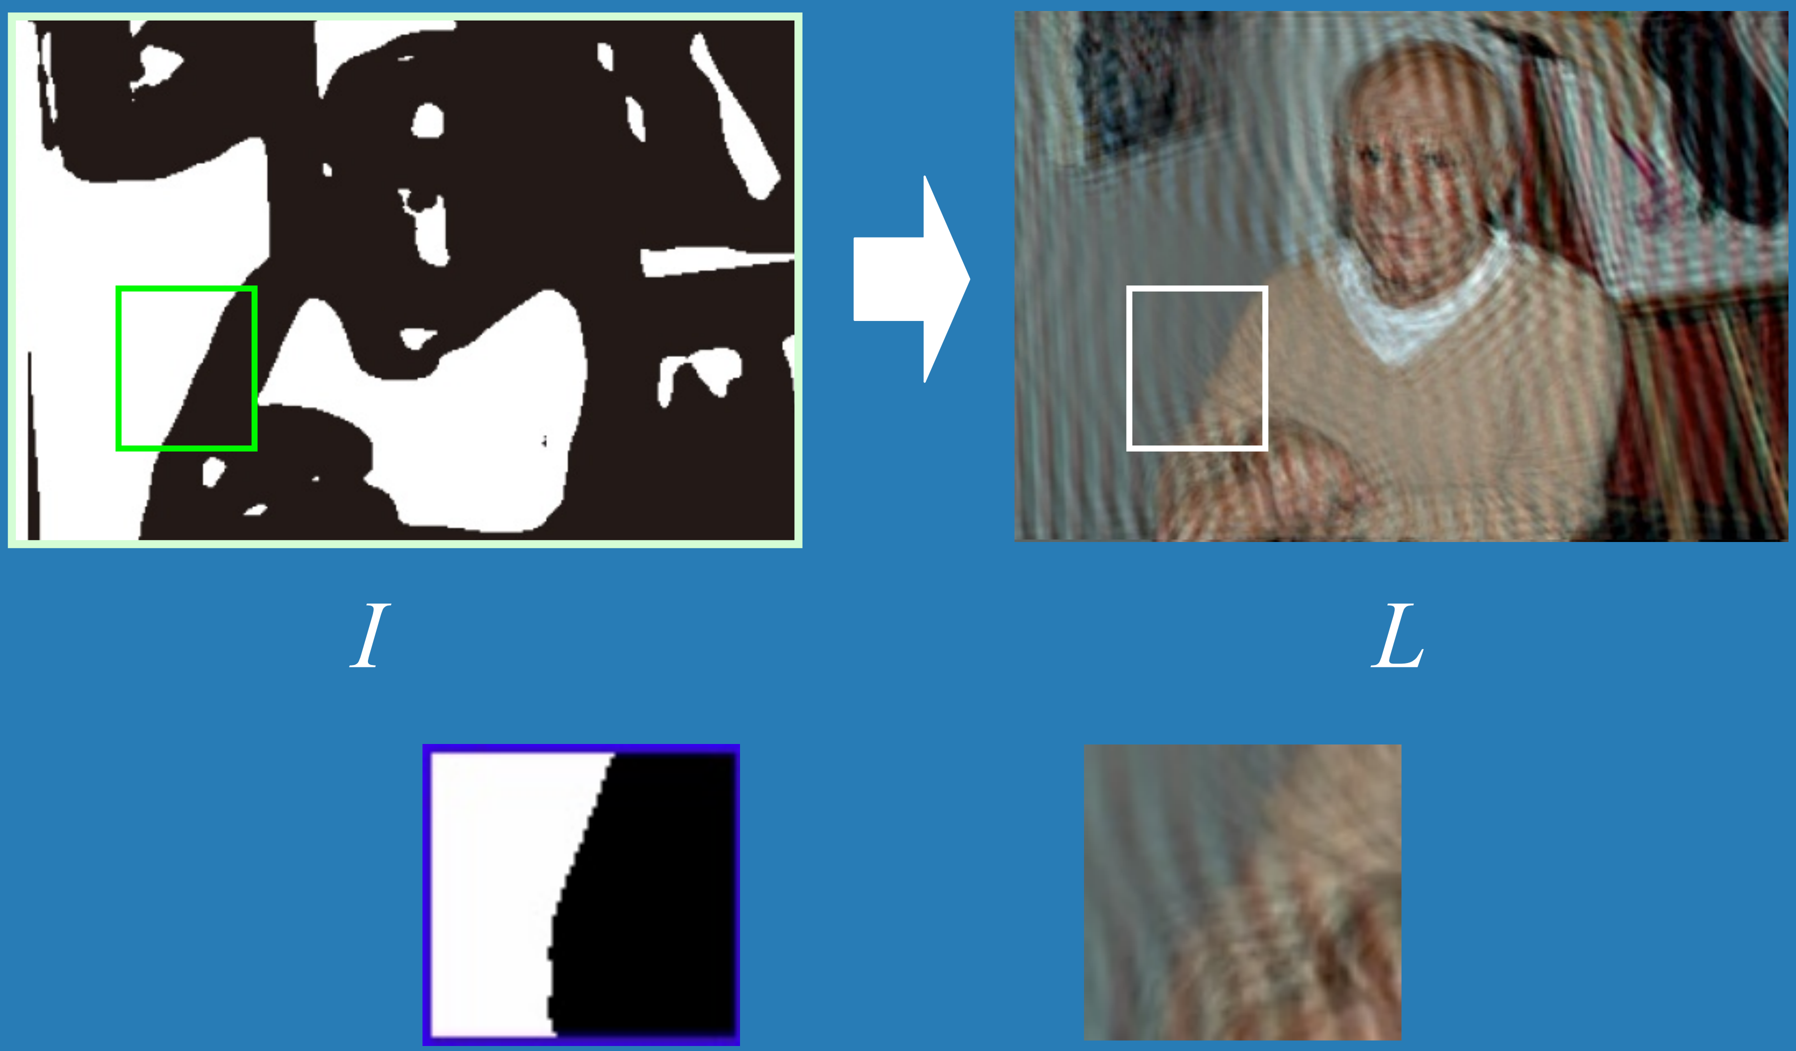

[Slides by Qi Shan]

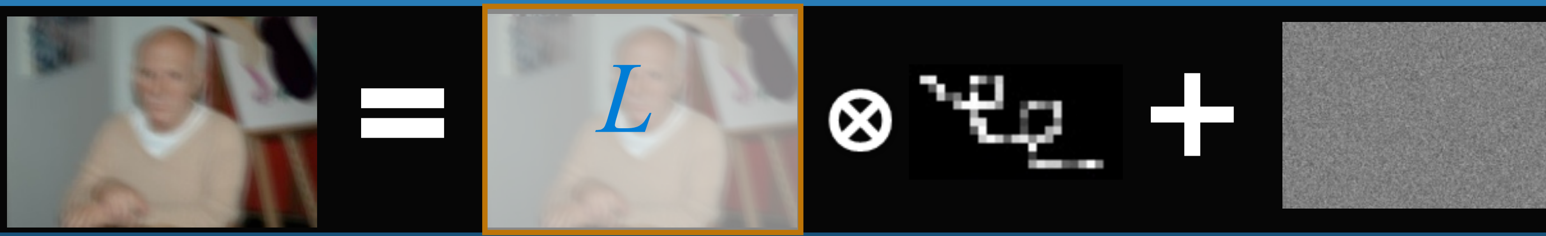

SIGGRAPH2008

Image Local Constraint

$$p_2(L) = \prod_{i \in white} N(\nabla L_i - \nabla I_i \mid 0, \sigma_1)$$

[Slides by Qi Shan]

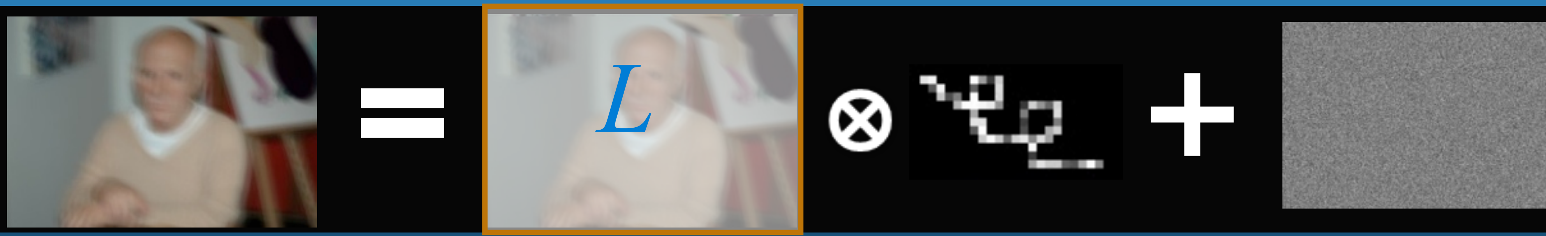

SIGGRAPH2008

Image Local Constraint

I $\qquad\qquad$ L

$$p_2(L) = \prod_{i \in white} N(\nabla L_i - \nabla I_i \mid 0, \sigma_1)$$

[Slides by Qi Shan]

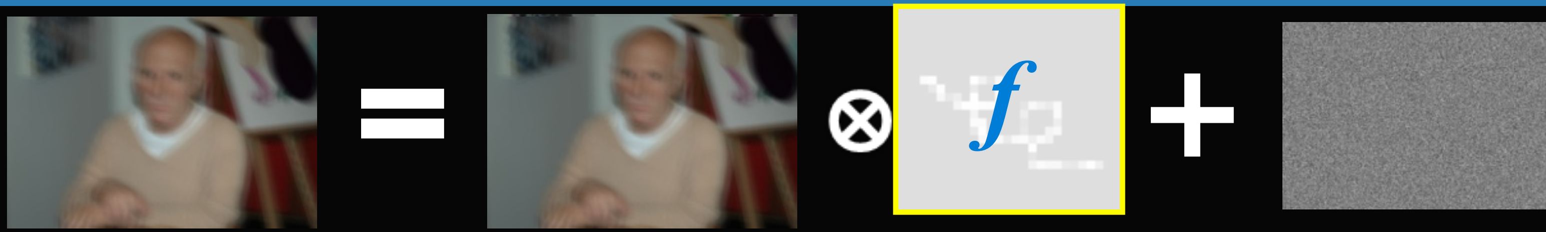

Kernel Statistics

exponentially distributed

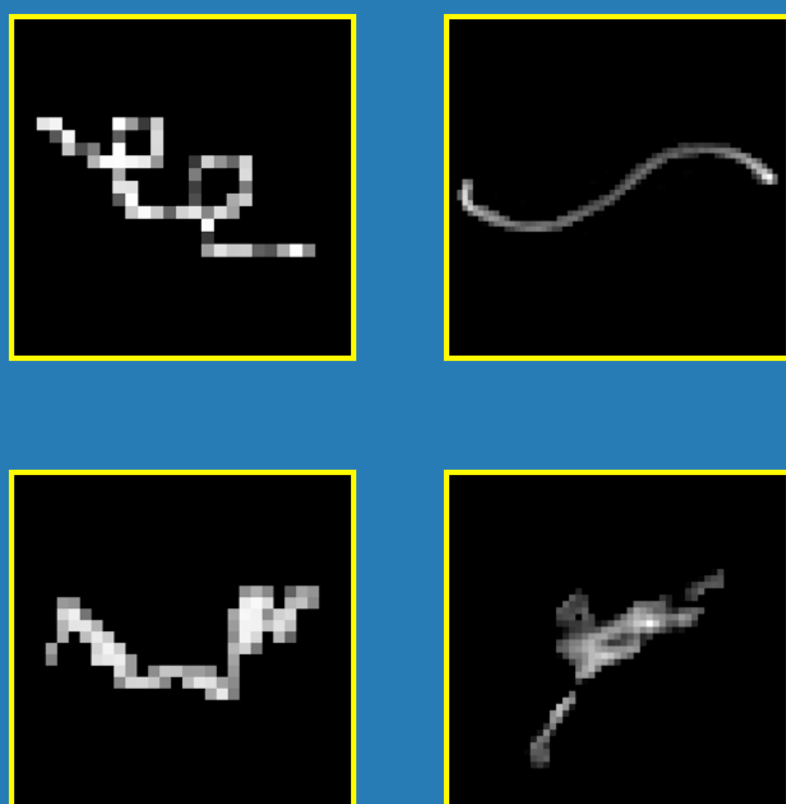

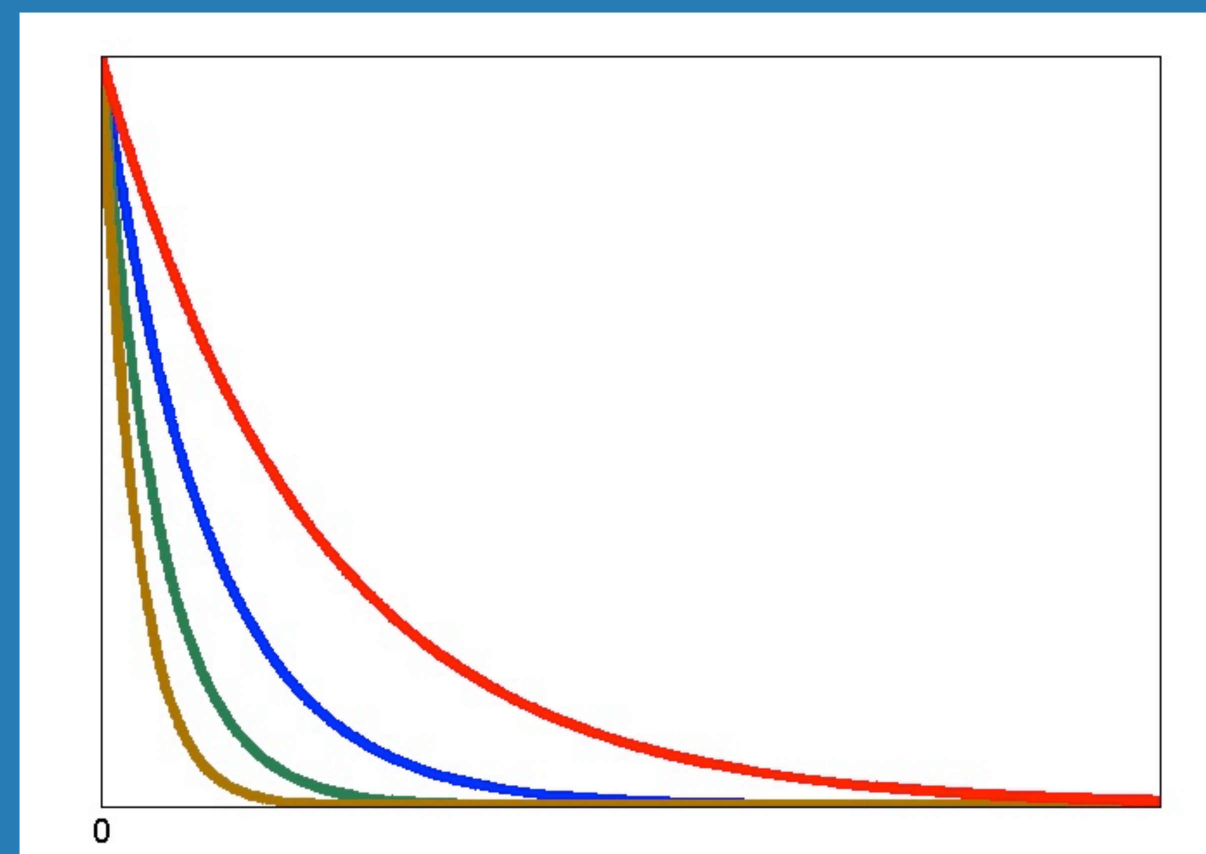

$$p(f) = \prod_i e^{-\tau f_j}, \qquad f_j \geq 0$$

[Slides by Qi Shan]

Combining All constraints

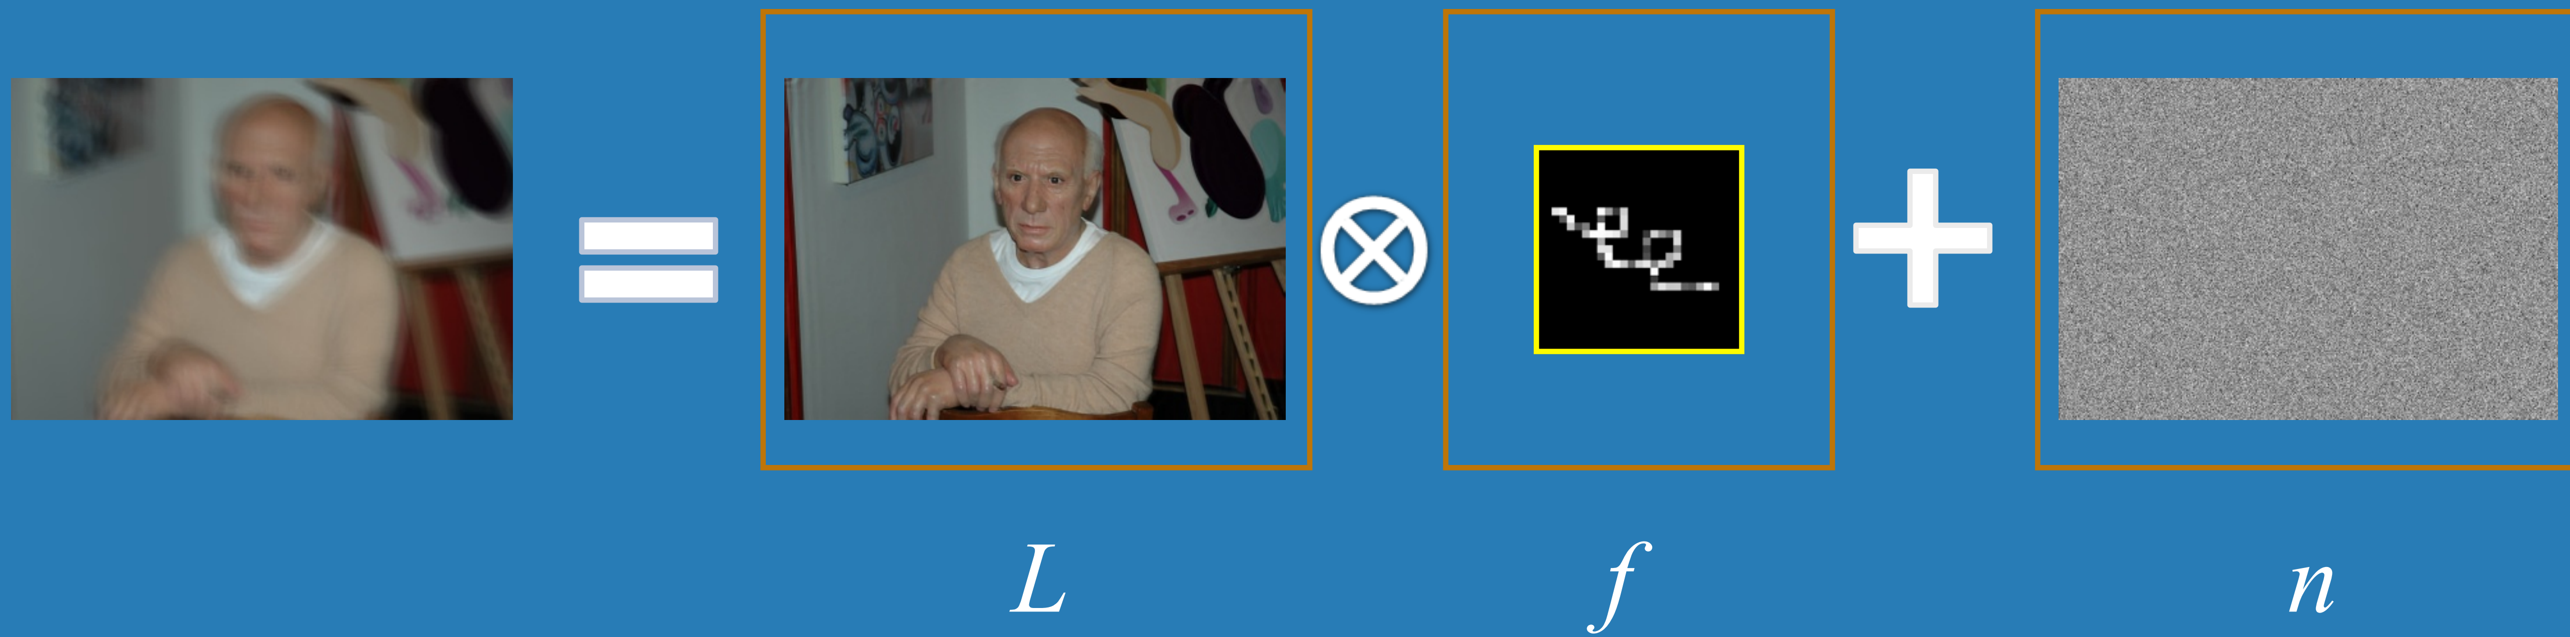

$$\min E(L, f) = \min \log[p(n)\, p_1(\nabla L)\, p_2(L)\, p(f)]$$

[Slides by Qi Shan]

Two-step iterative optimization

- Optimize L
- Optimize f

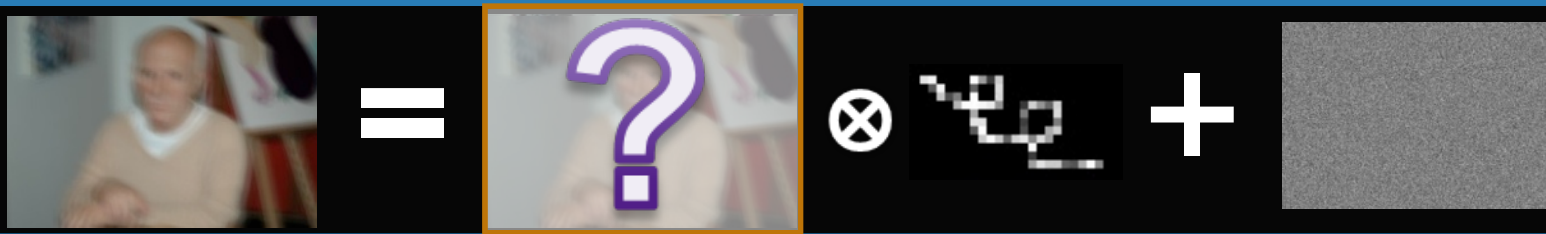

SIGGRAPH2008

Iteratively optimize L:

Update L

Update Ψ

Iteration [illegible]

(converge)

[Slides by Qi Shan]

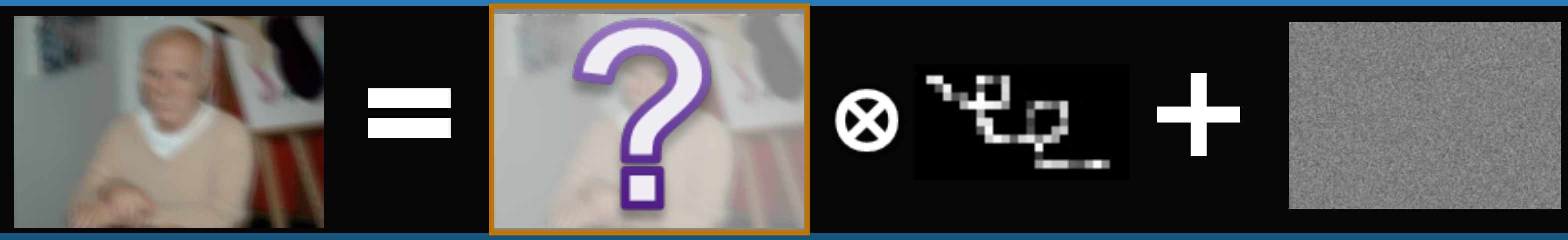

Time: about 30 seconds for an 800x600 image

Iteration 8
(converge)

[Slides by Qi Shan]

