b
UNIVERSITÄT
BERN

Blind Deconvolution

From Model-Based to Deep Learning

Paolo Favaro
Computer Vision Group — University of Bern


Meiguang Jin


Givi Meishvili


Stefan Roth


Zhe Hu


Daniele Perrone

NTIRE 2019 — Long Beach, CA

Motion Blur

Motion blur is caused by object and/or camera motion during the exposure interval

Motion Blur

Motion blur is caused by object and/or camera motion during the exposure interval


Short Exposure Images

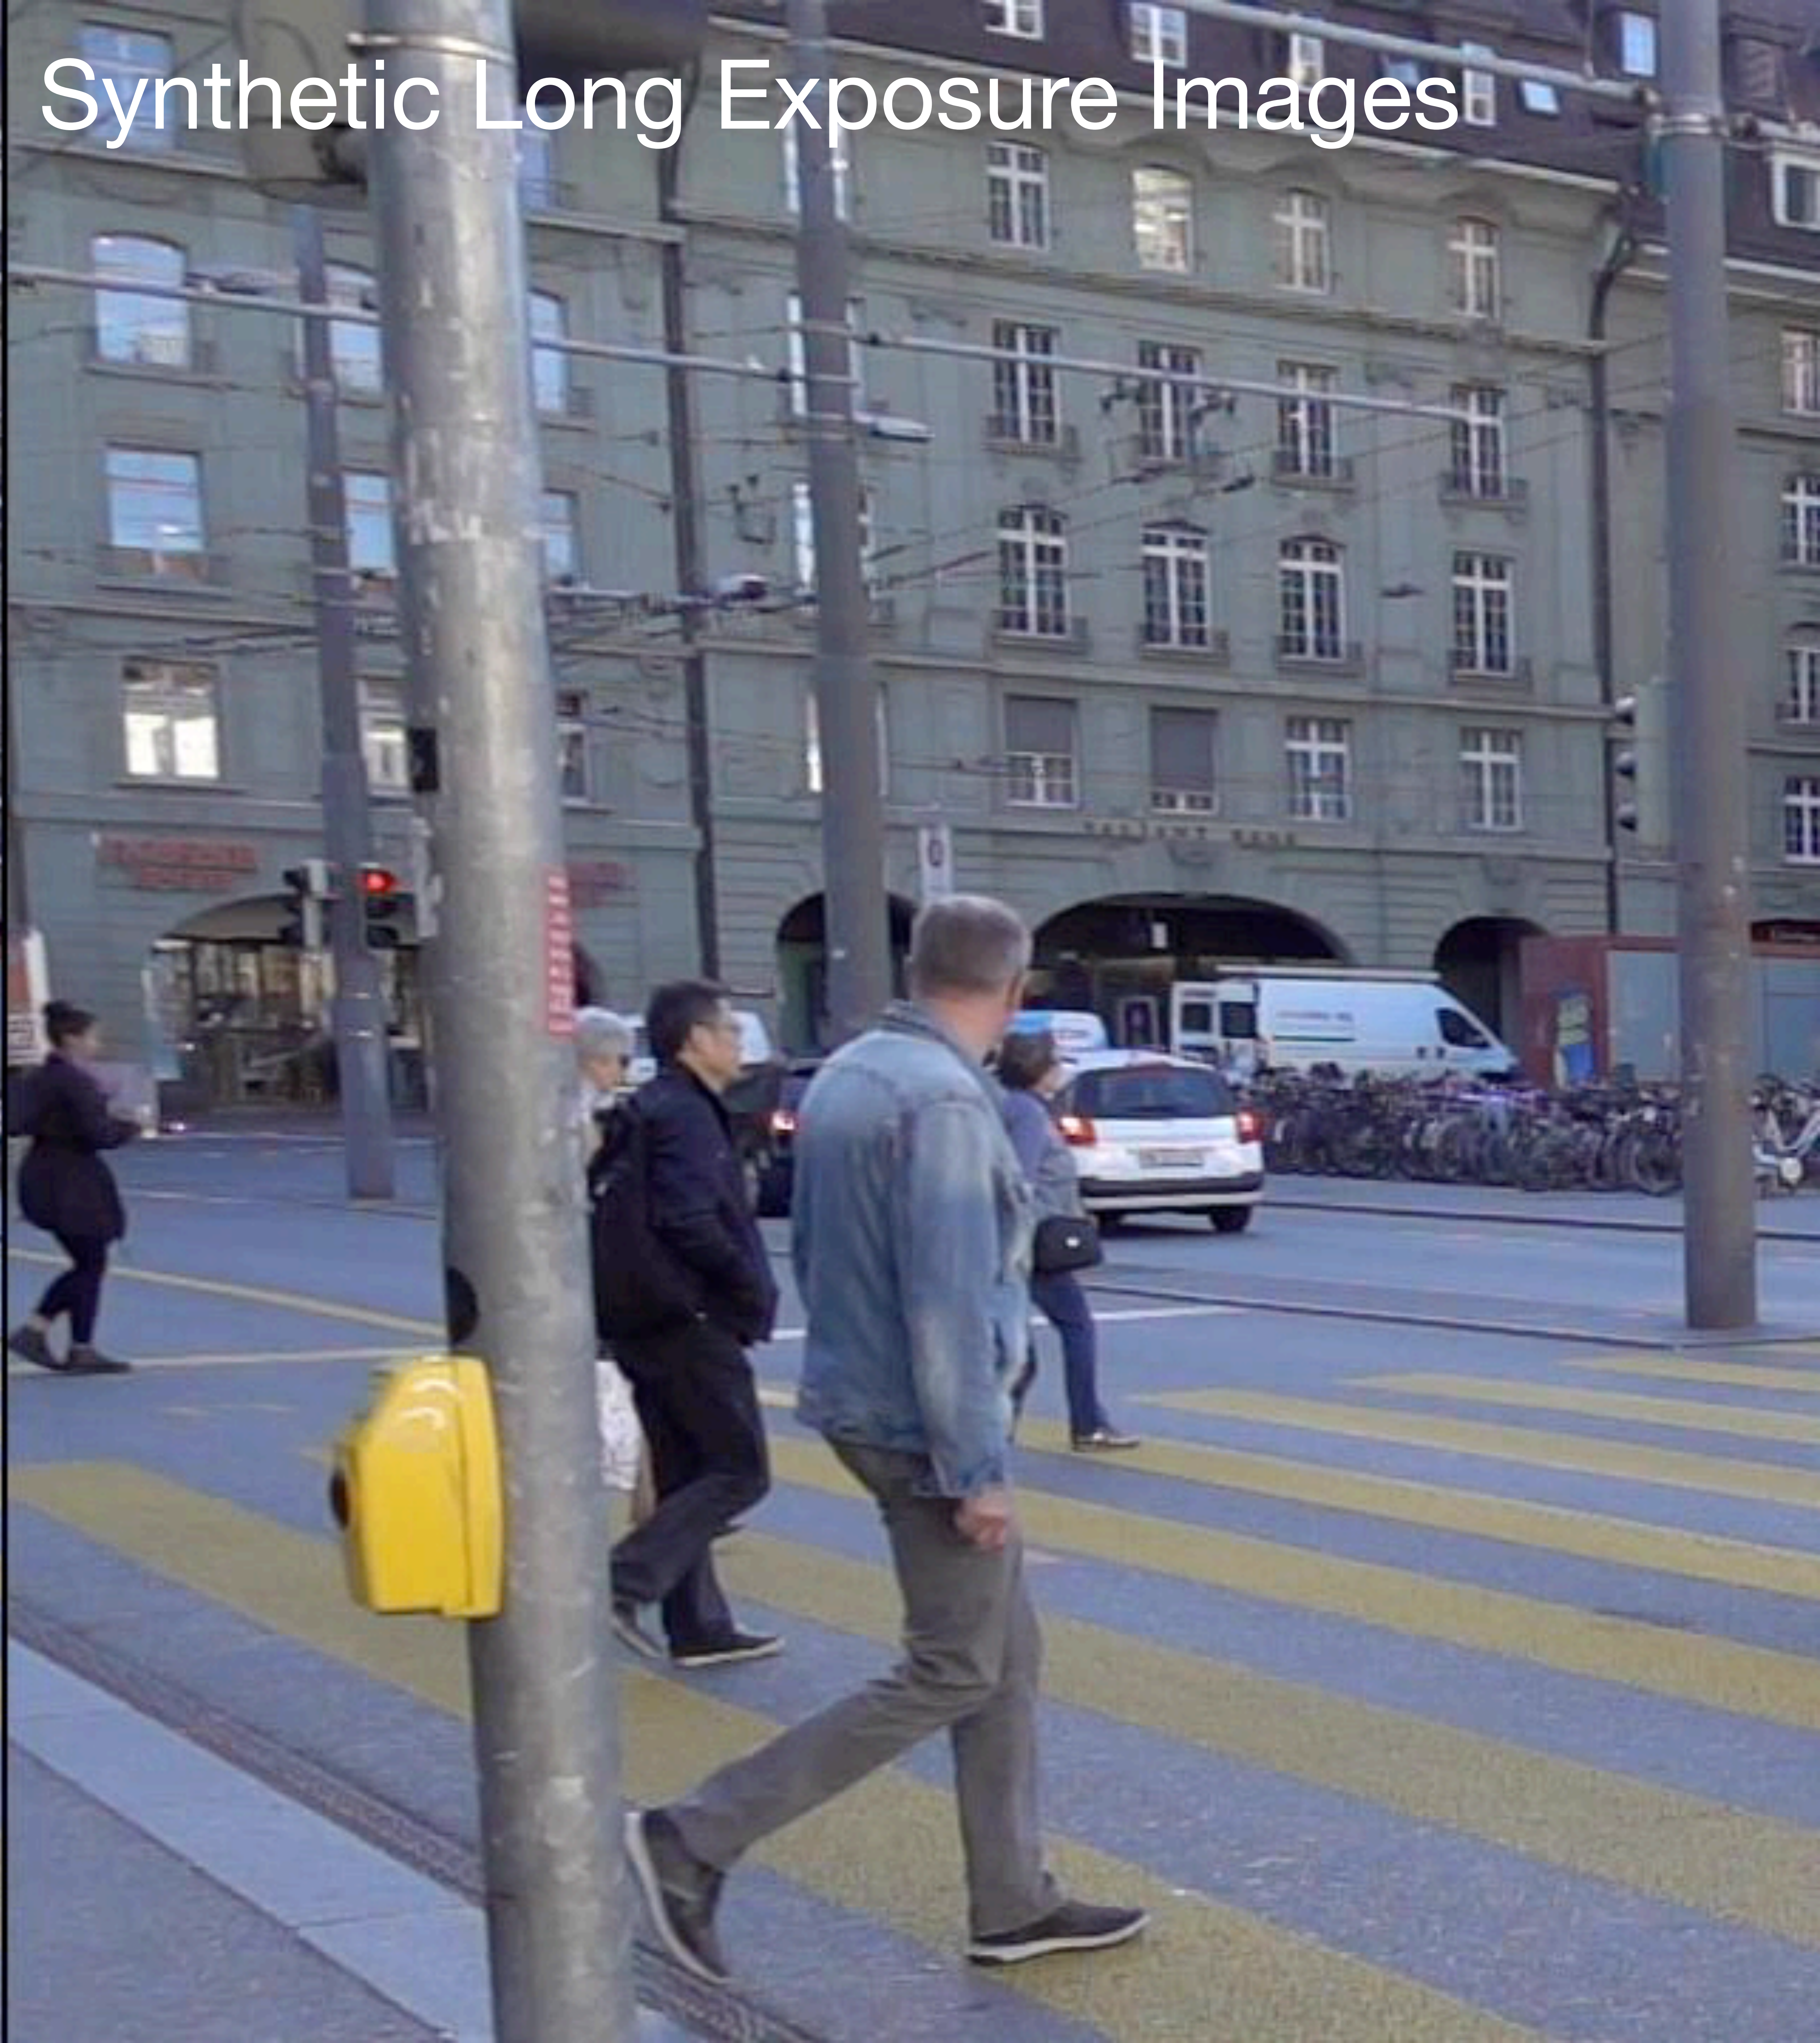
Synthetic Long Exposure Images


Short Exposure Images
Synthetic Long Exposure Images


Blurry Input
Sharp Ground Truth


Blurry Input
Sharp Frames


Blurry Input
Sharp Frames

Deep learning approach

- Need to collect ground truth data: (blur image, sharp video sequence)
  - Use high frame rate cameras, average frames to simulate blurry image, use the average as input and the sharp frames as output
- Need to address temporal ambiguities (eg forward or backward ordering yields the same blurry image), otherwise learning cannot succeed
  - Use a sequence order-invariant loss function