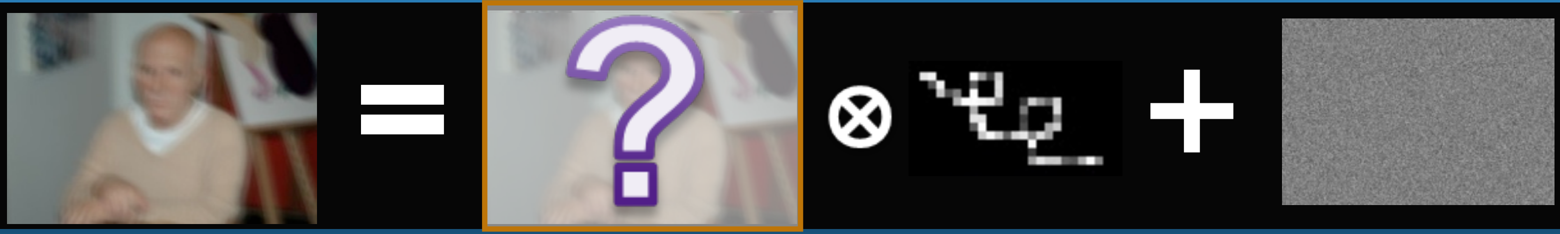

A comparison

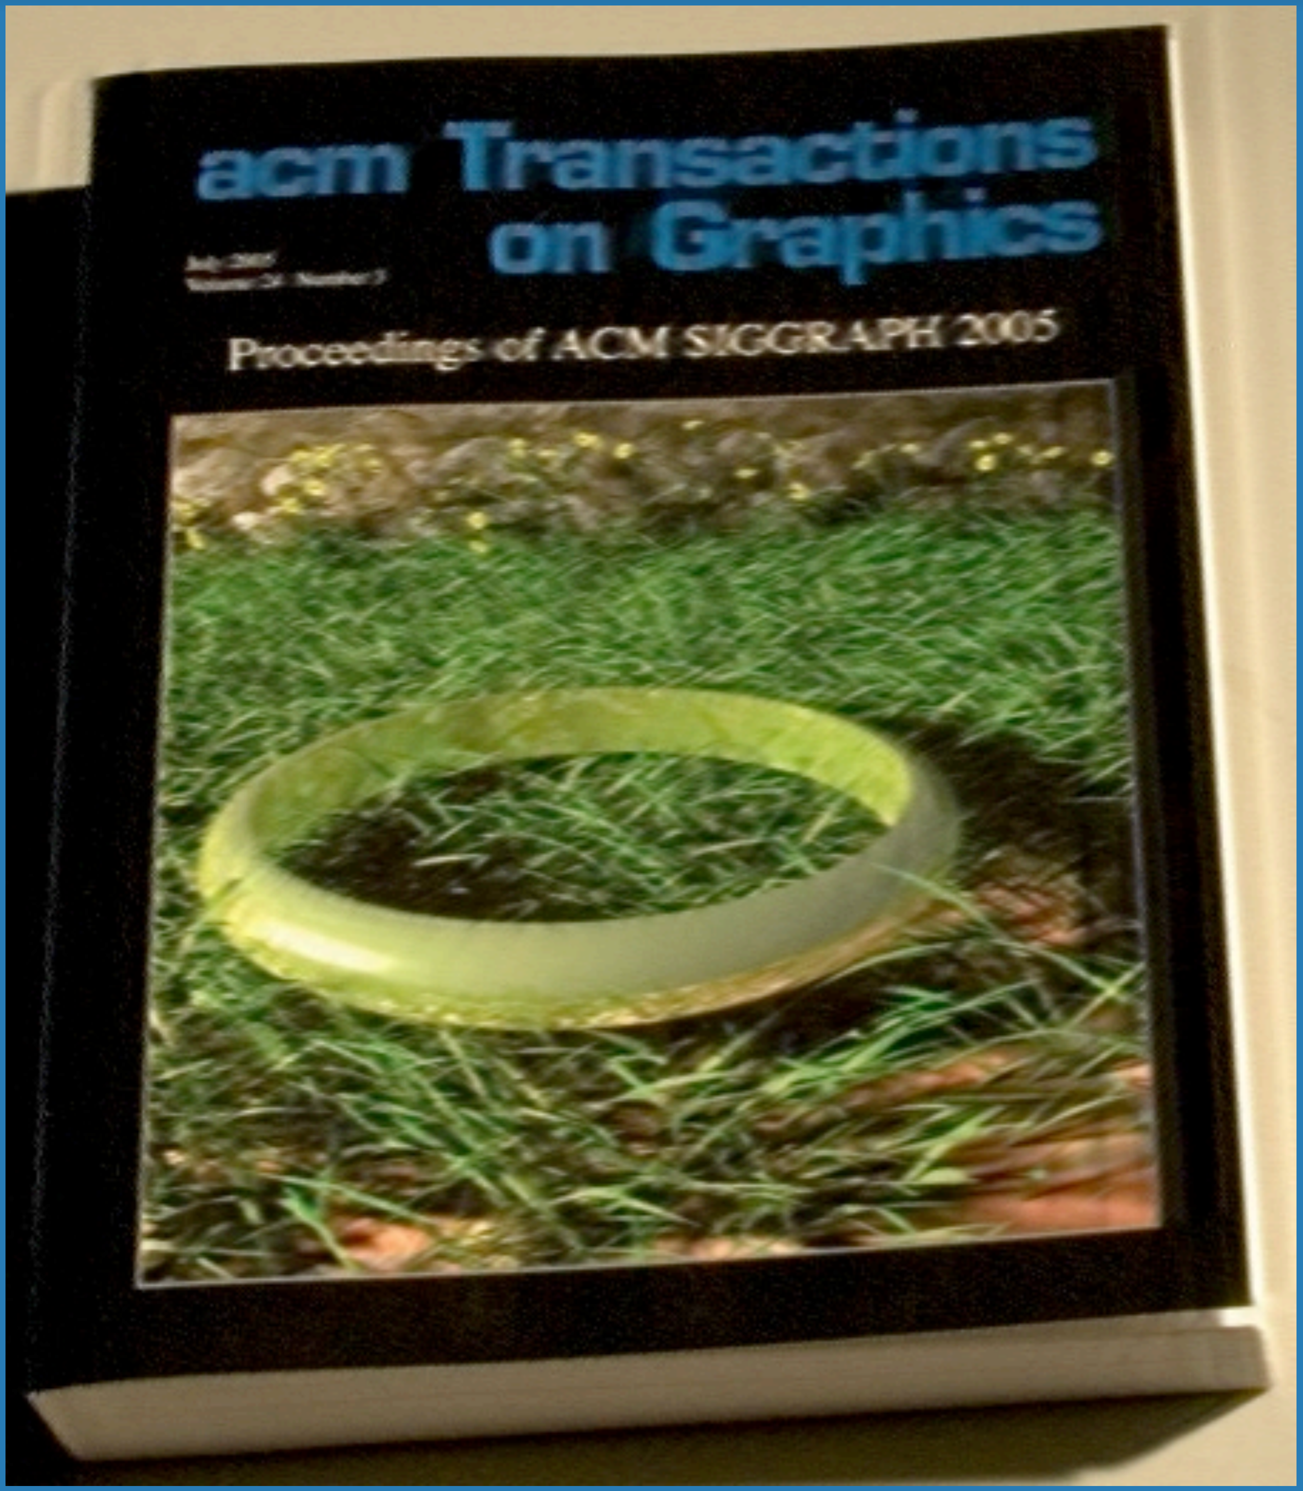

RL deconvolution

[Slides by Qi Shan]

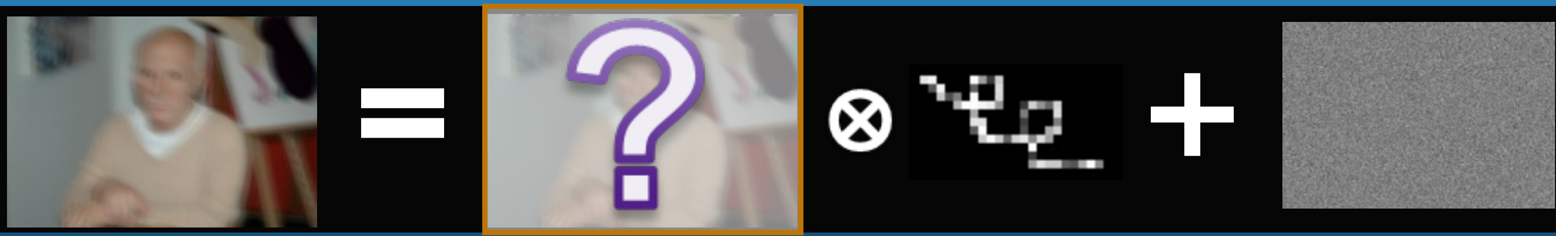

SIGGRAPH2008

A comparison

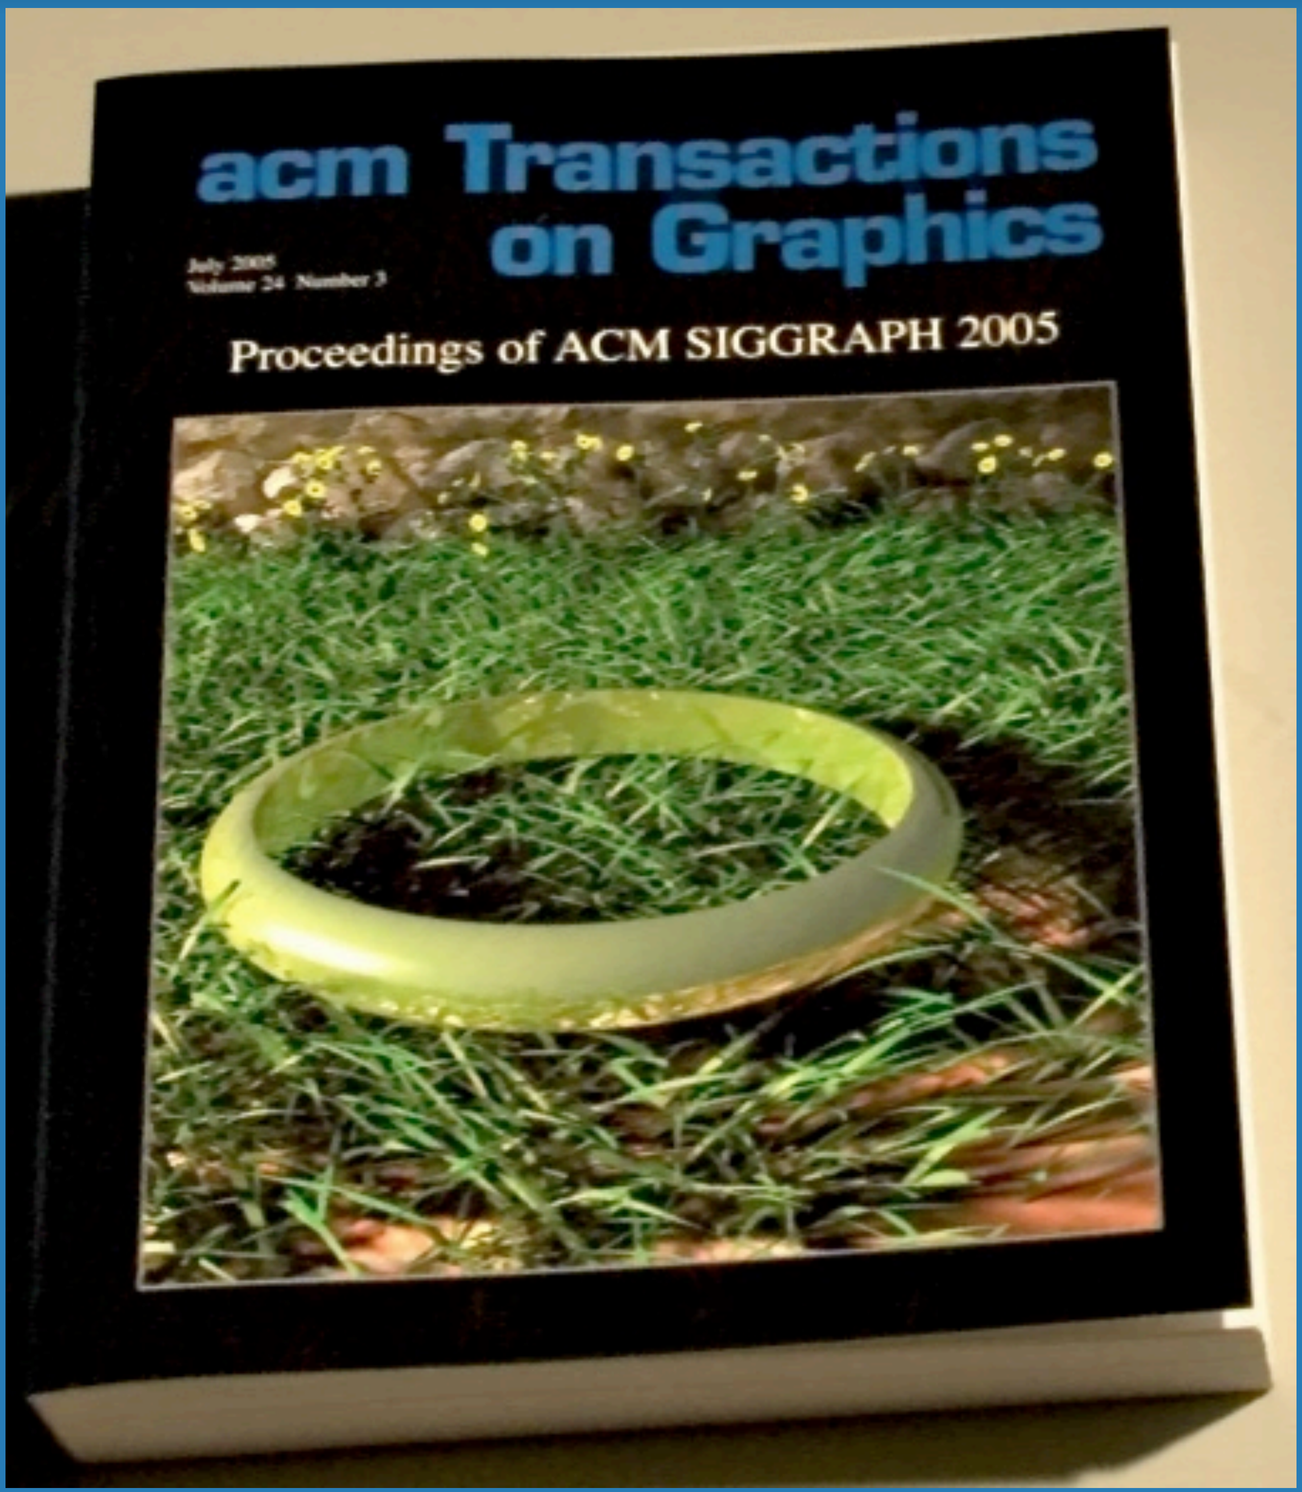

Our deconvolution

[Slides by Qi Shan]

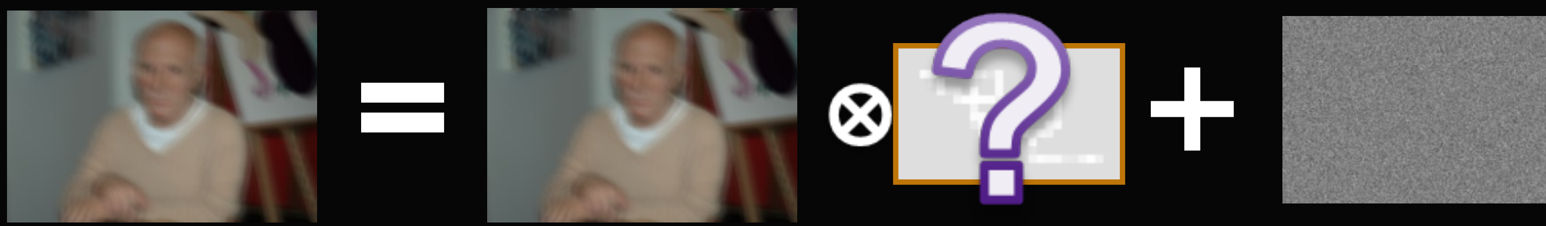

Two-step iterative optimization

- Optimize L
- Optimize f

$$\min E(L, f) = \min \log[p(n)\, p_1(\nabla L)\, p_2(L)\, p(f)]$$

$$E(f) = \left(\sum_{\nabla^*} w_{\nabla^*} \| \nabla^* L \otimes f - \nabla^* I \|_2^2 \right) + \| f \|_1$$

A form of L1-norm regularized problem and is solved using an interior point method

[Slides by Qi Shan]

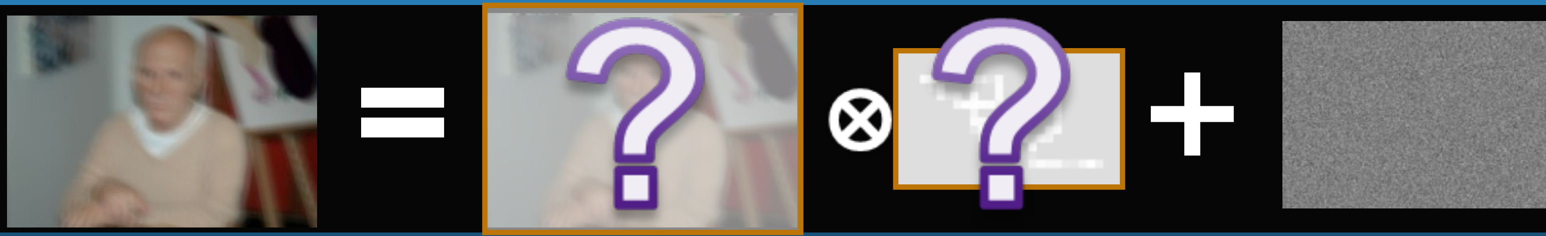

SIGGRAPH2008

[Slides by Qi Shan]

SIGGRAPH2008

[Slides by Qi Shan]

[Slides by Qi Shan]

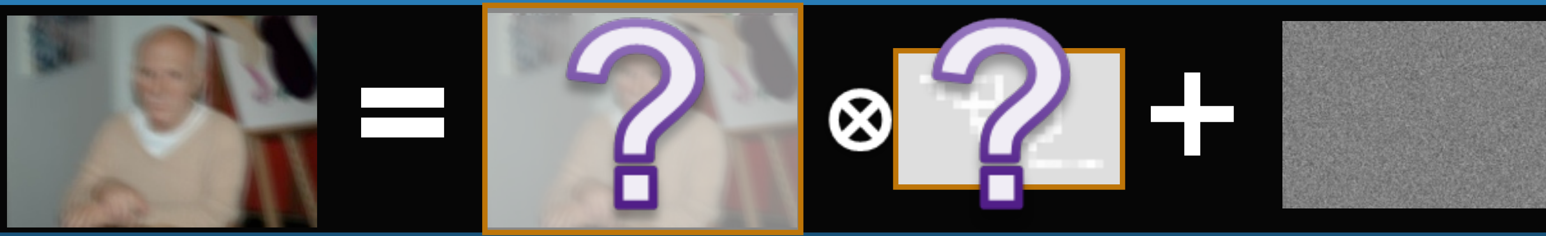

[Slides by Qi Shan]

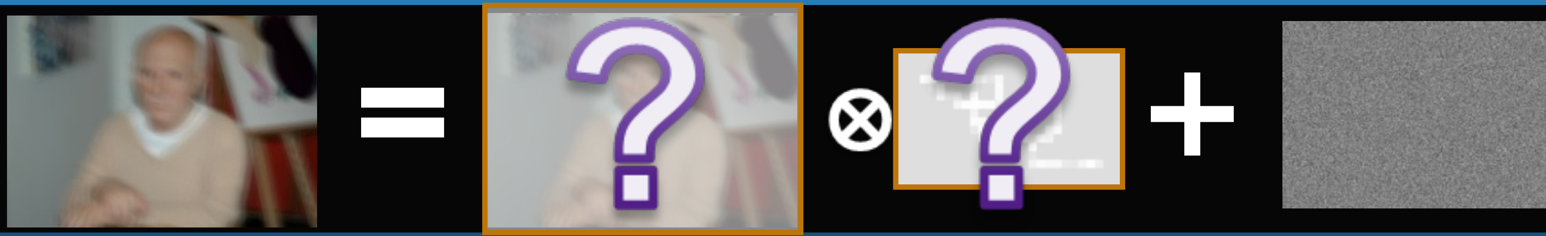

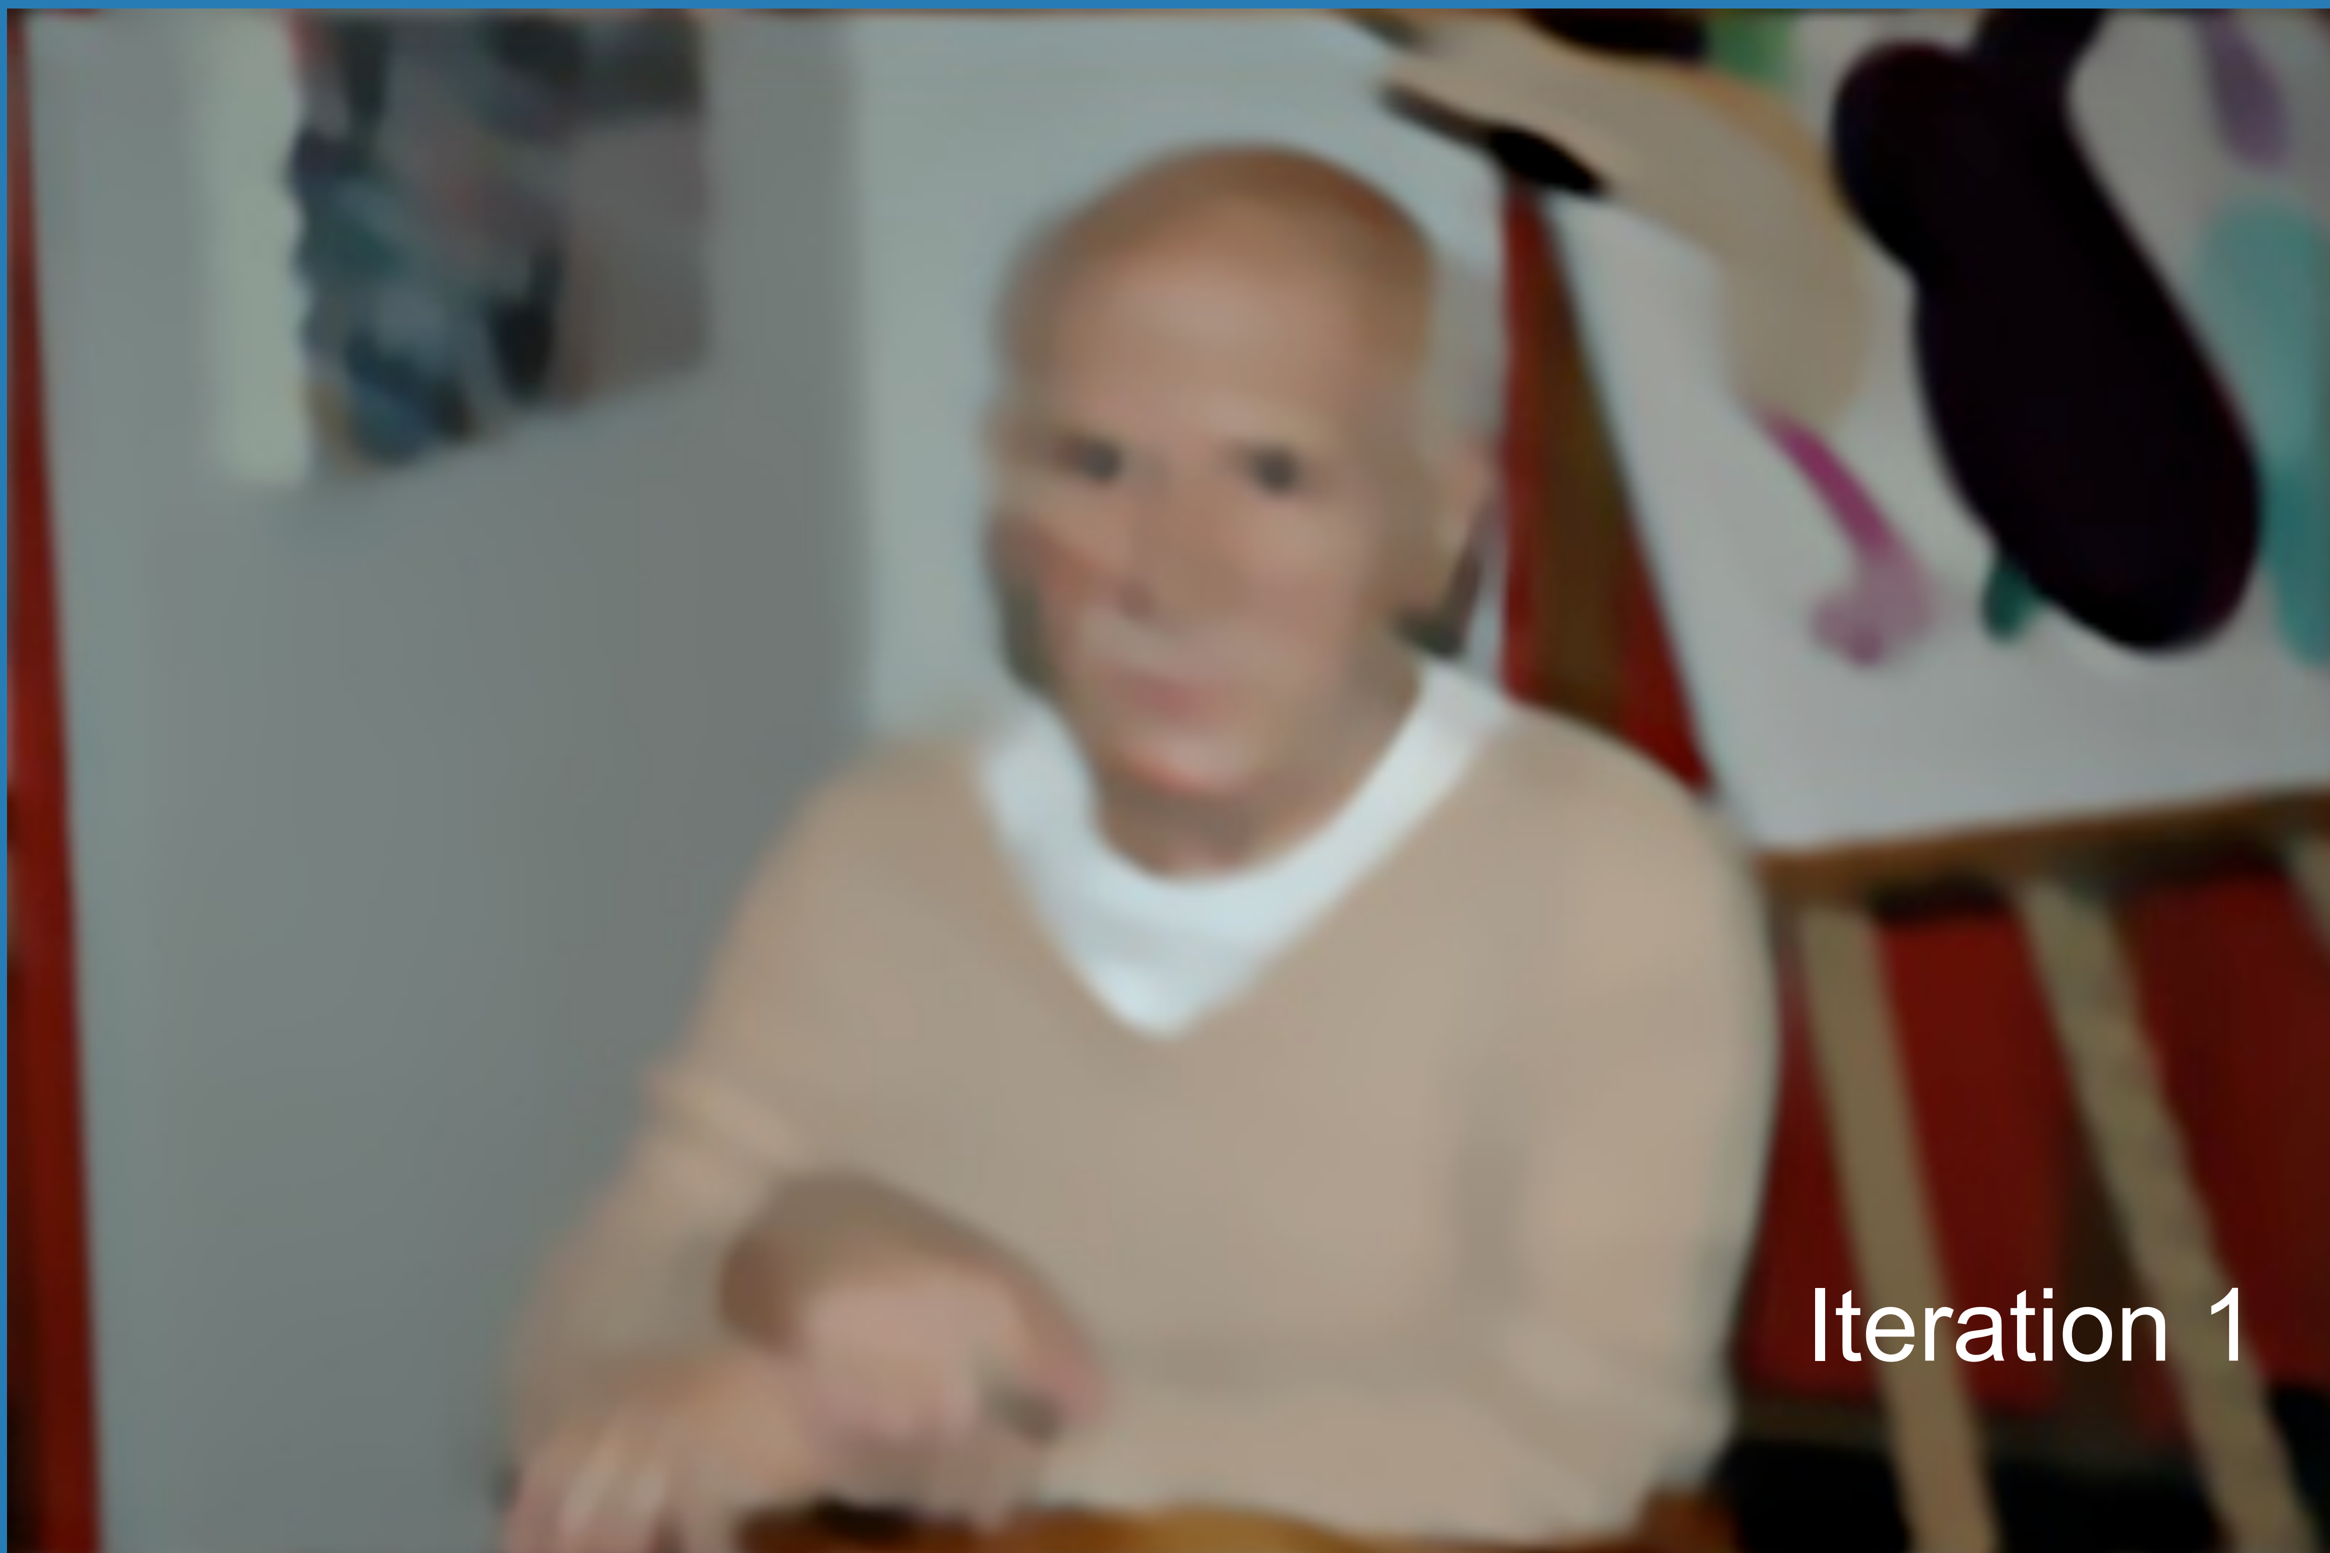

[Slides by Qi Shan]

SIGGRAPH2008

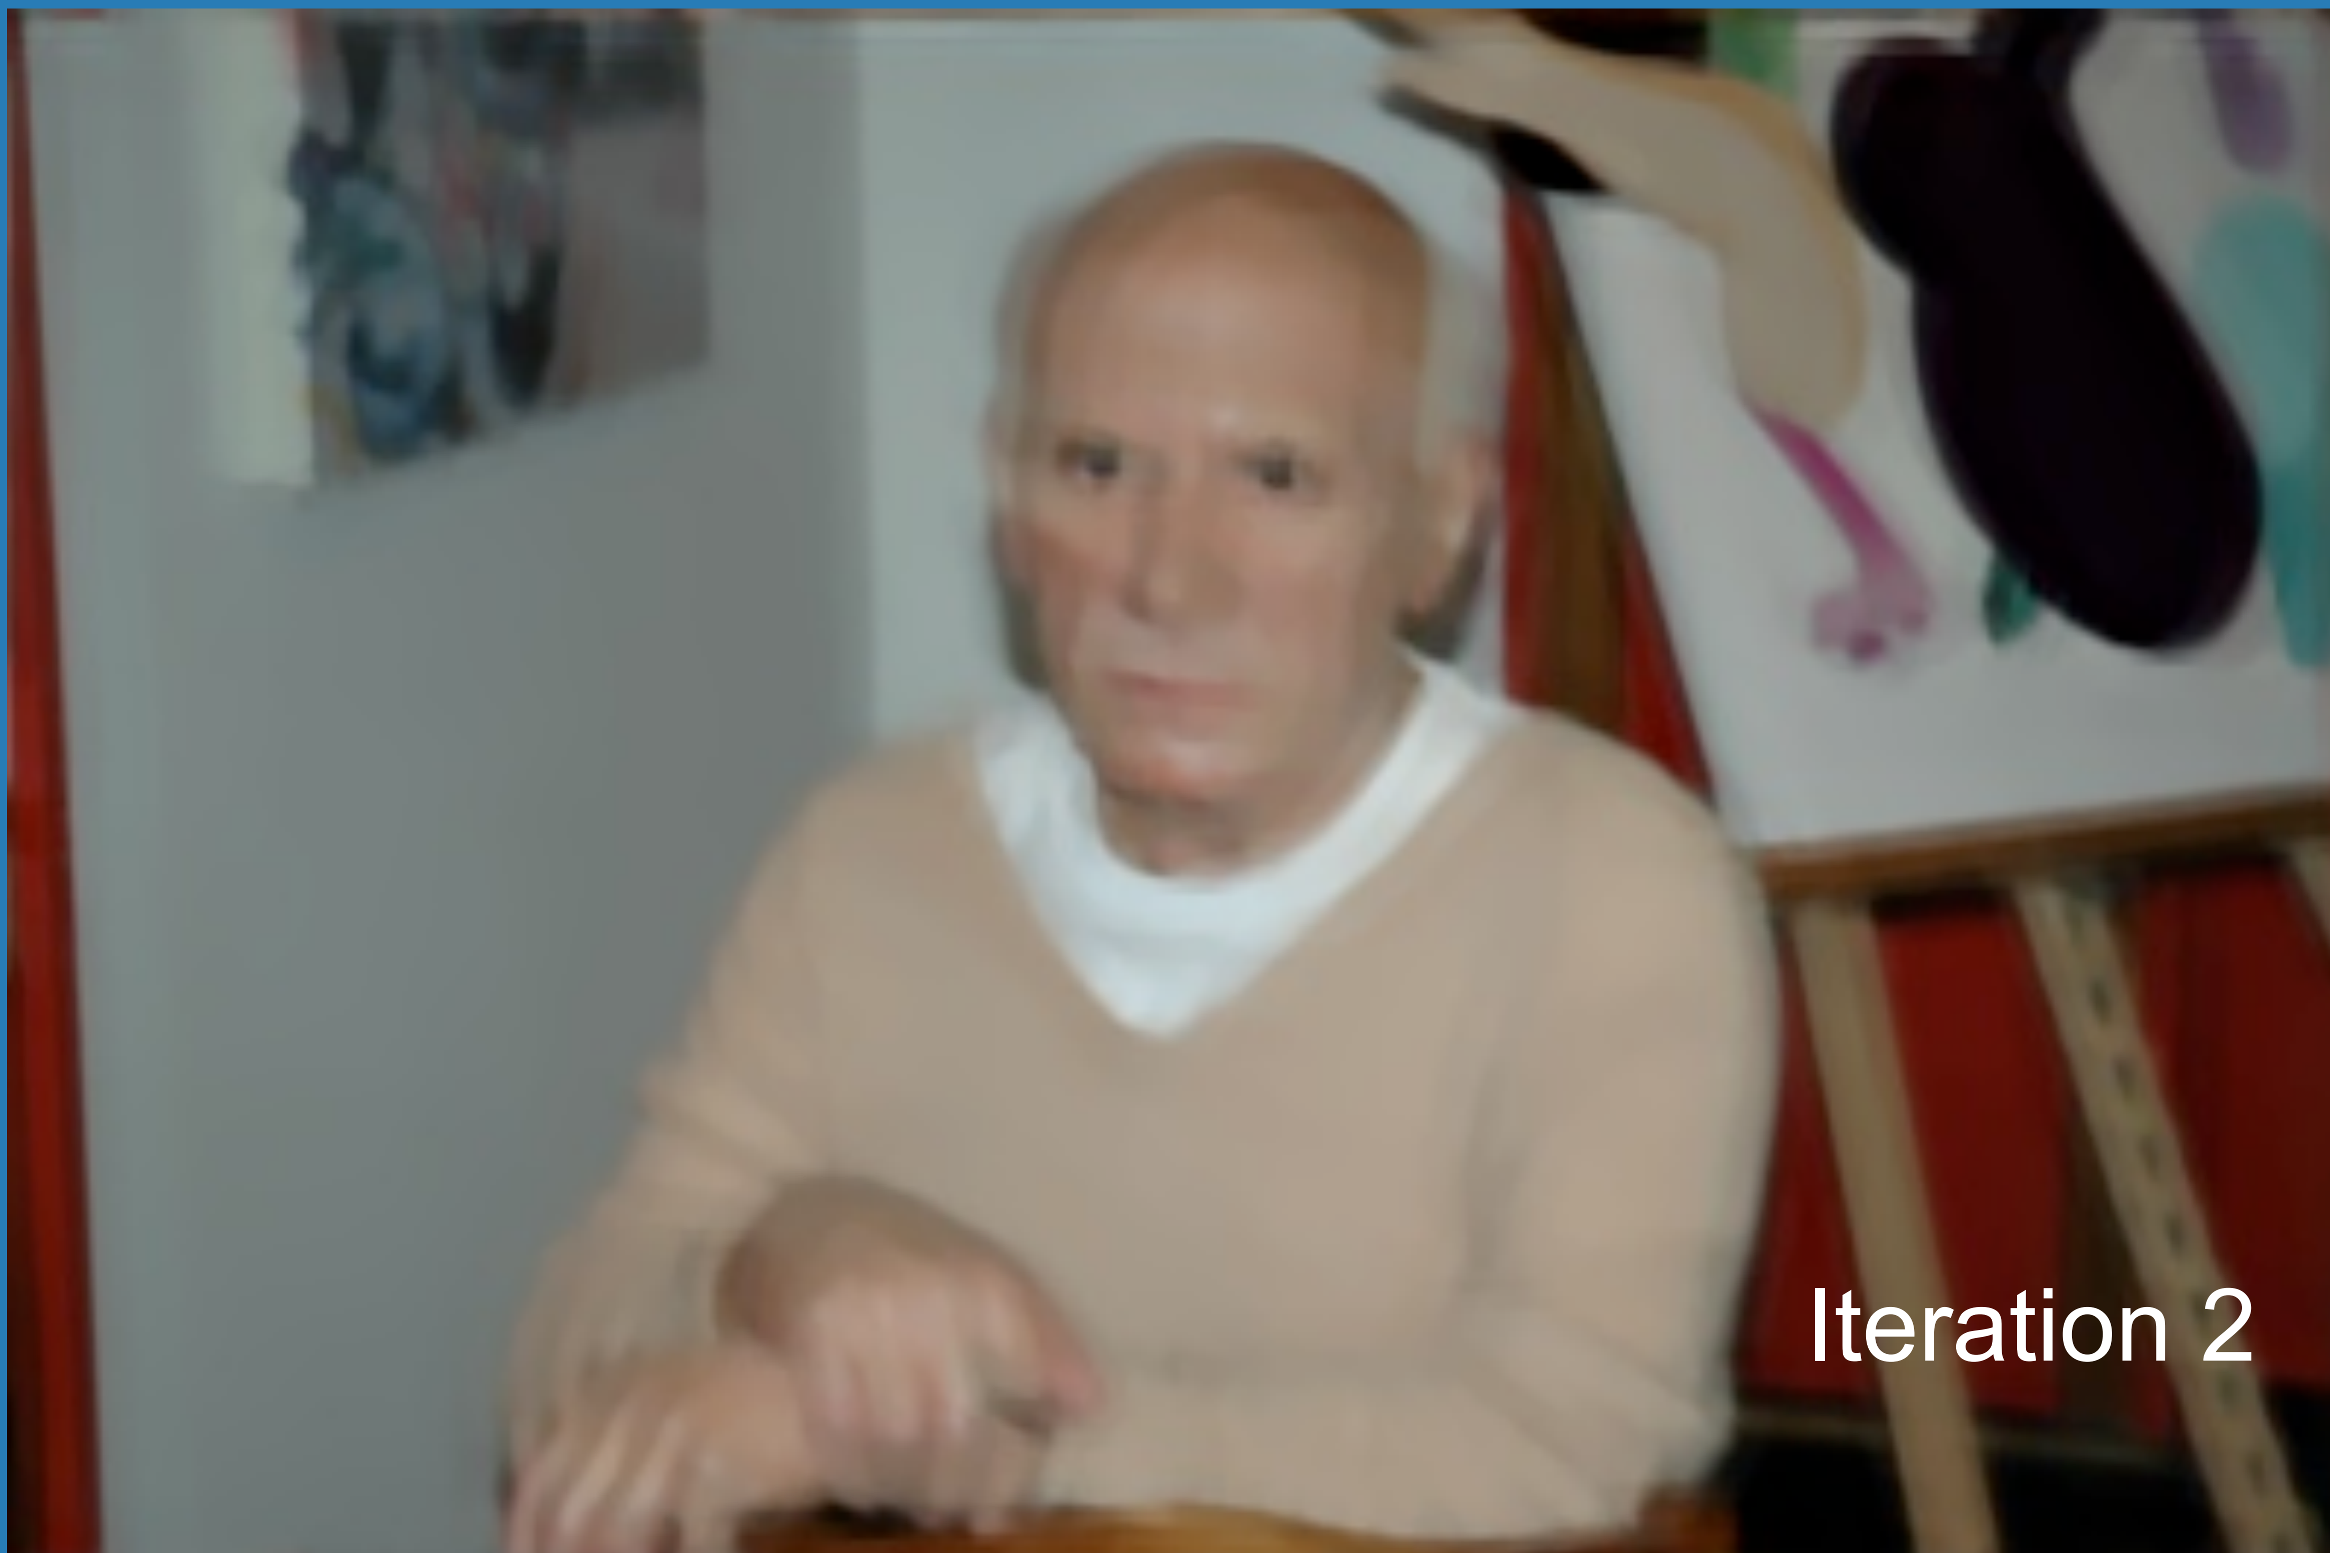

[Slides by Qi Shan]

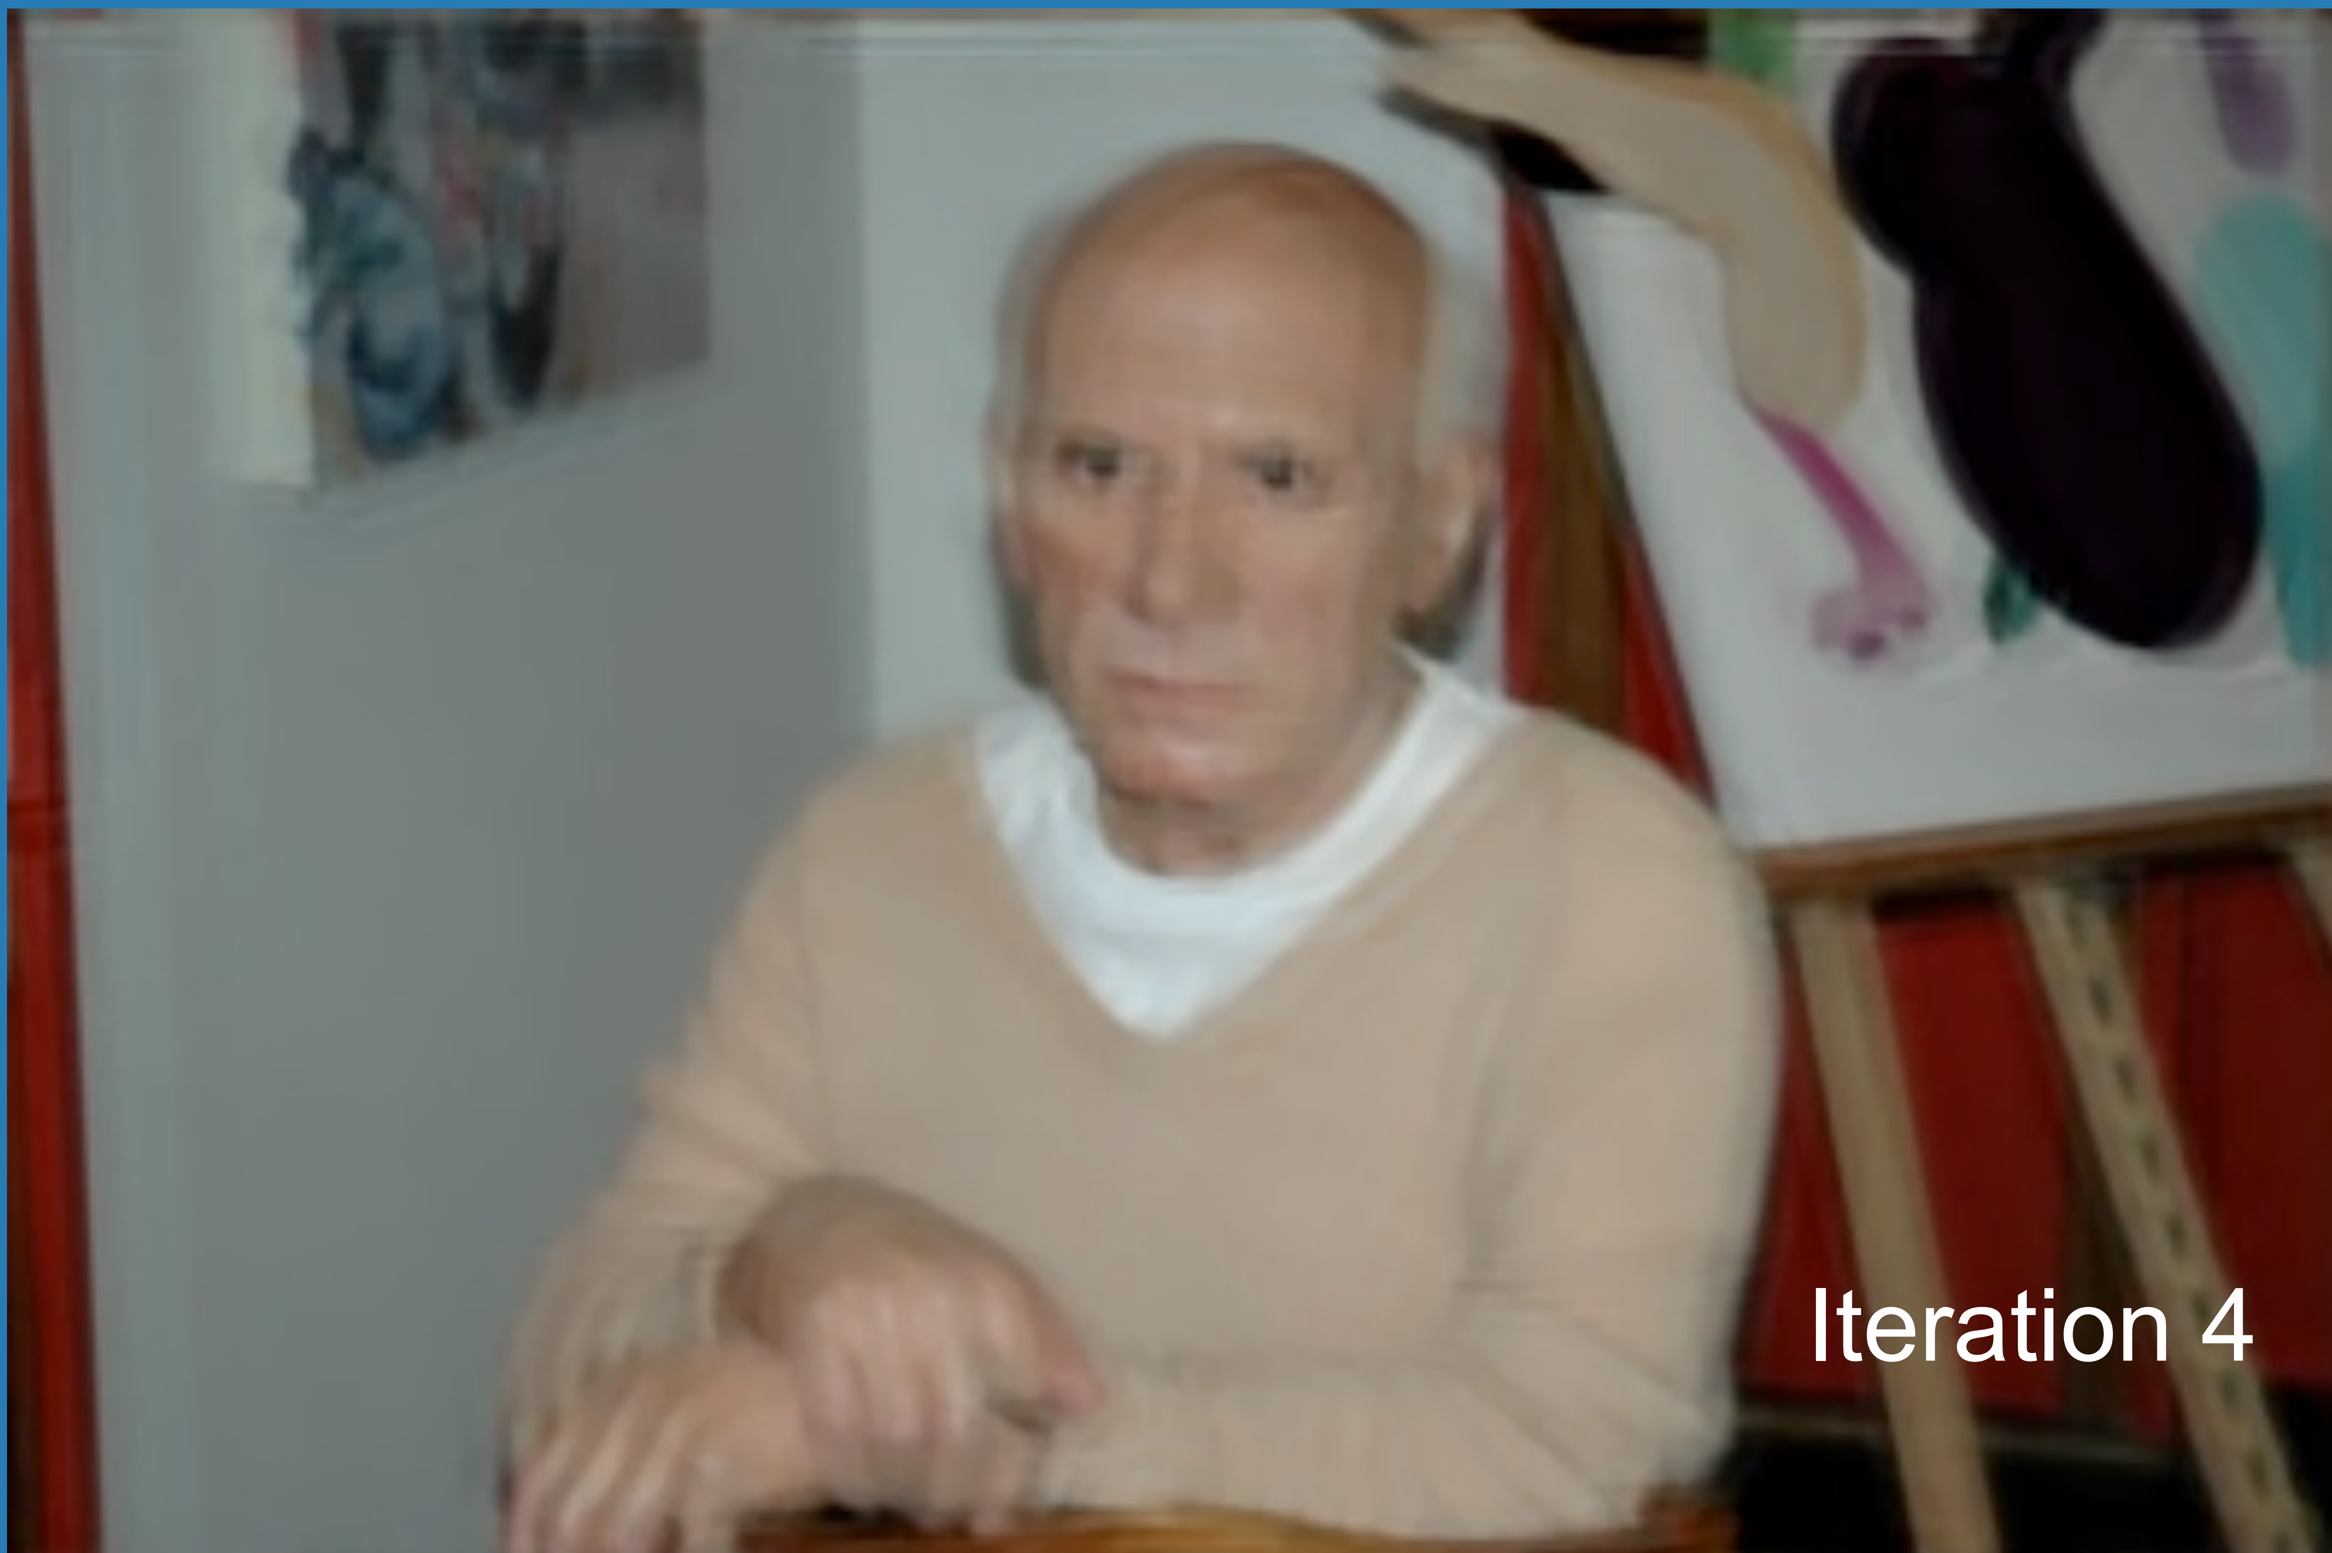

[Slides by Qi Shan]

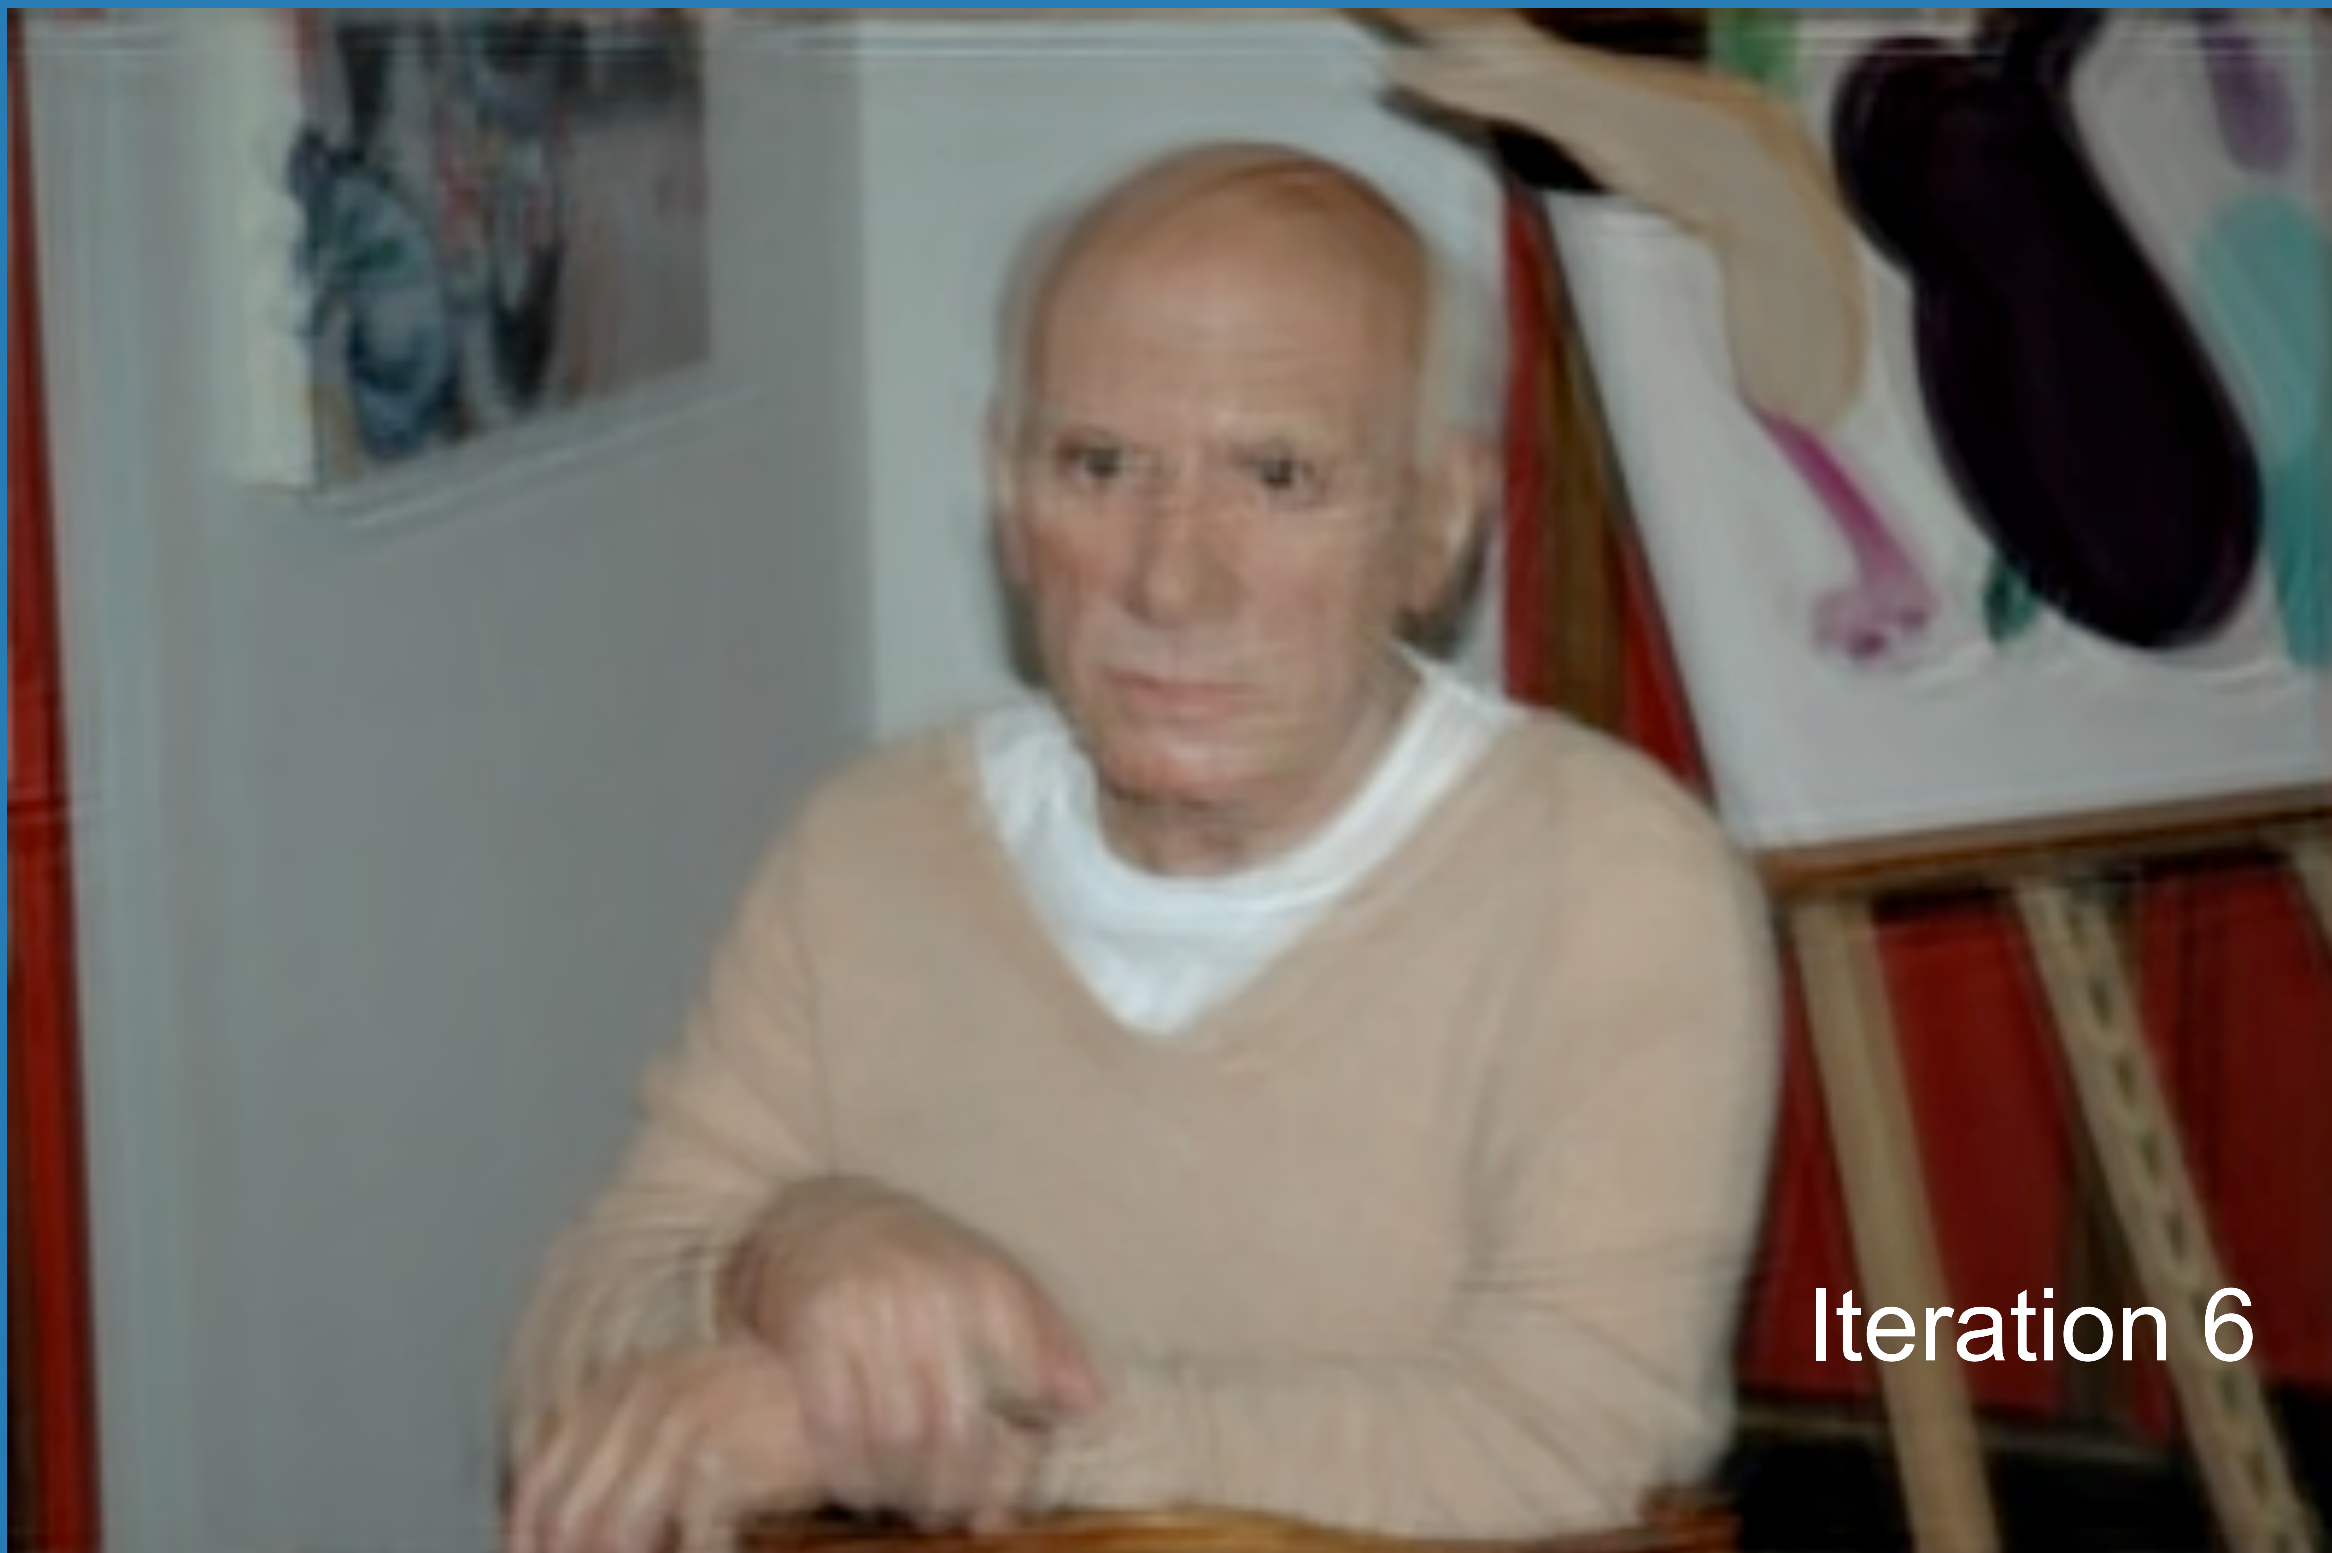

SIGGRAPH2008

[Slides by Qi Shan]

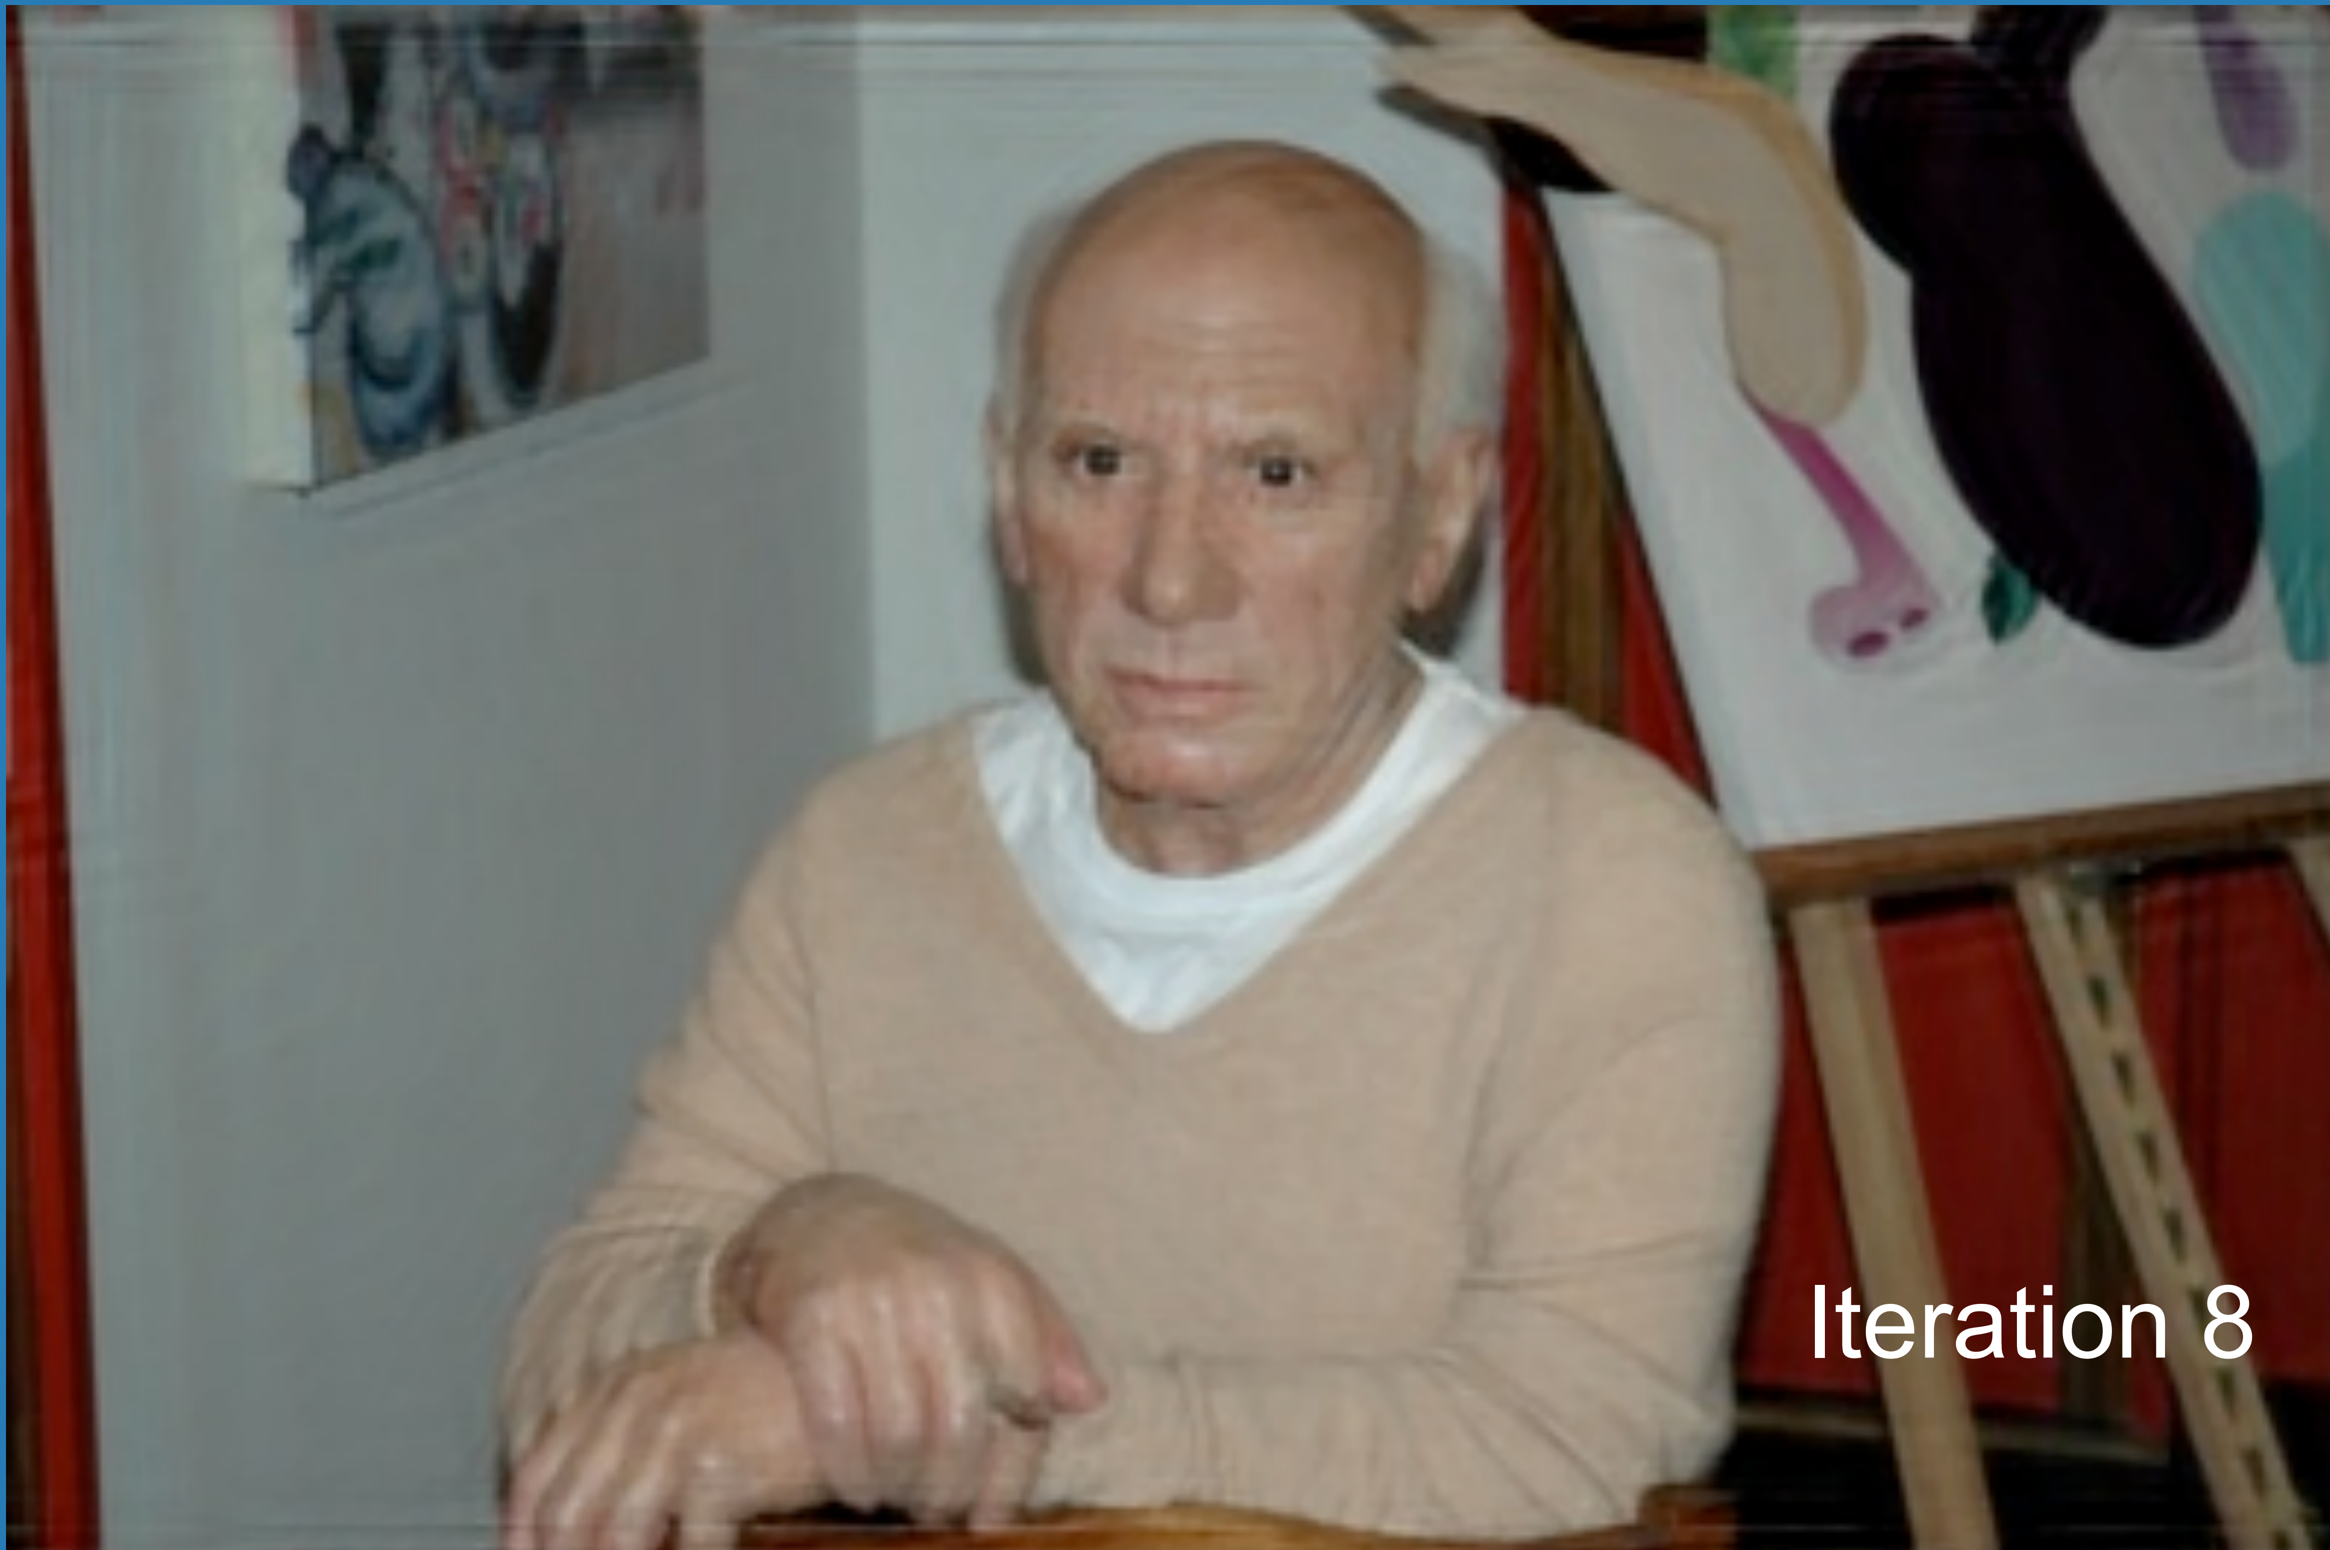

[Slides by Qi Shan]

Time: about 350 seconds for an 800x600 image

Convergence

[Slides by Qi Shan]

SIGGRAPH2008

Results

[Slides by Qi Shan]

SIGGRAPH2008

Results

[Slides by Qi Shan]

Results

[Slides by Qi Shan]

Results

[Slides by Qi Shan]

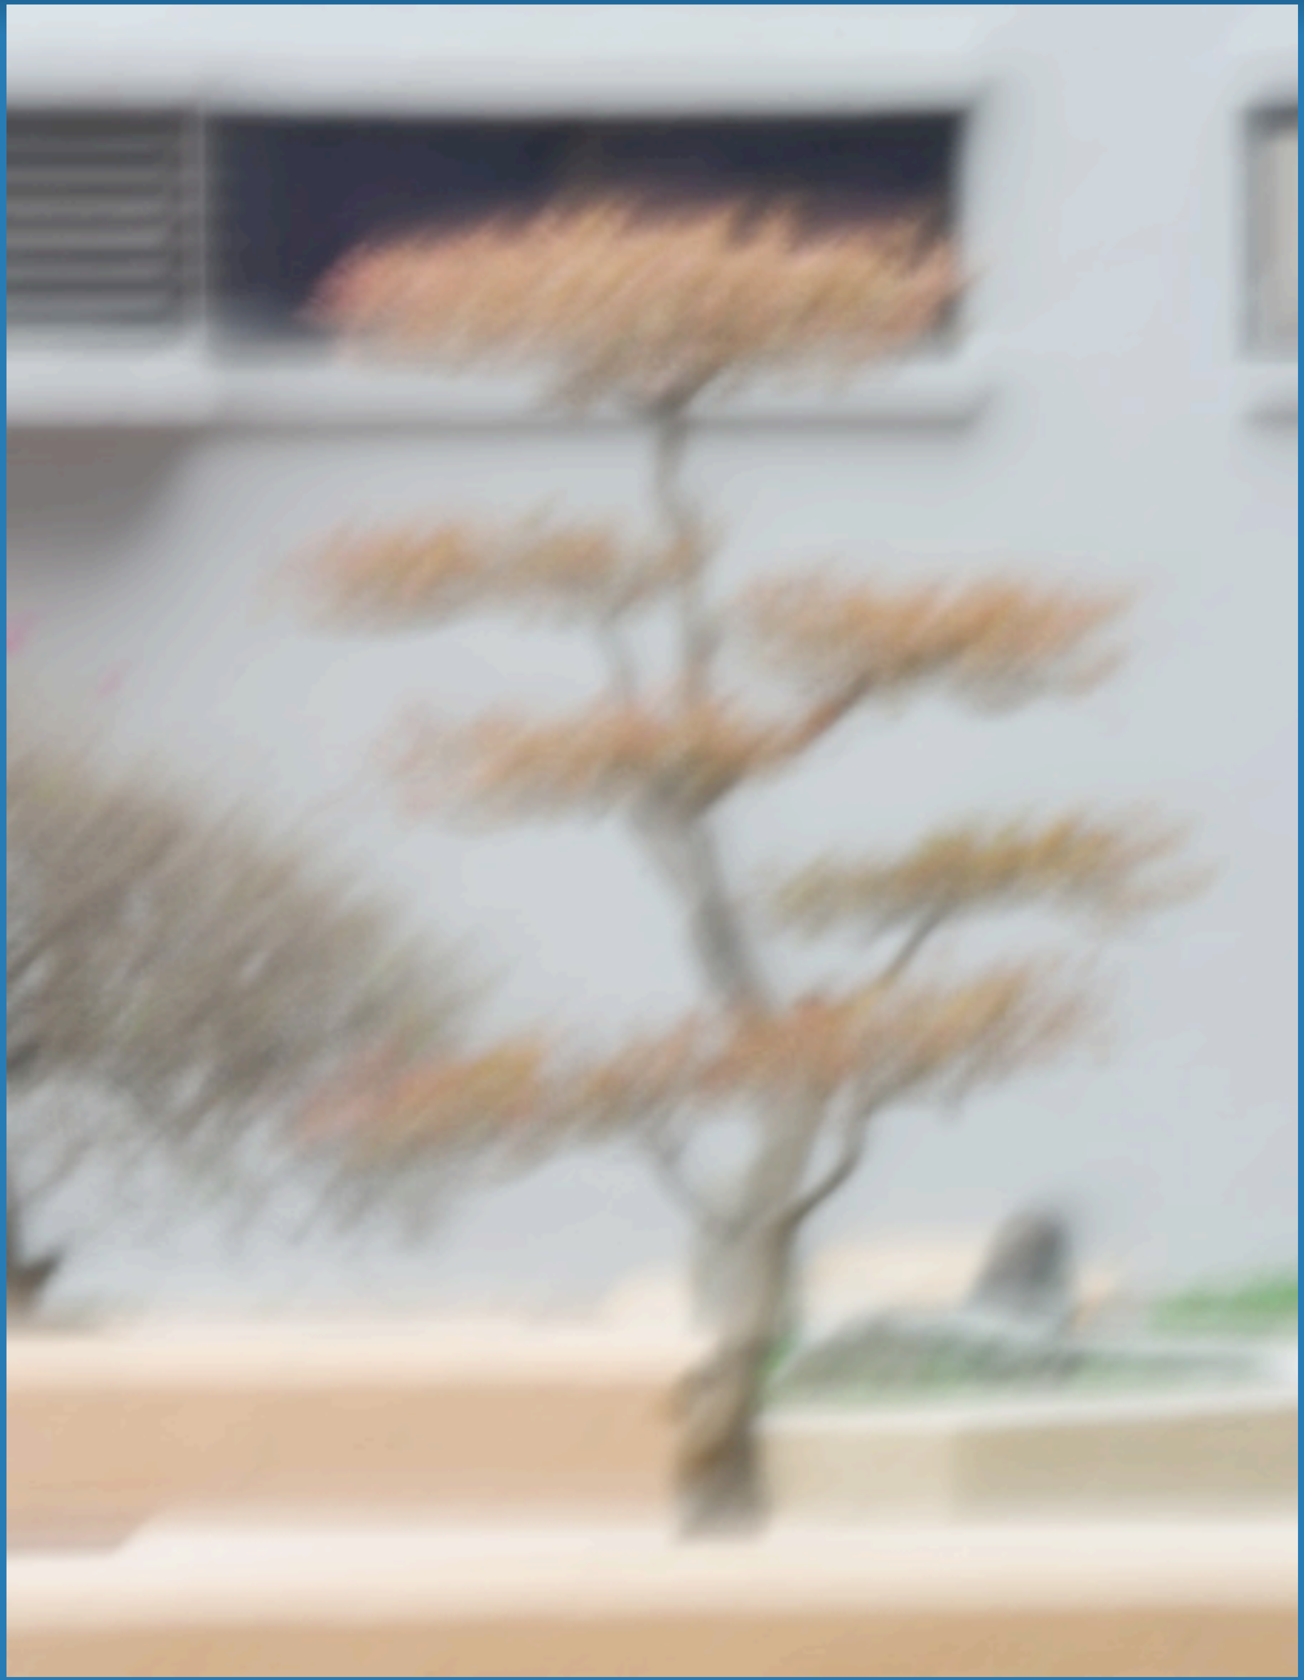

[Slides by Qi Shan]

[Slides by Qi Shan]

More results

[Slides by Qi Shan]

More results

[Slides by Qi Shan]

More results

[Slides by Qi Shan]

More results

[Slides by Qi Shan]

More results

[Slides by Qi Shan]

More results

[Slides by Qi Shan]

SIGGRAPH2008

More results

[Slides by Qi Shan]

More results

[Slides by Qi Shan]

Joshi et al.

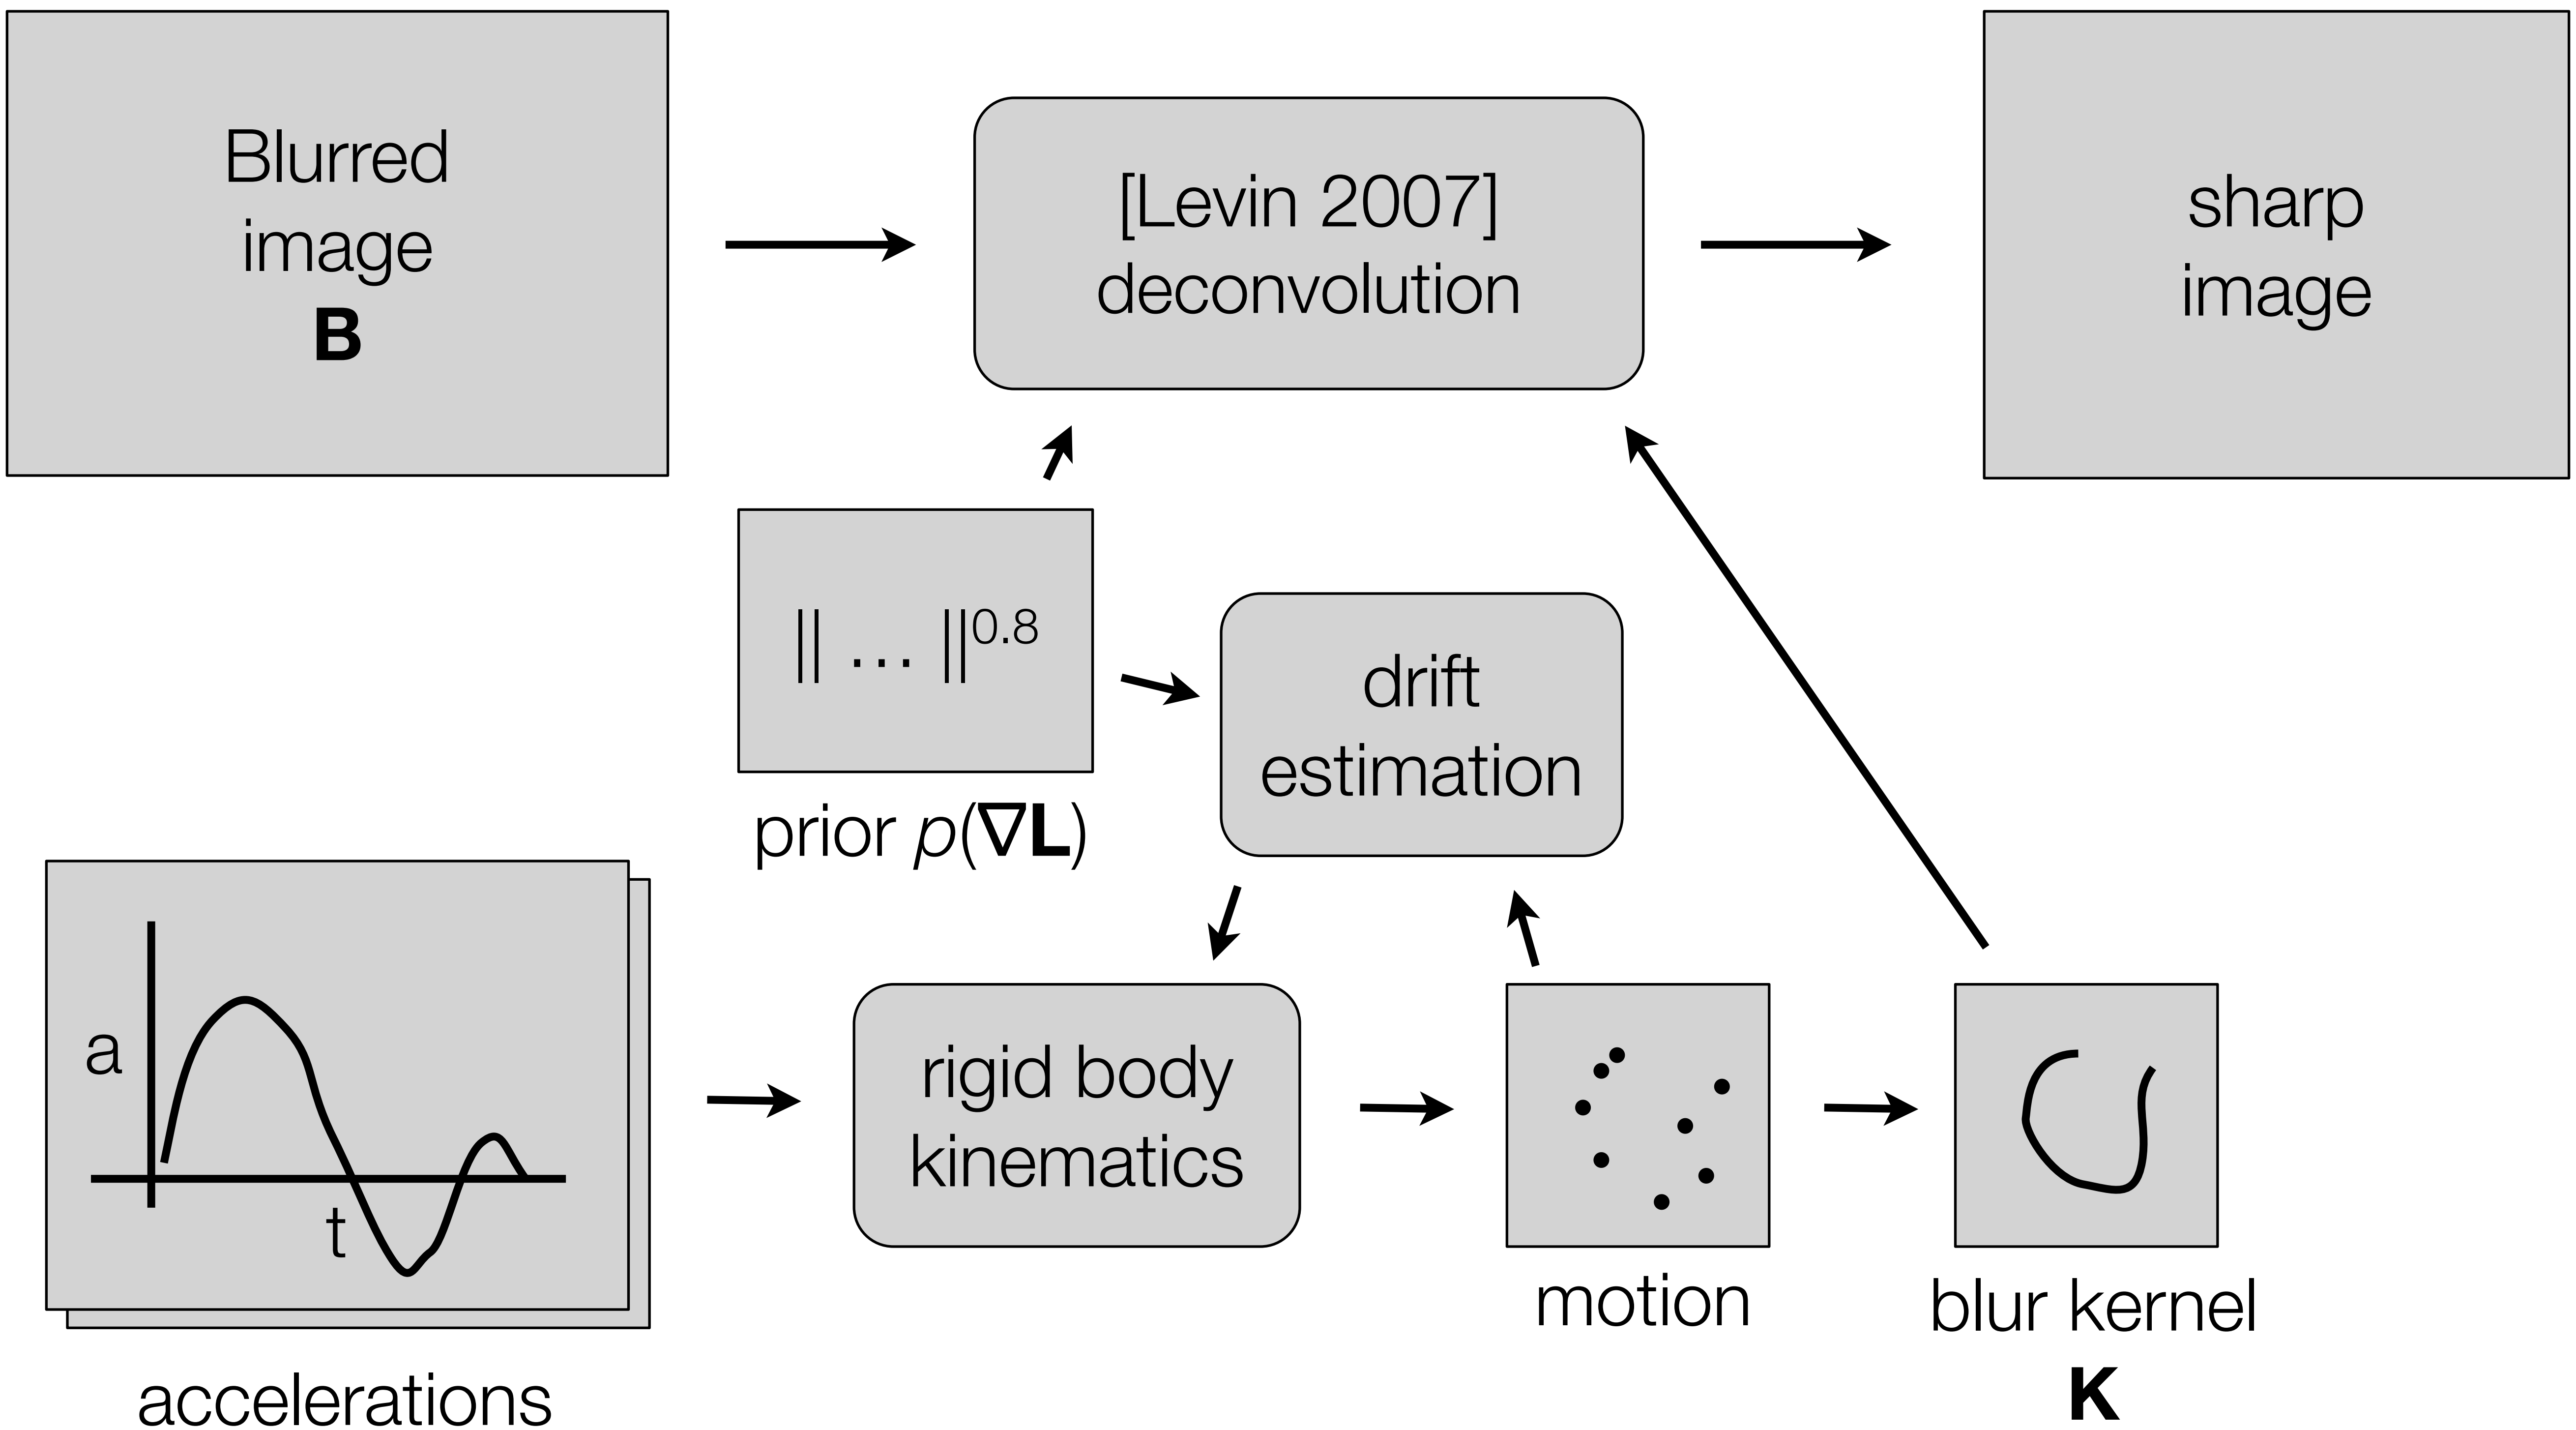

Blurred image **B**
[Levin 2007] deconvolution
sharp image
|| … ||0.8
prior p(∇L)
drift estimation
a
t
accelerations
rigid body kinematics
motion
blur kernel **K**

System Overview

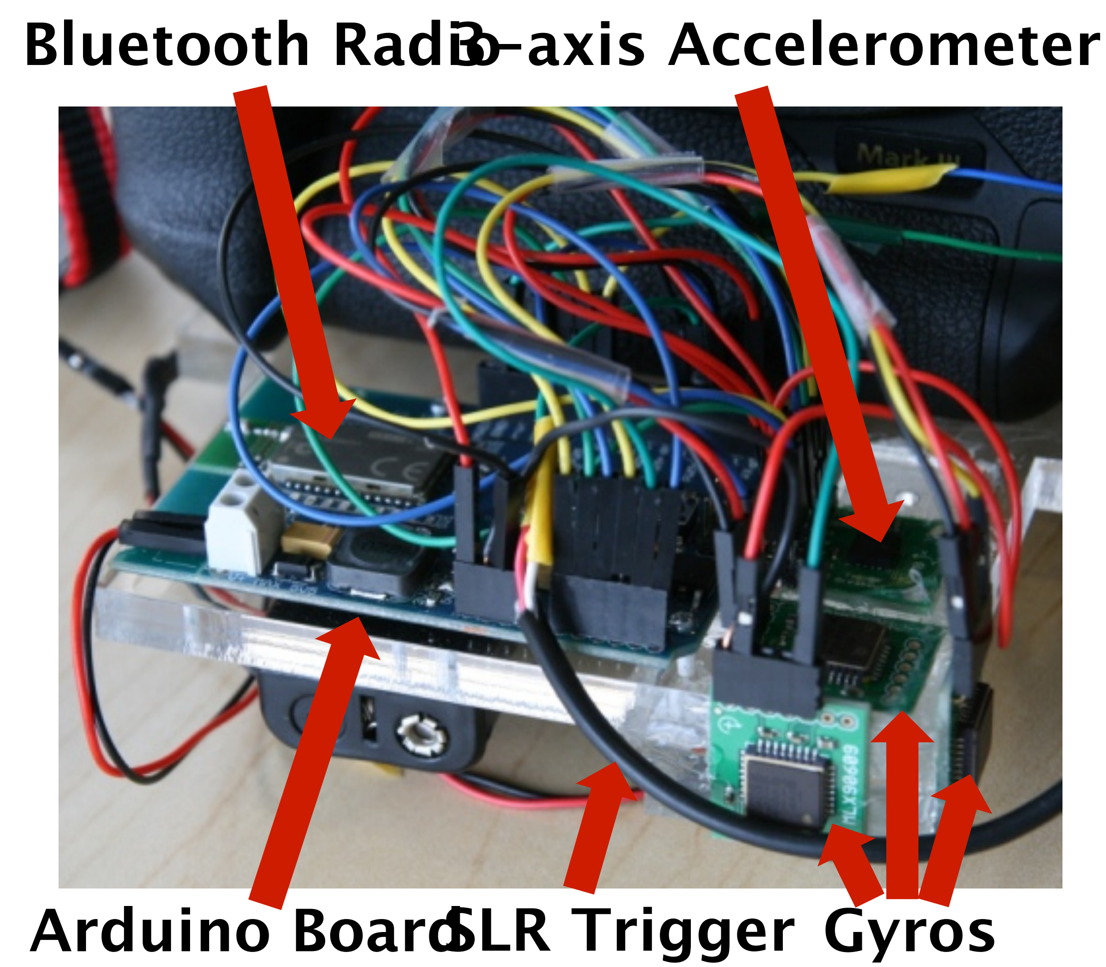

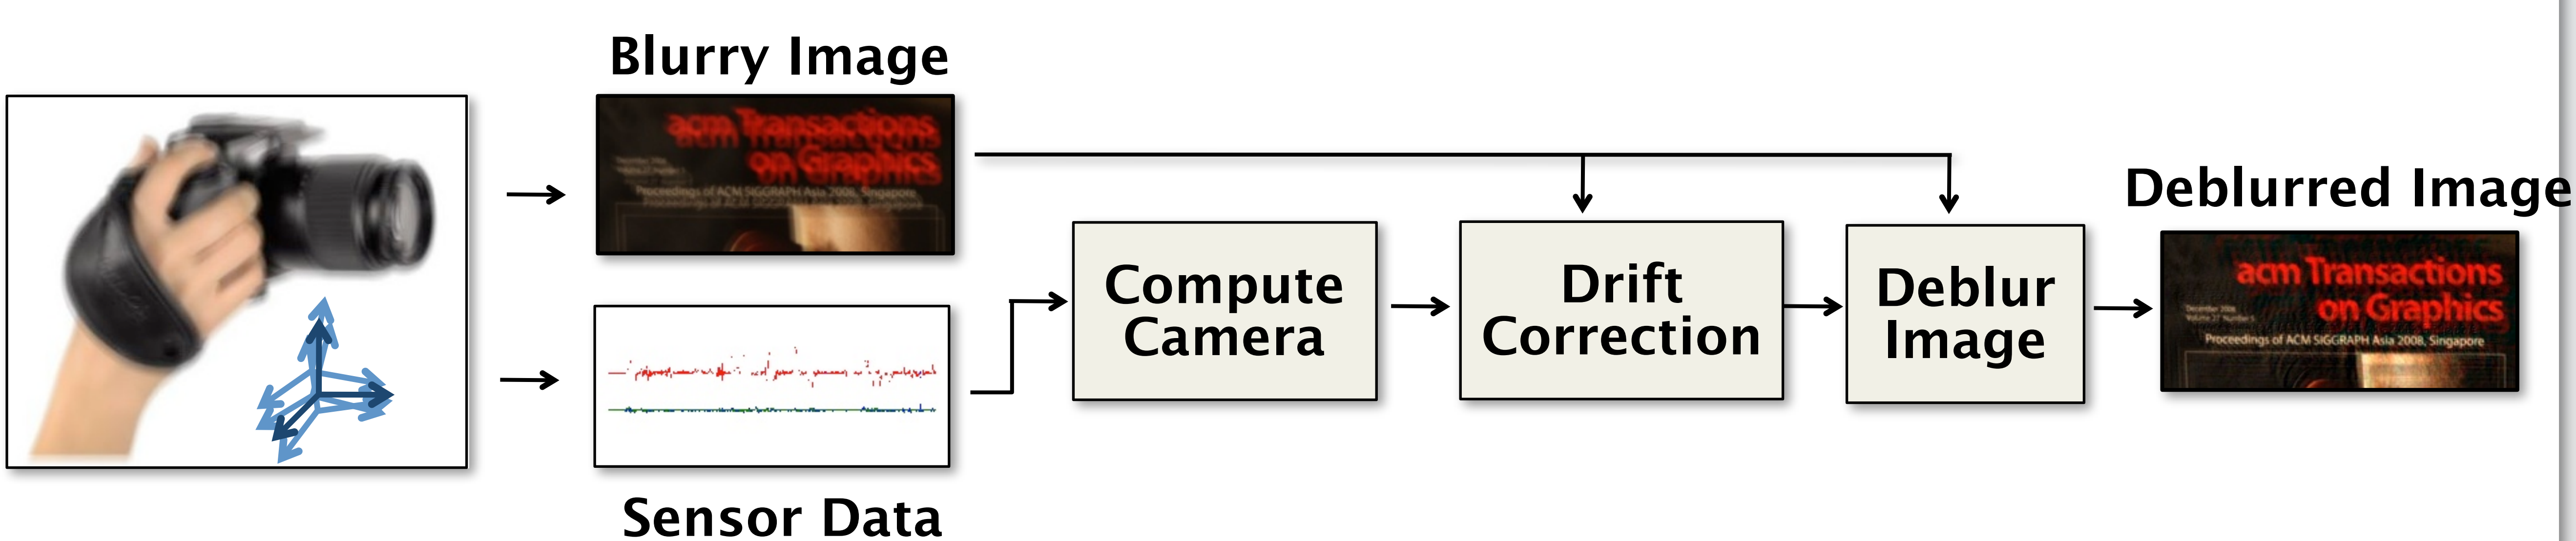

Camera and Blur

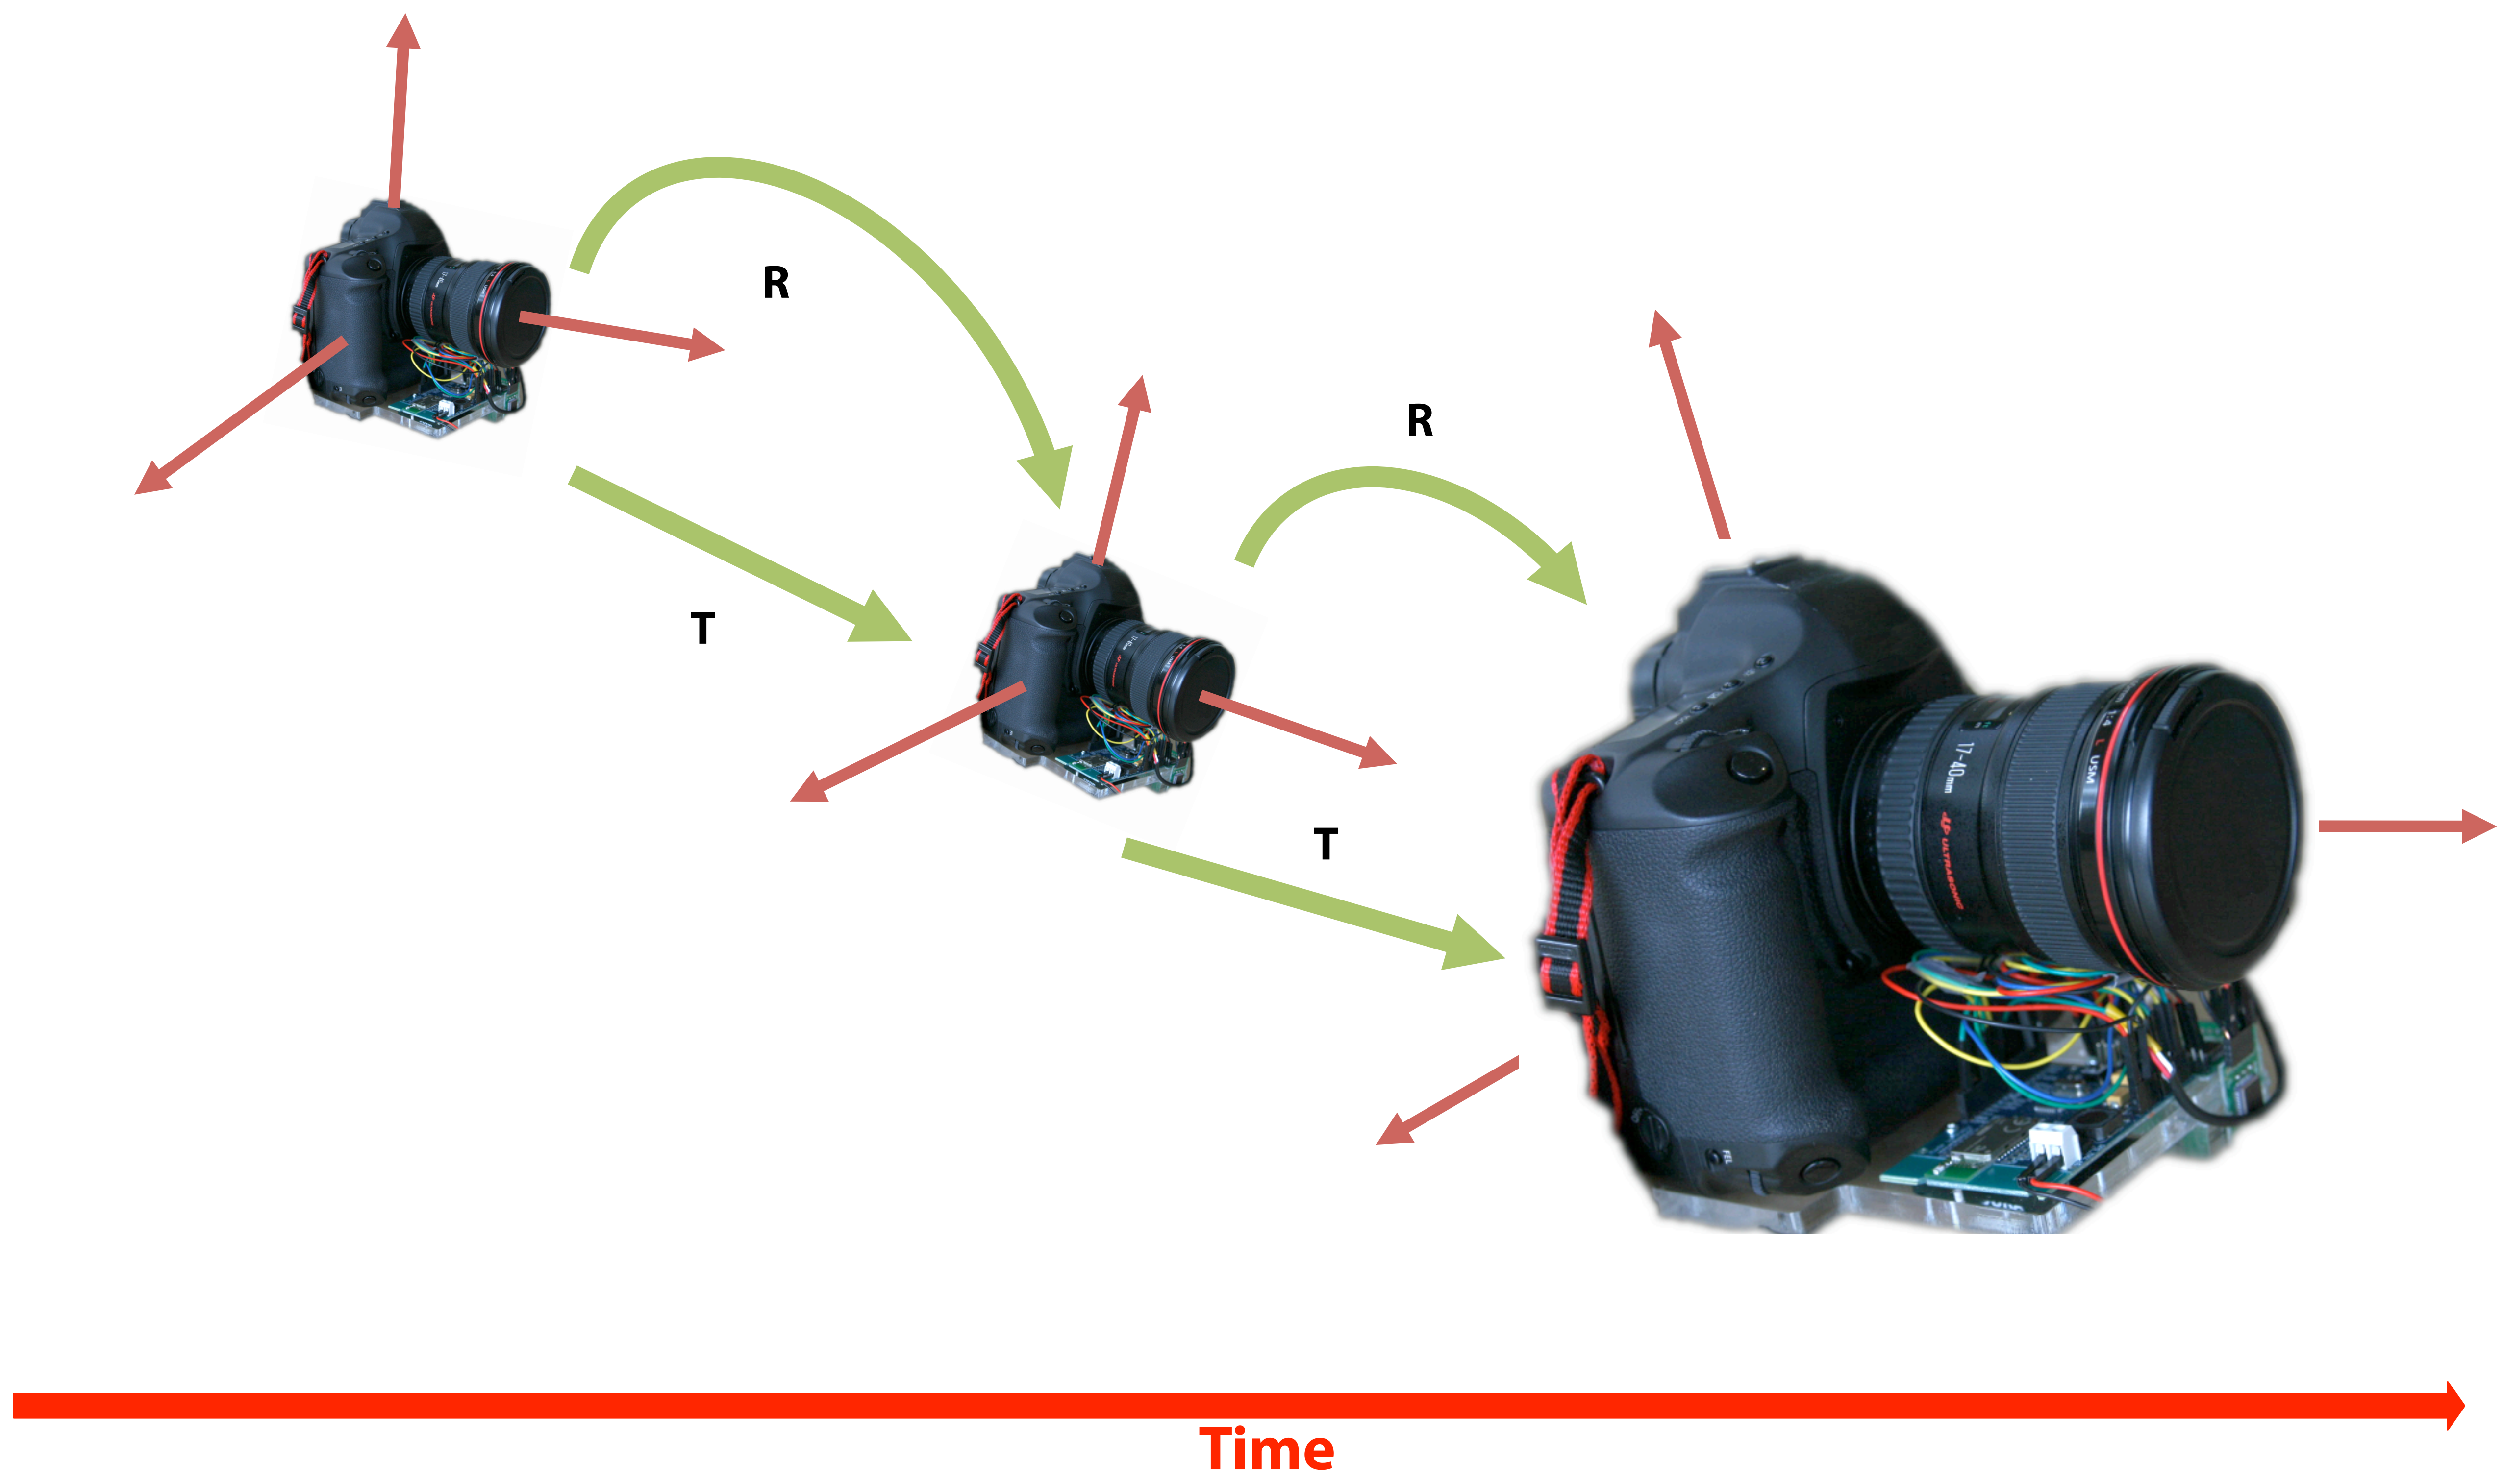

Neel Joshi, SIGGRAPH 2010

Recovering Motion from Inertial Sensors

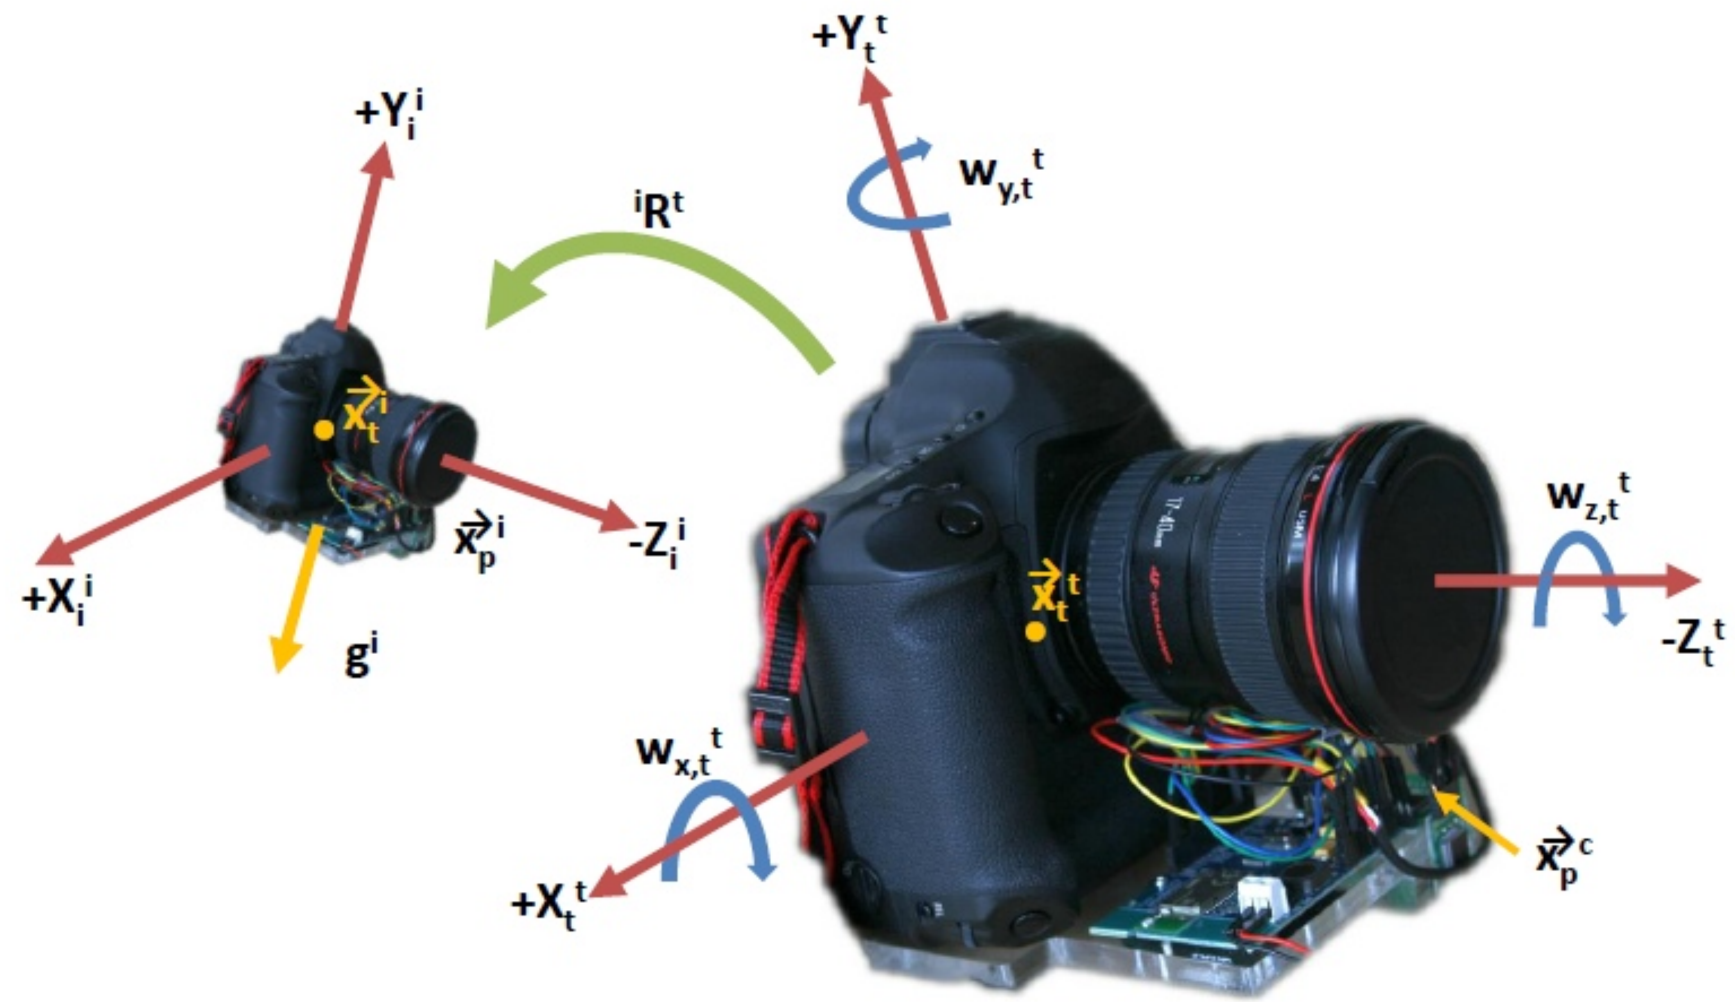

$$\vec{\omega}_t^t = {}^tR^i * \vec{\omega}_t^i$$

$$\vec{a}_p^t = {}^tR^i \left(\vec{a}_t^i + \vec{g}^i + (\vec{\omega}_t^i \times (\vec{\omega}_t^i \times \vec{r}_p^q)) + (\vec{\alpha}_t^i \times \vec{r}_p^q)\right)$$

Measured by accelerometers and gyros

Integrate to Recover Camera Rotation/Translation

Spatially Varying Deblurring

Blurry

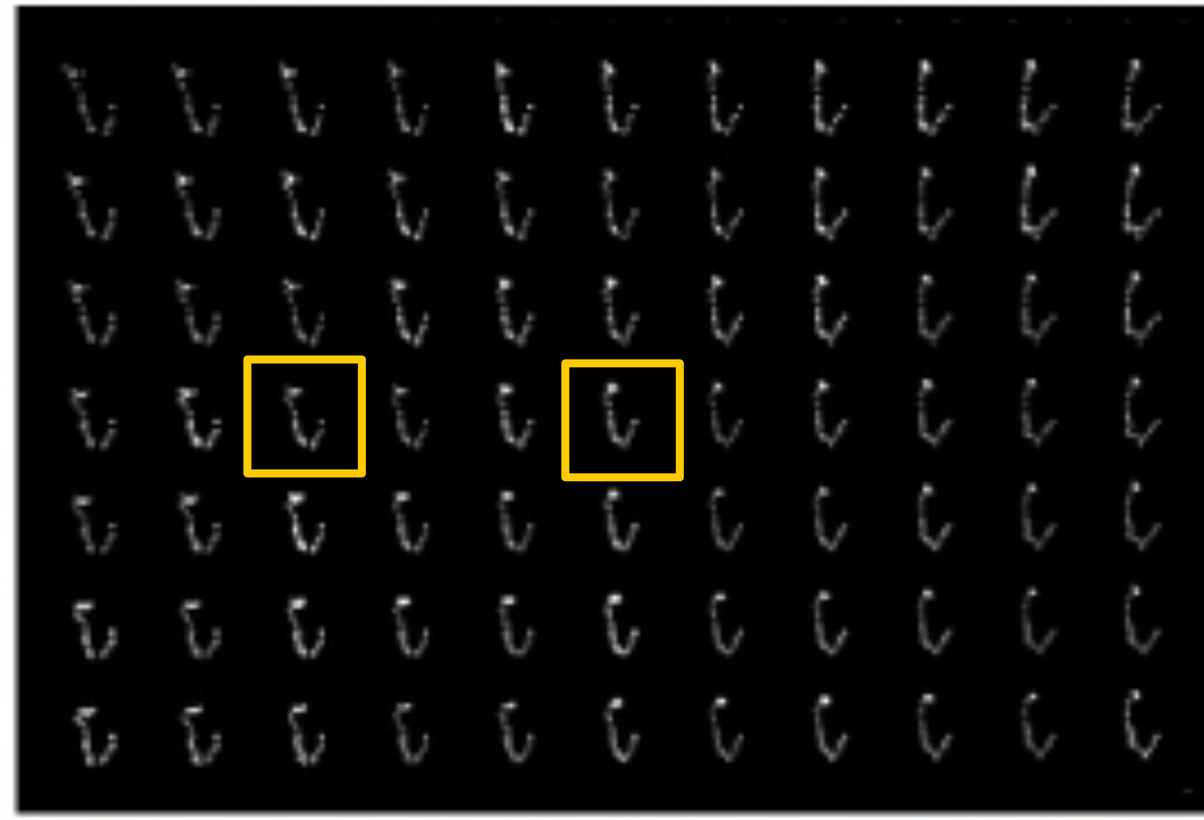

Spatially-Varying Kernels
(Single Depth Plane)

Deblurred Using Correct Kernel

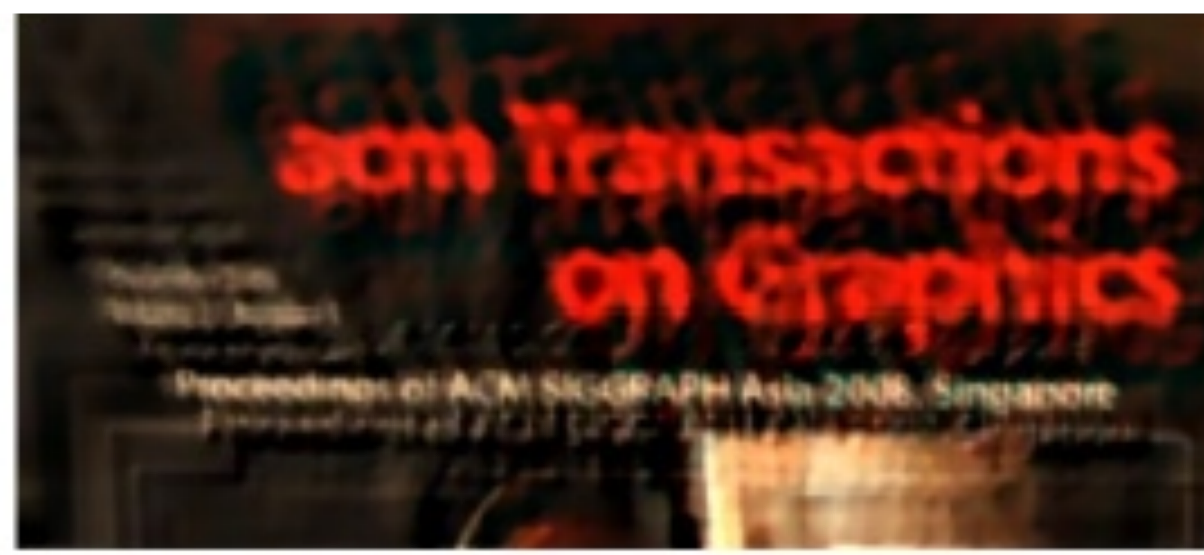

Deblurred Using Center Kernel

How accurate are the sensors

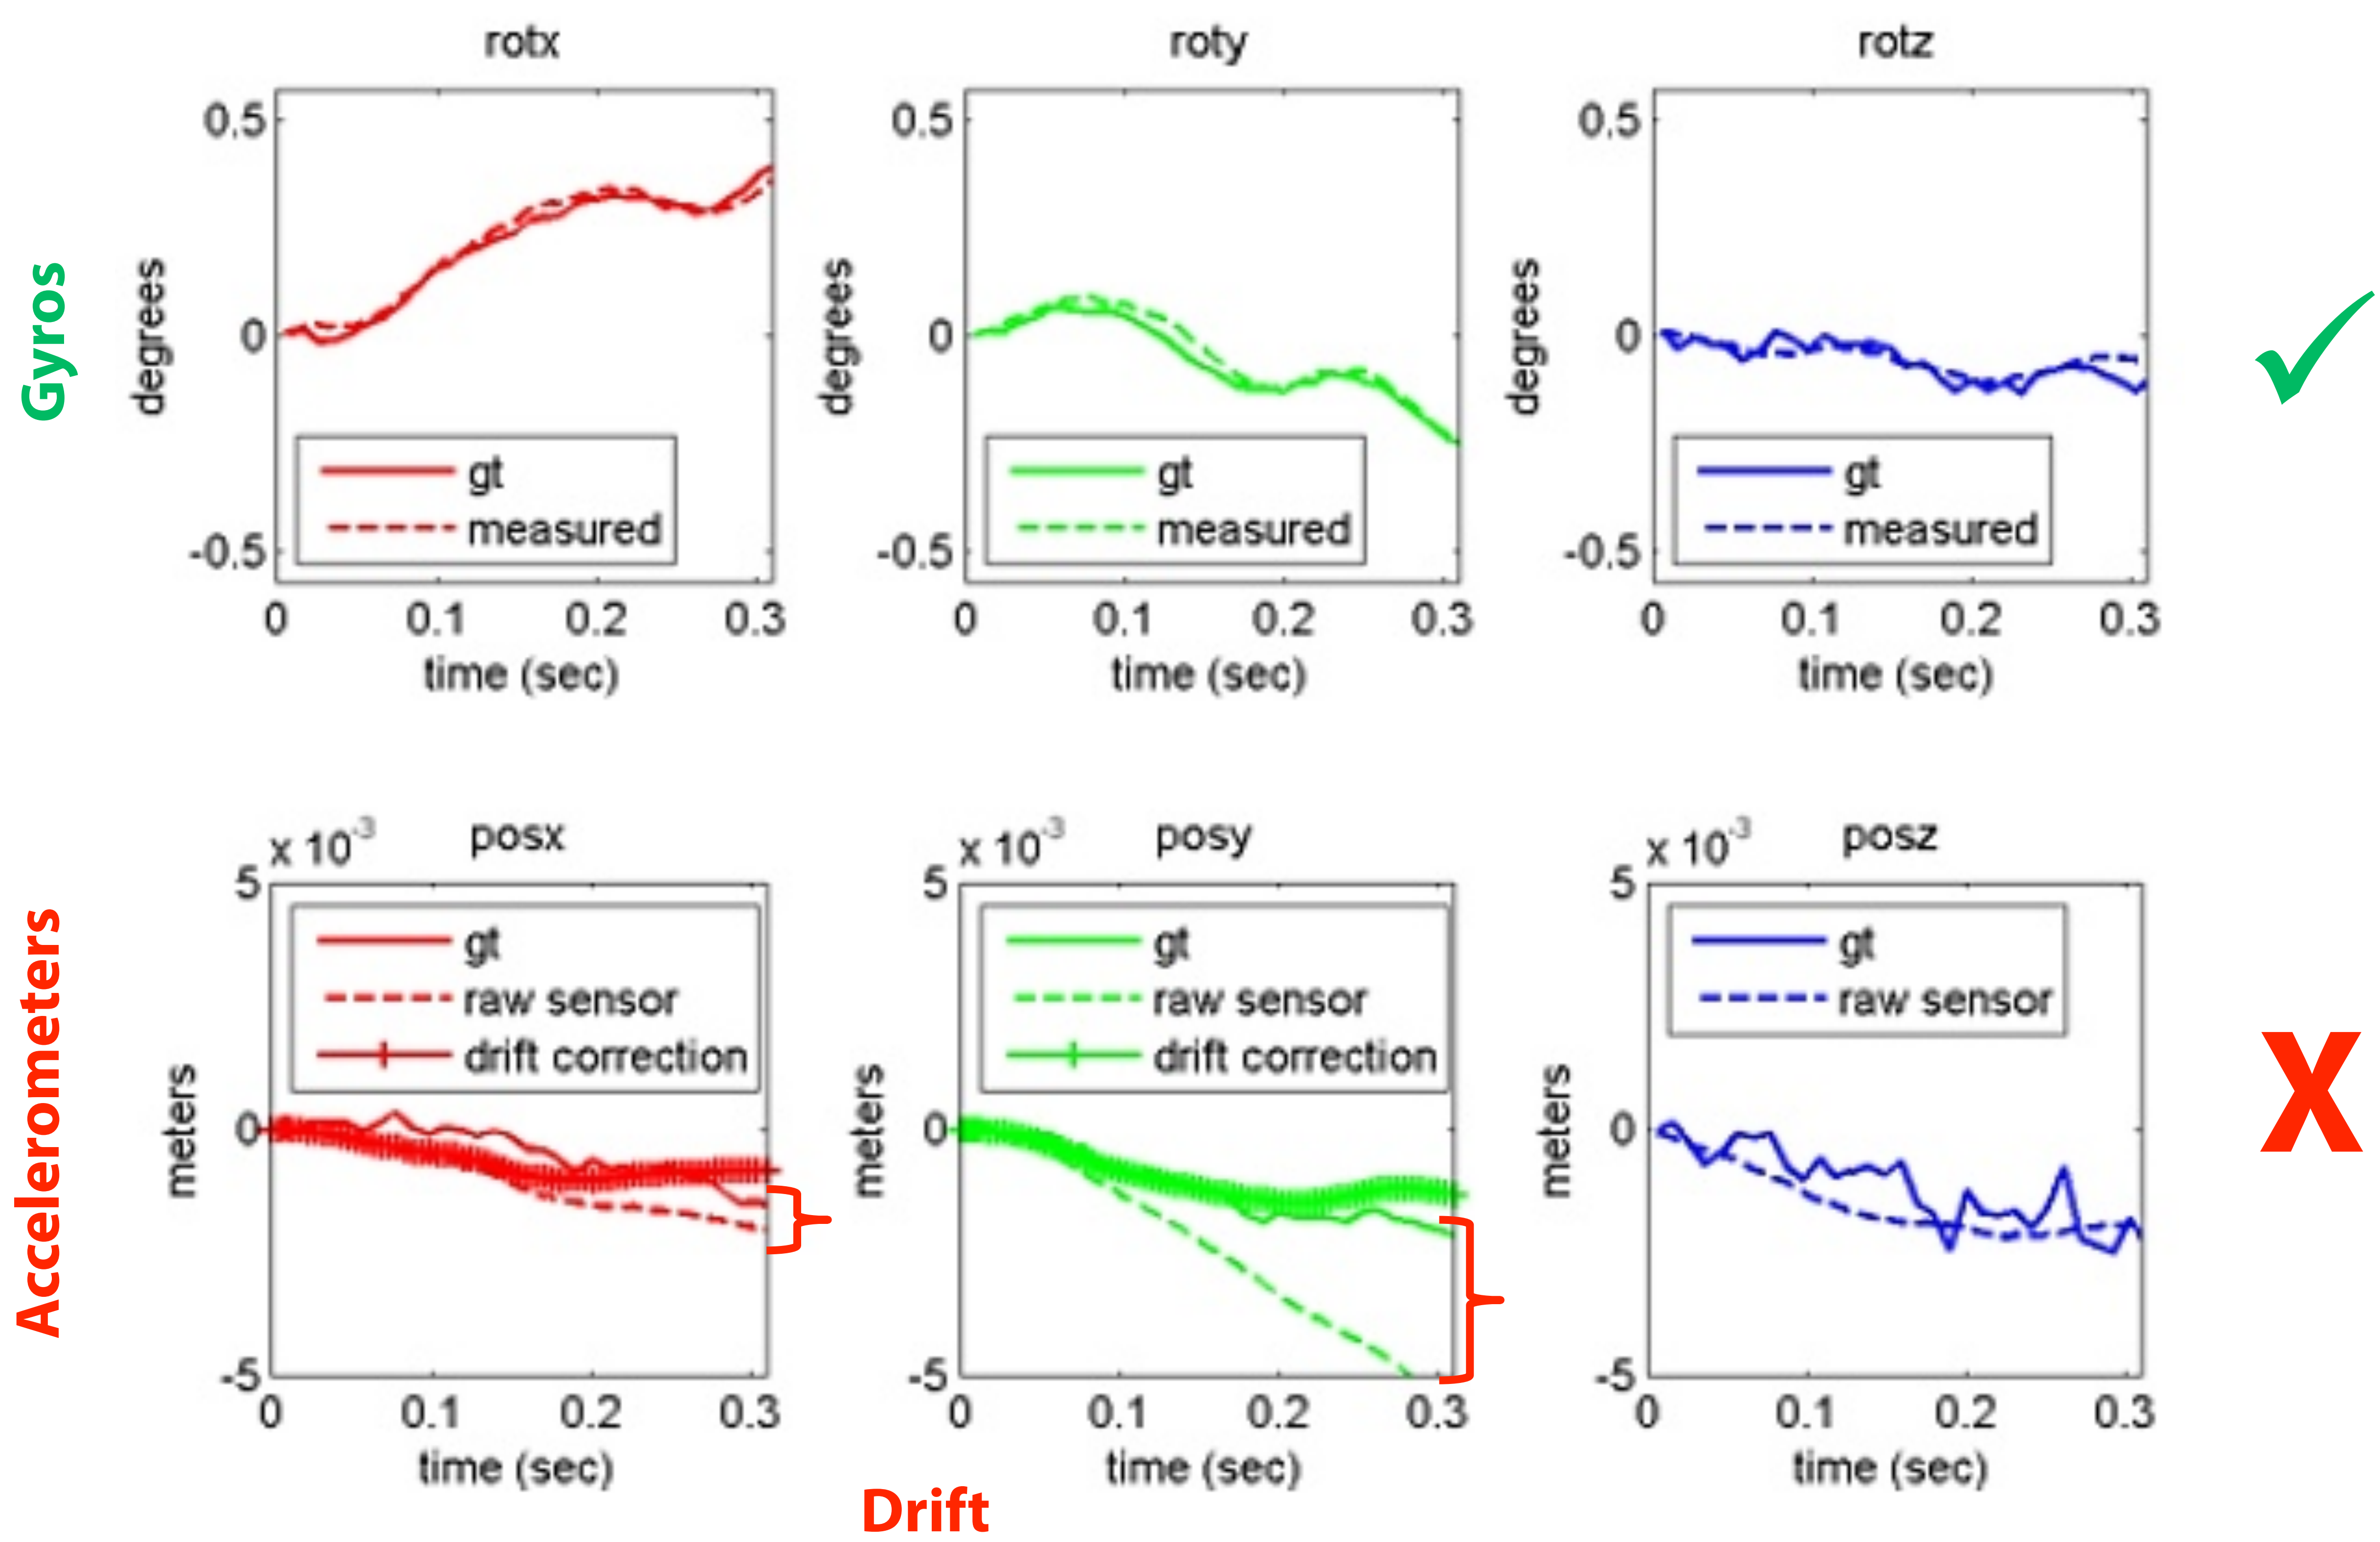

Drift Correction

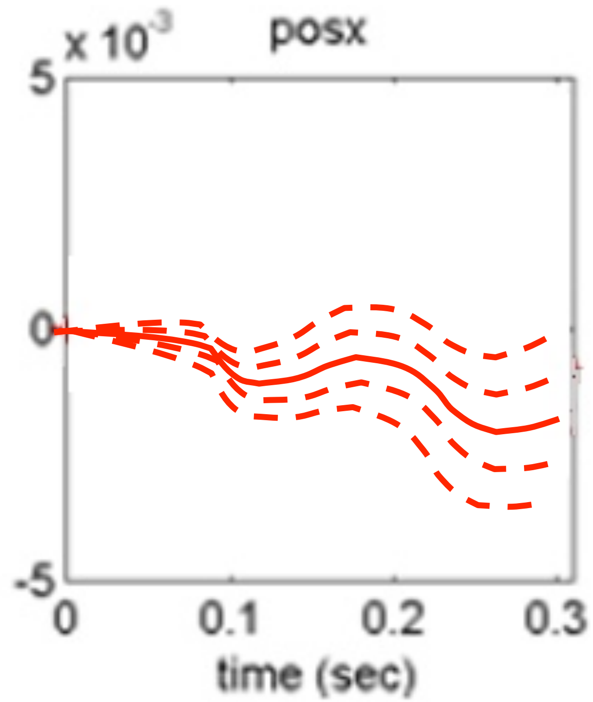

- Assume drift is linear
- Solve for x,y endpoint (u,v) (and planar depth) using **sensors as a constraint** and **maximize image prior**

$$I = \underset{I,d,u,v}{\operatorname{argmin}}[||\vec{B} - A(d,u,v)\vec{I}||^2/\sigma^2 + \lambda||\nabla I||^{0.8}]$$

- Nelder-Mead Simplex Optimization

Large blur kernels (>20 pixels)

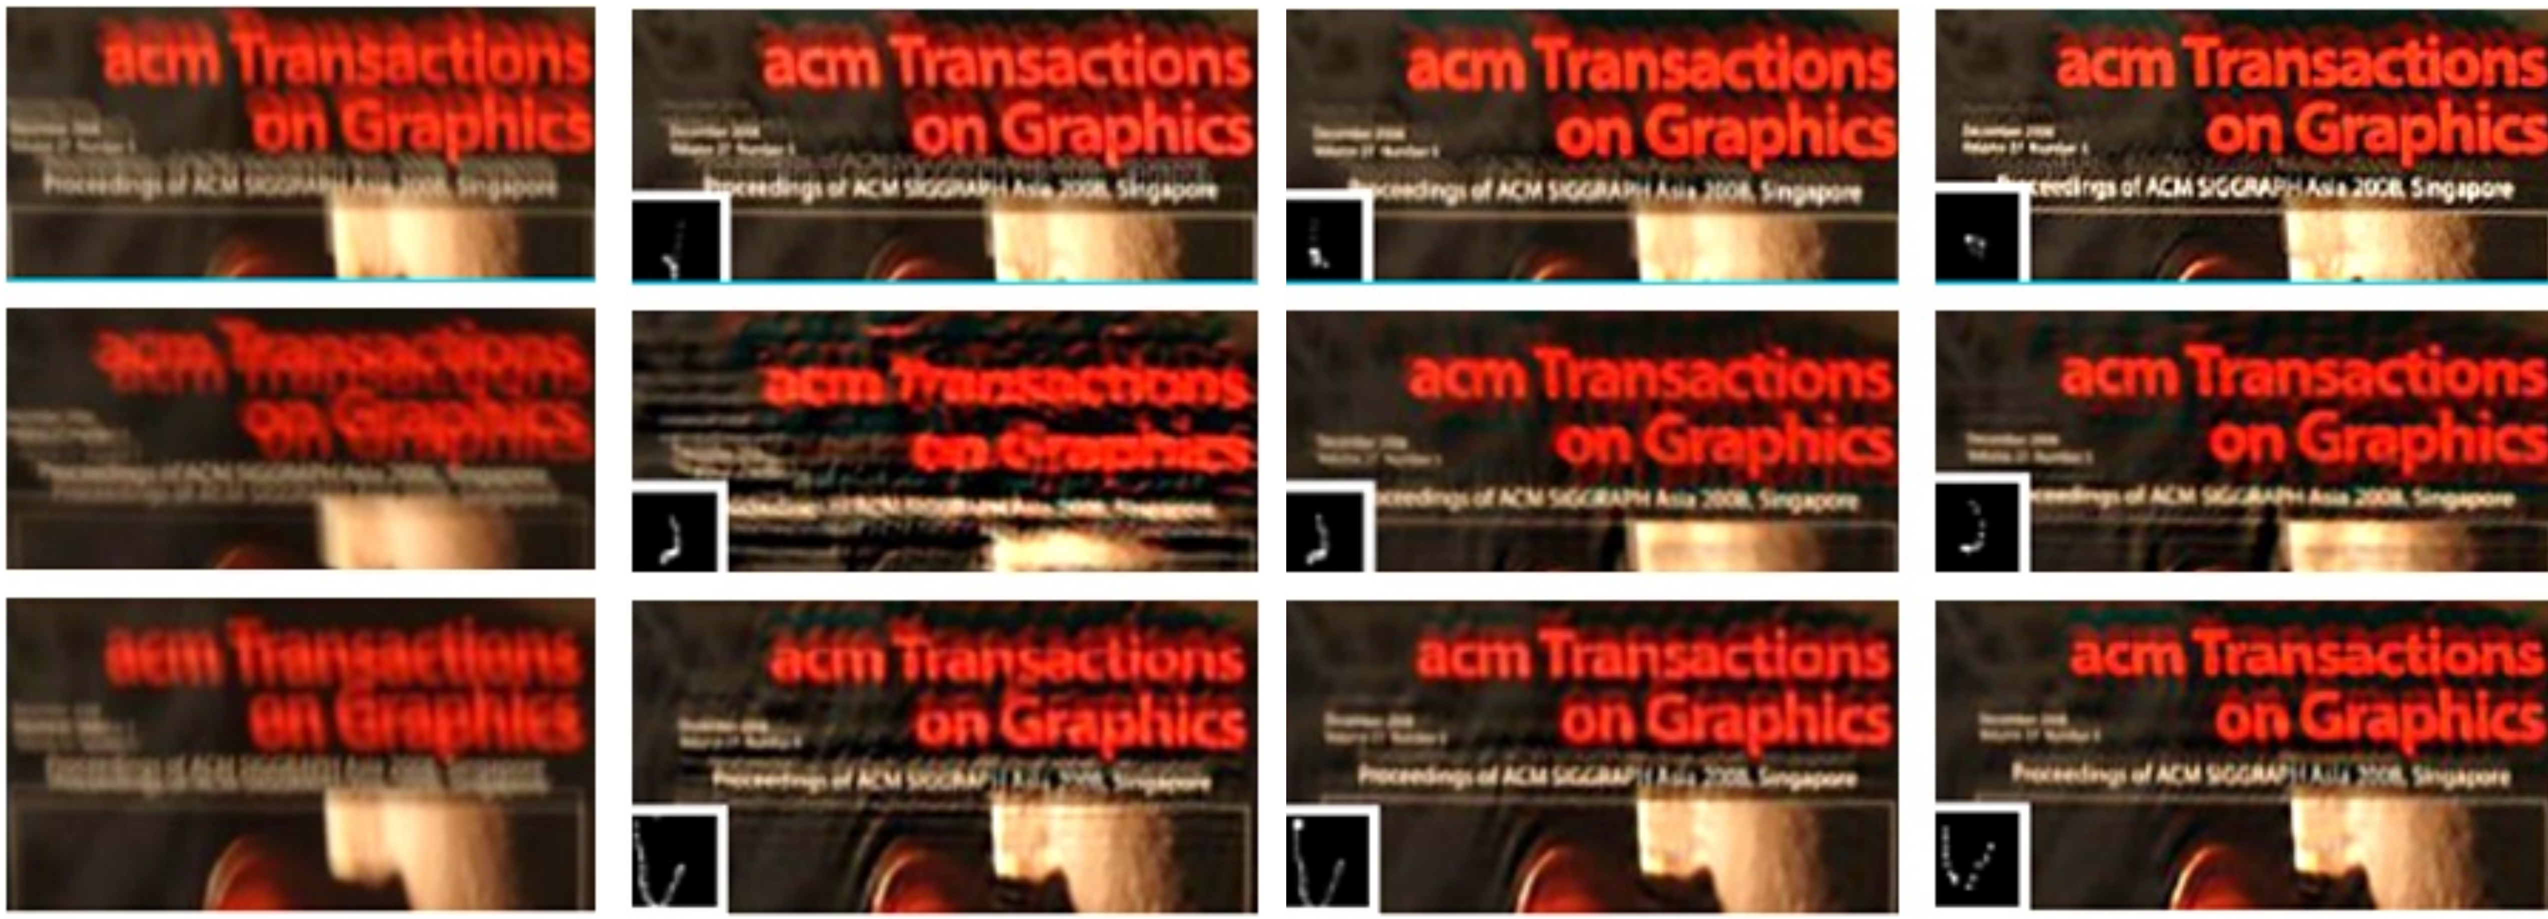

Blurry Image | Using PSFs from the raw sensor values | Our Output (after drift correction) | Using Groundtruth Motion

Large blur kernels (>20 pixels)

Microsoft Research

Blurry Image | Our Output | Shan et al. | Fergus et al.

Results: Deblurred

Results: Deblurred

Results: Deblurred

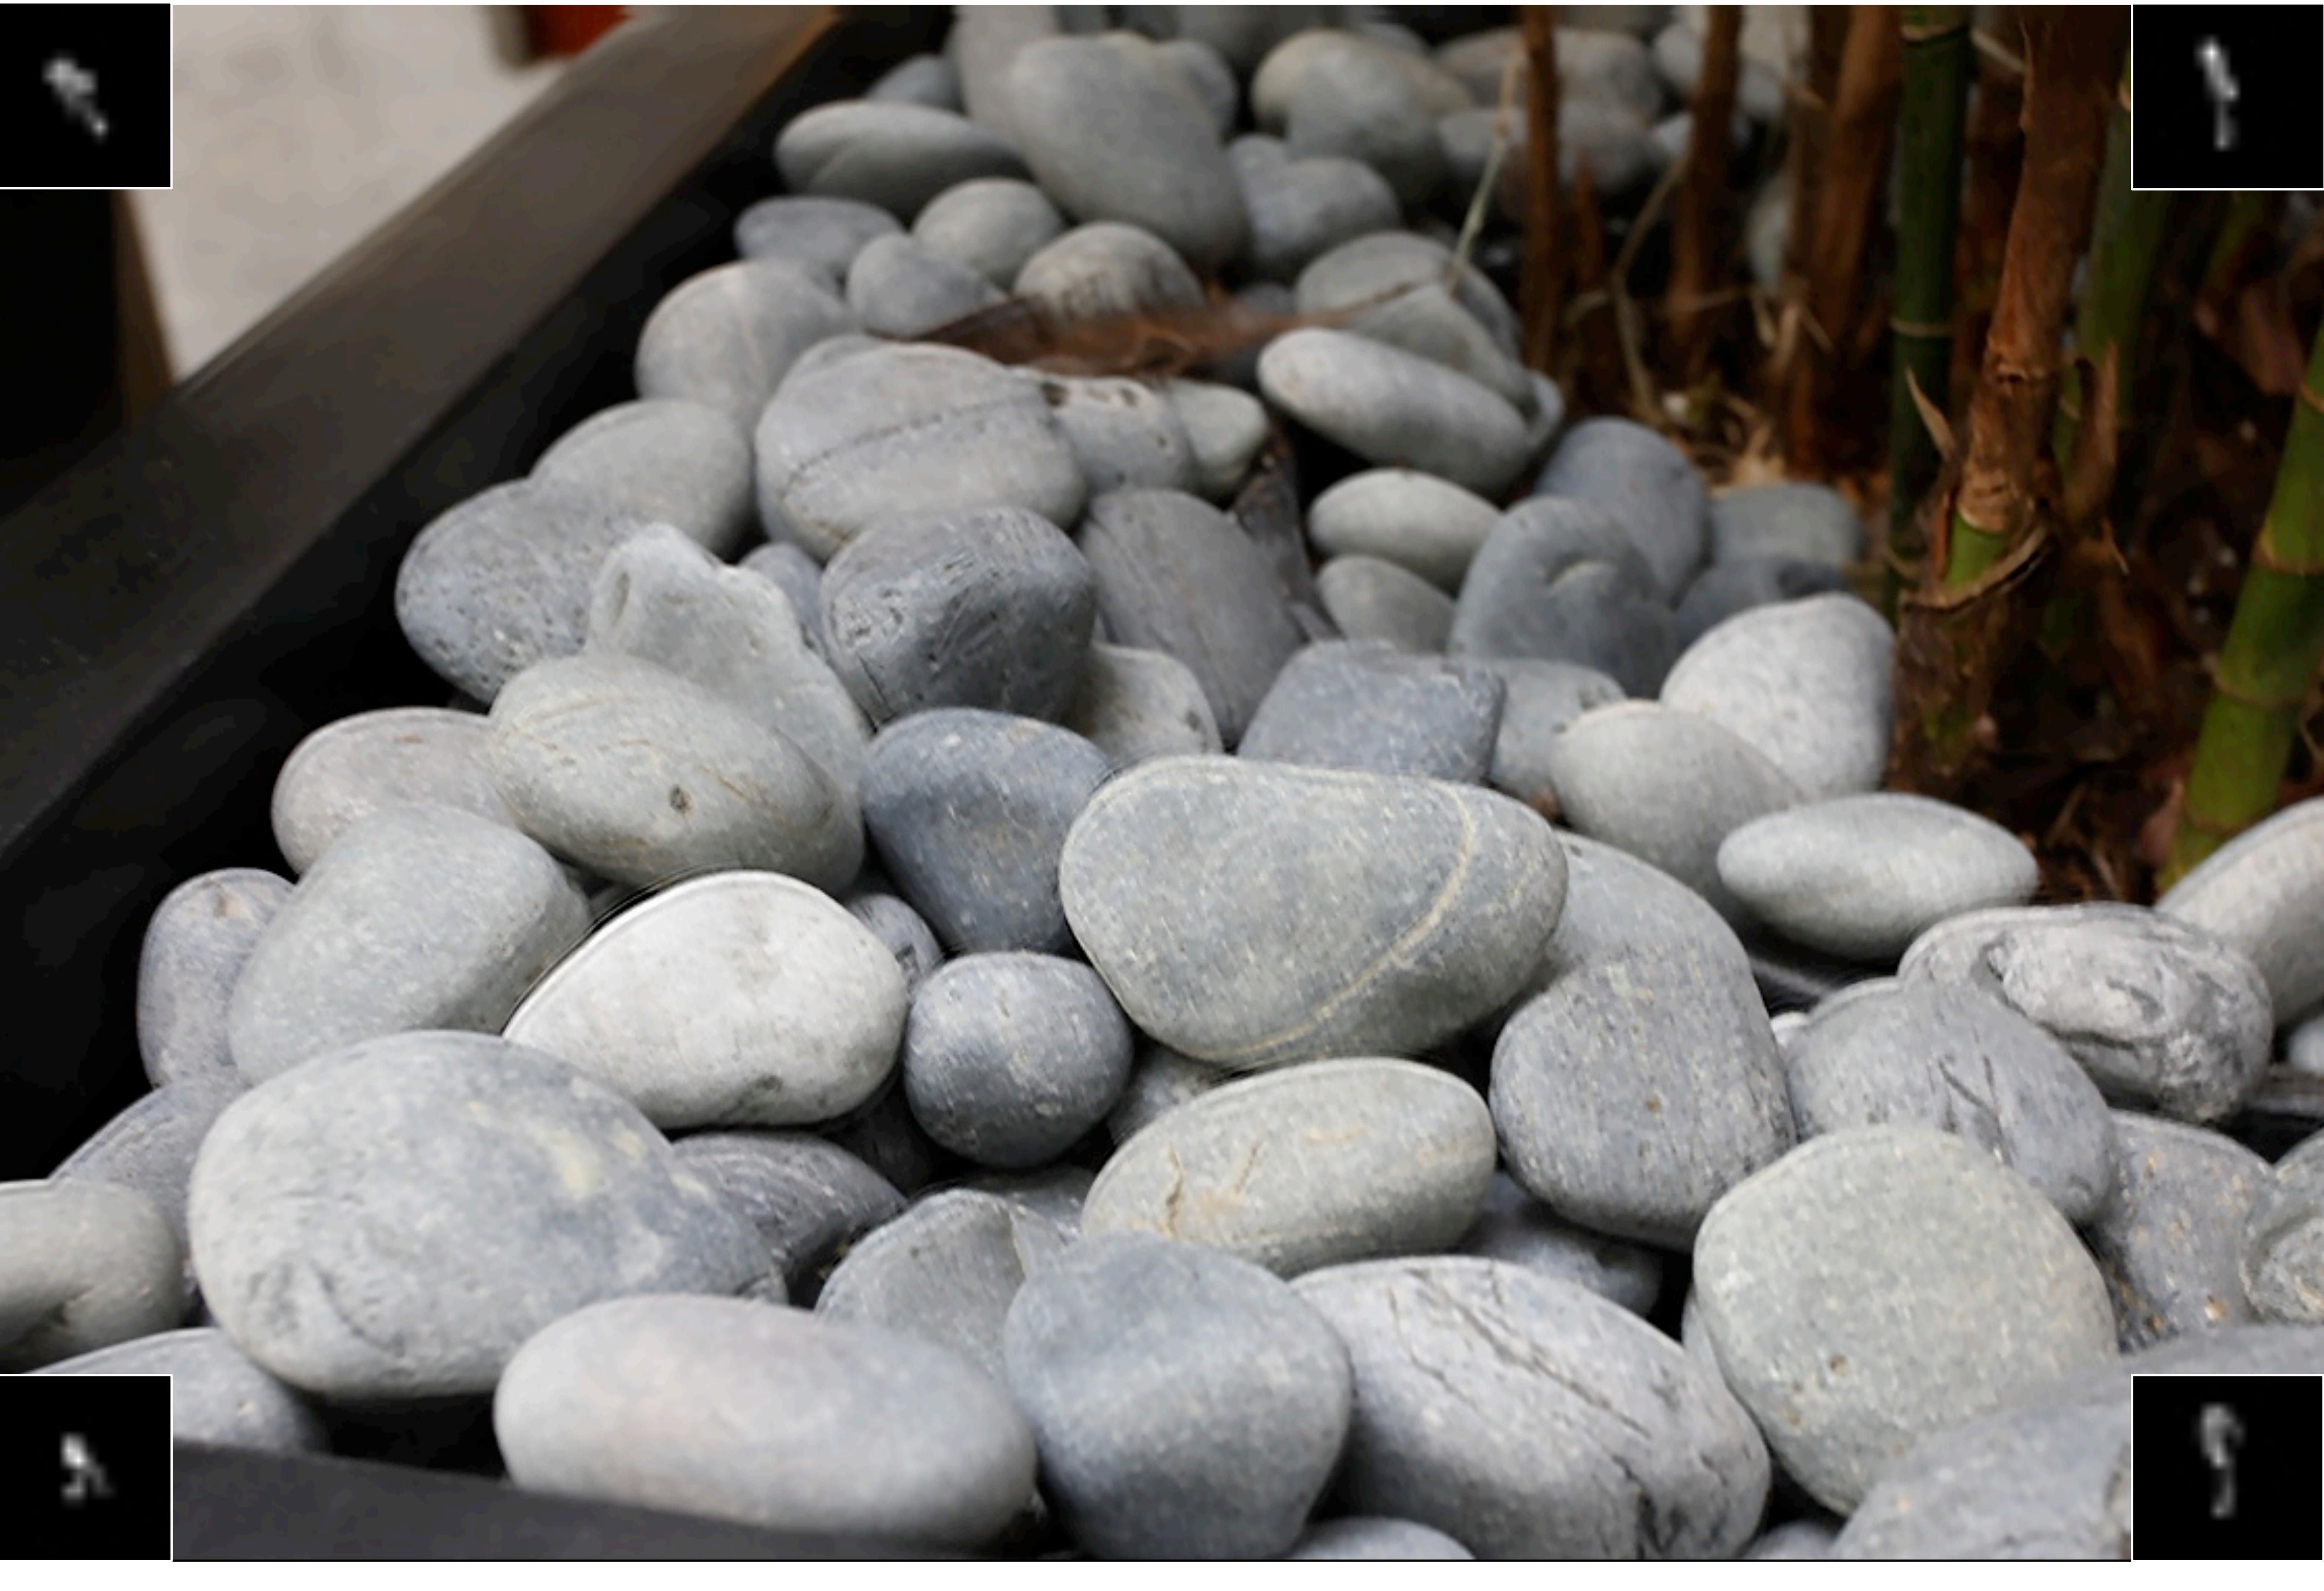

Comparison to Spatially Invariant Deblurring

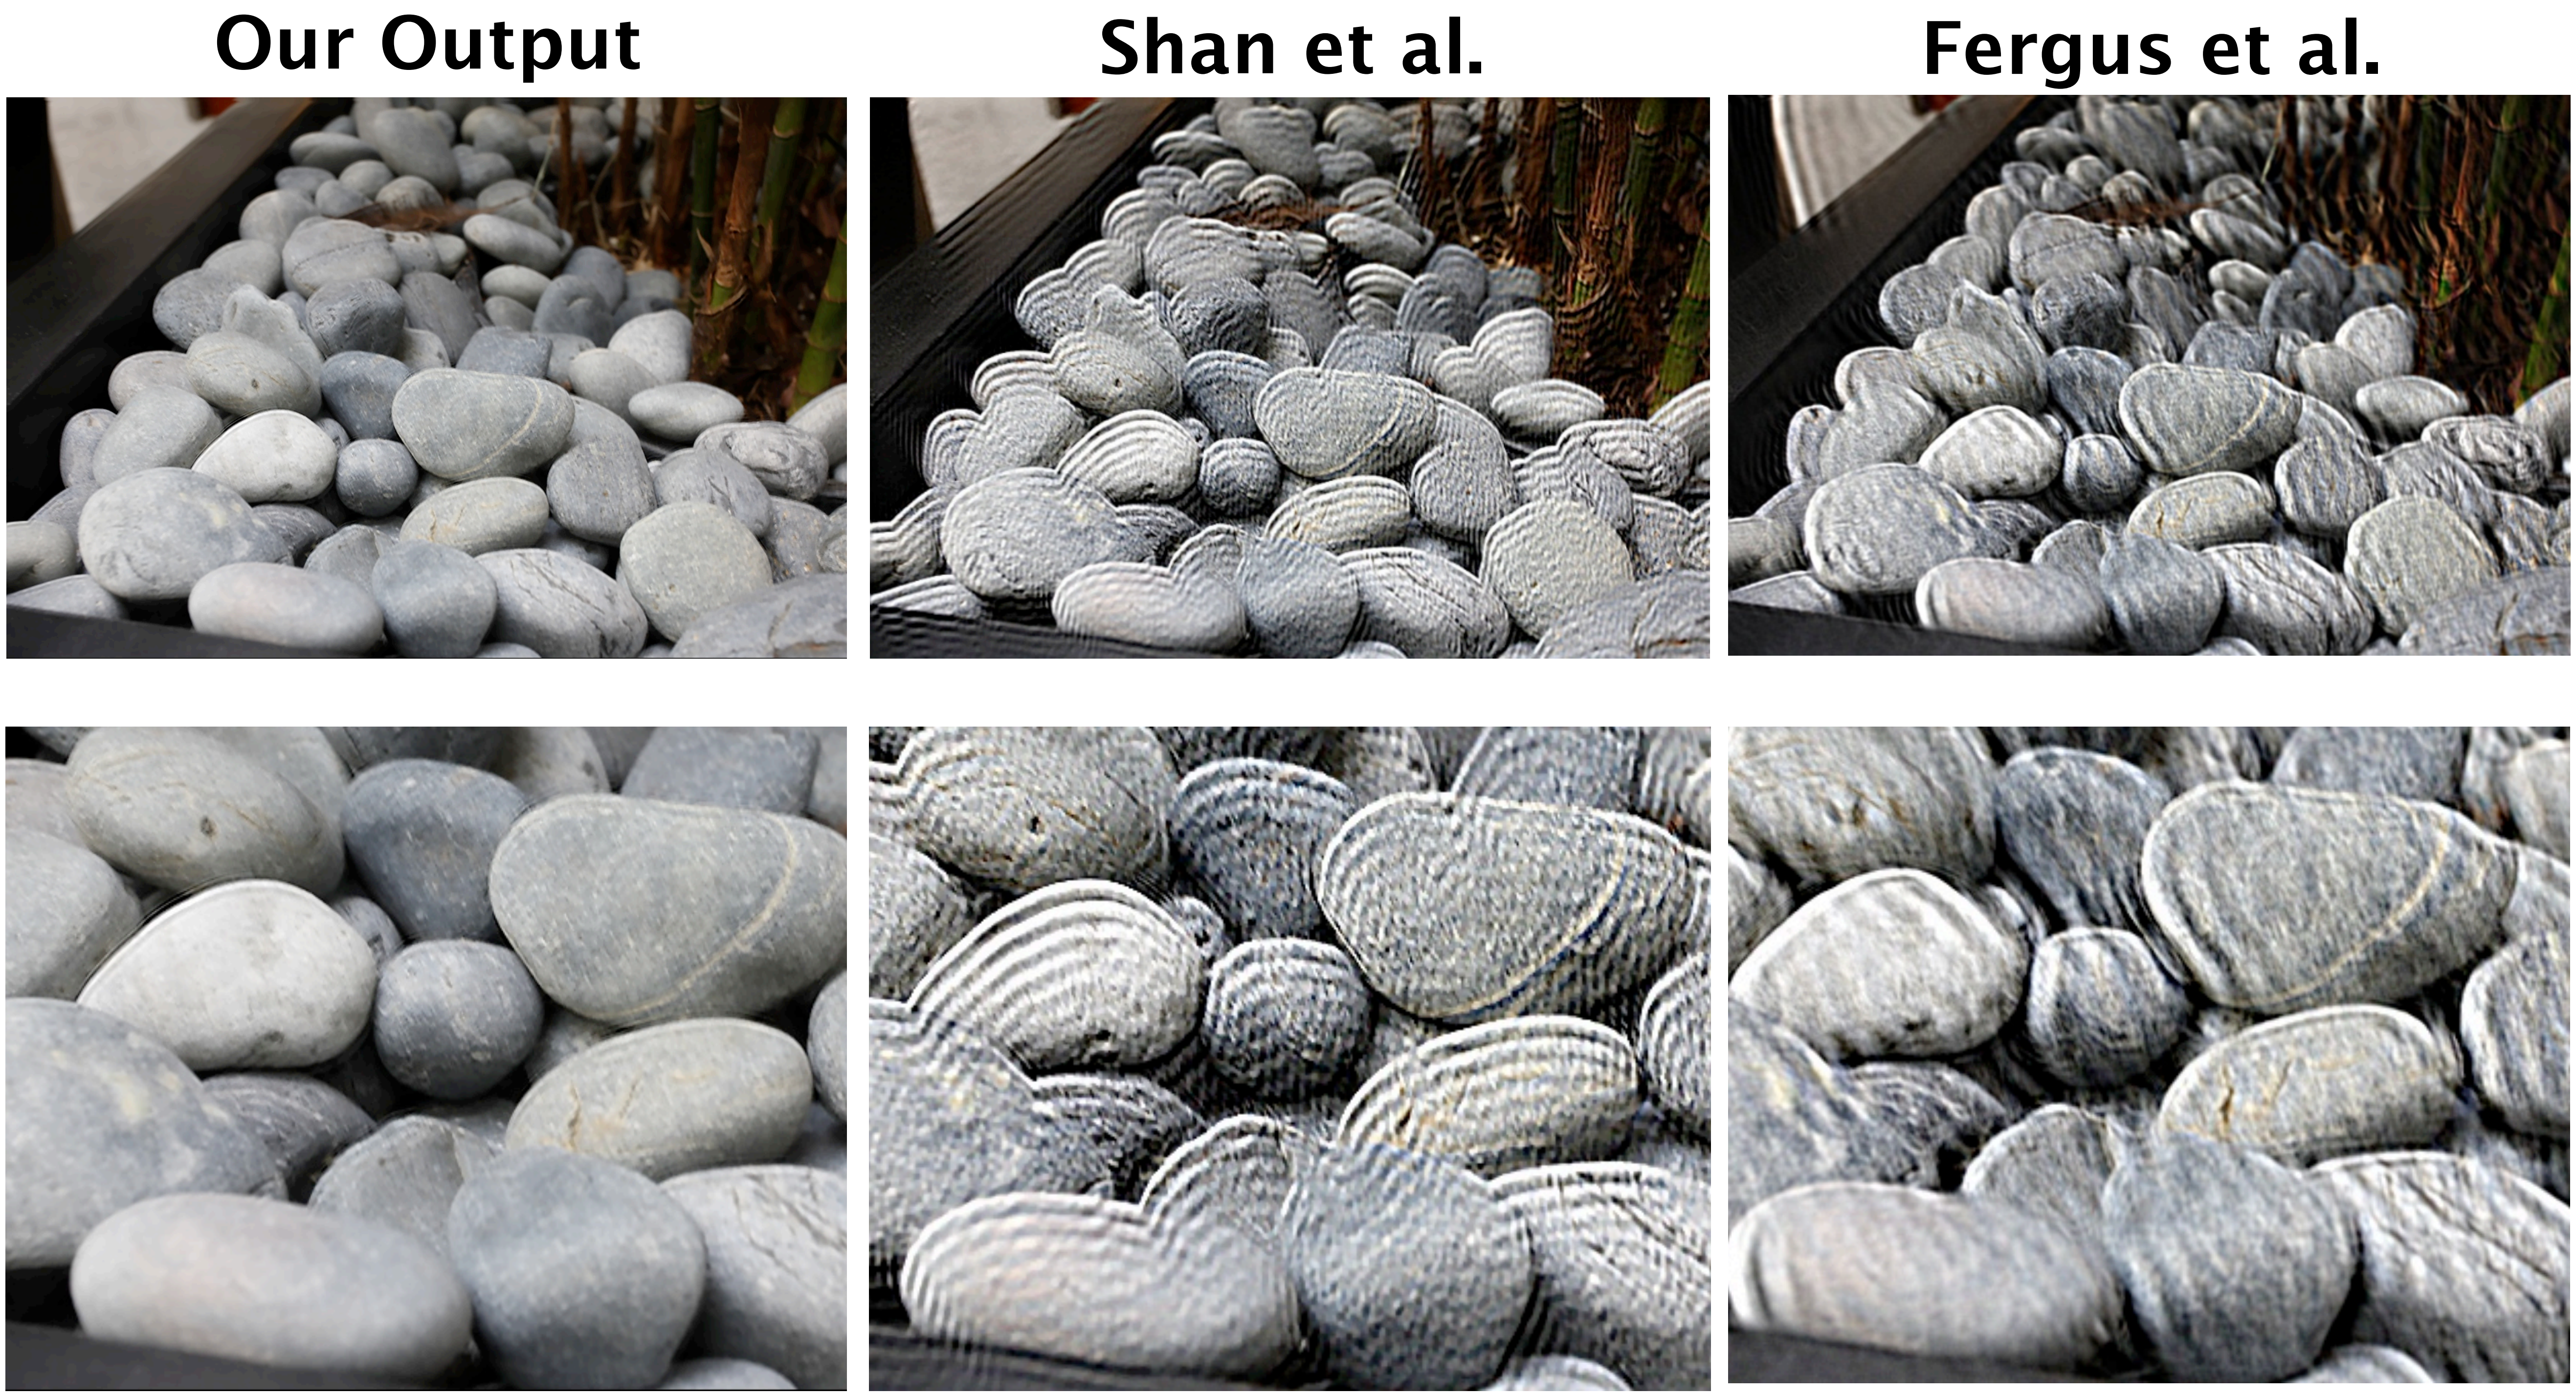

Results: Deblurred

Bibliography

R. Fergus, B Singh, A Hertzmann, S T Roweis, and W.T. Freeman, **Removing camera shake from a single photograph**, *SIGGRAPH 2006*.

N Joshi, S.B. Kang, C L Zitnick, and R Szeliski, **Image deblurring using inertial measurement sensors**, *SIGGRAPH 2010*.

S.K. Nayar and M Ben-Ezra, **Motion-based motion deblurring**, *PAMI* 26:6 (2004).

Q Shan, J Jia, and A Agarwala, **High-quality motion deblurring from a single image**, *SIGGRAPH 2008*.

W.H. Richardson, **Bayesian-Based Iterative Method of Image Restoration**, *JOSA* 62:1 (1972).

L.B. Lucy, **An iterative technique for the rectification of observed distributions**, *The Astronomical Journal* 79:6 (1974).