Blurry Input

7 Frame Estimates

Jin, Meishvili, Favaro Learning to Extract a Video Sequence from a Single Motion-Blurred Image CVPR 2018

Blurry Input

7 Frame Estimates

Jin, Meishvili, Favaro Learning to Extract a Video Sequence from a Single Motion-Blurred Image CVPR 2018

Blurry Input

7 Frame Estimates

Jin, Meishvili, Favaro Learning to Extract a Video Sequence from a Single Motion-Blurred Image CVPR 2018

Blurry Input

7 Frame Estimates

Jin, Meishvili, Favaro Learning to Extract a Video Sequence from a Single Motion-Blurred Image CVPR 2018

Slow motion & deblurring from a blurry video

input (30 FPS)

Jin, Zhe, Favaro Learning to Extract Flawless Slow Motion from Blurry Videos CVPR 2019

Slow motion & deblurring from a blurry video

output (300 FPS)

Jin, Zhe, Favaro Learning to Extract Flawless Slow Motion from Blurry Videos CVPR 2019

Slow motion & deblurring from a blurry video

input (30 FPS)

Poster #157 - Wednesday, June 19, 15.20 – 18.00

Jin, Zhe, Favaro Learning to Extract Flawless Slow Motion from Blurry Videos CVPR 2019

Slow motion & deblurring from a blurry video

output (300 FPS)

Poster #157 - Wednesday, June 19, 15.20 – 18.00

Jin, Zhe, Favaro Learning to Extract Flawless Slow Motion from Blurry Videos CVPR 2019

Deep learning approaches

- **pros**
  - Can handle scenes of high complexity
  - No need to manually design models/priors
  - No need to design custom optimization procedures
  - Extremely fast execution
- **cons**
  - Not state of the art in existing datasets (Nah et al @ -2dB PSNR from best model-based)
  - No direct control/guarantees on the artifacts

Nah et al *Deep multi-scale convolutional neural network for dynamic scene deblurring* CVPR 2017

Model-based approaches

- If the camera translates along the X-Y axes and the scene is a fronto-parallel plane (or at infinity) a simple blur model is

$$f = k * u + n$$

blurry image

=

kernel

*

sharp image

+

noise

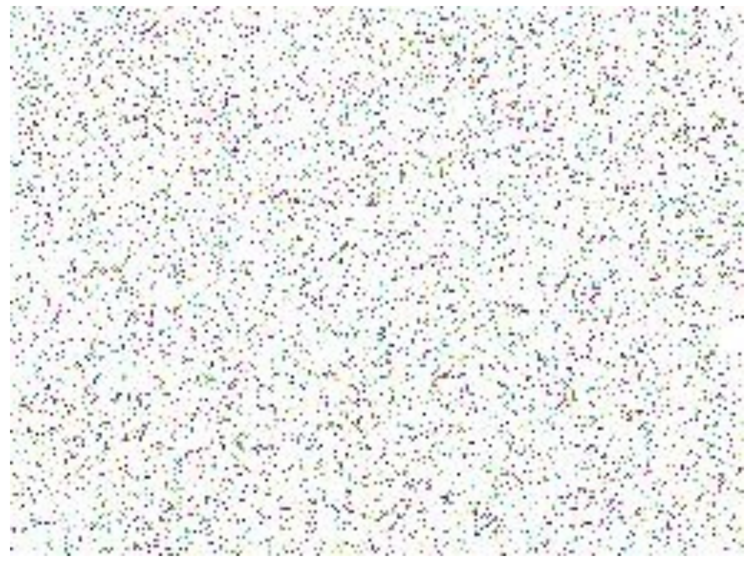

Model-based approaches

- If the camera translates along the X-Y axes and the scene is a fronto-parallel plane (or at infinity) a simple blur model is

$$f = k * u + n$$

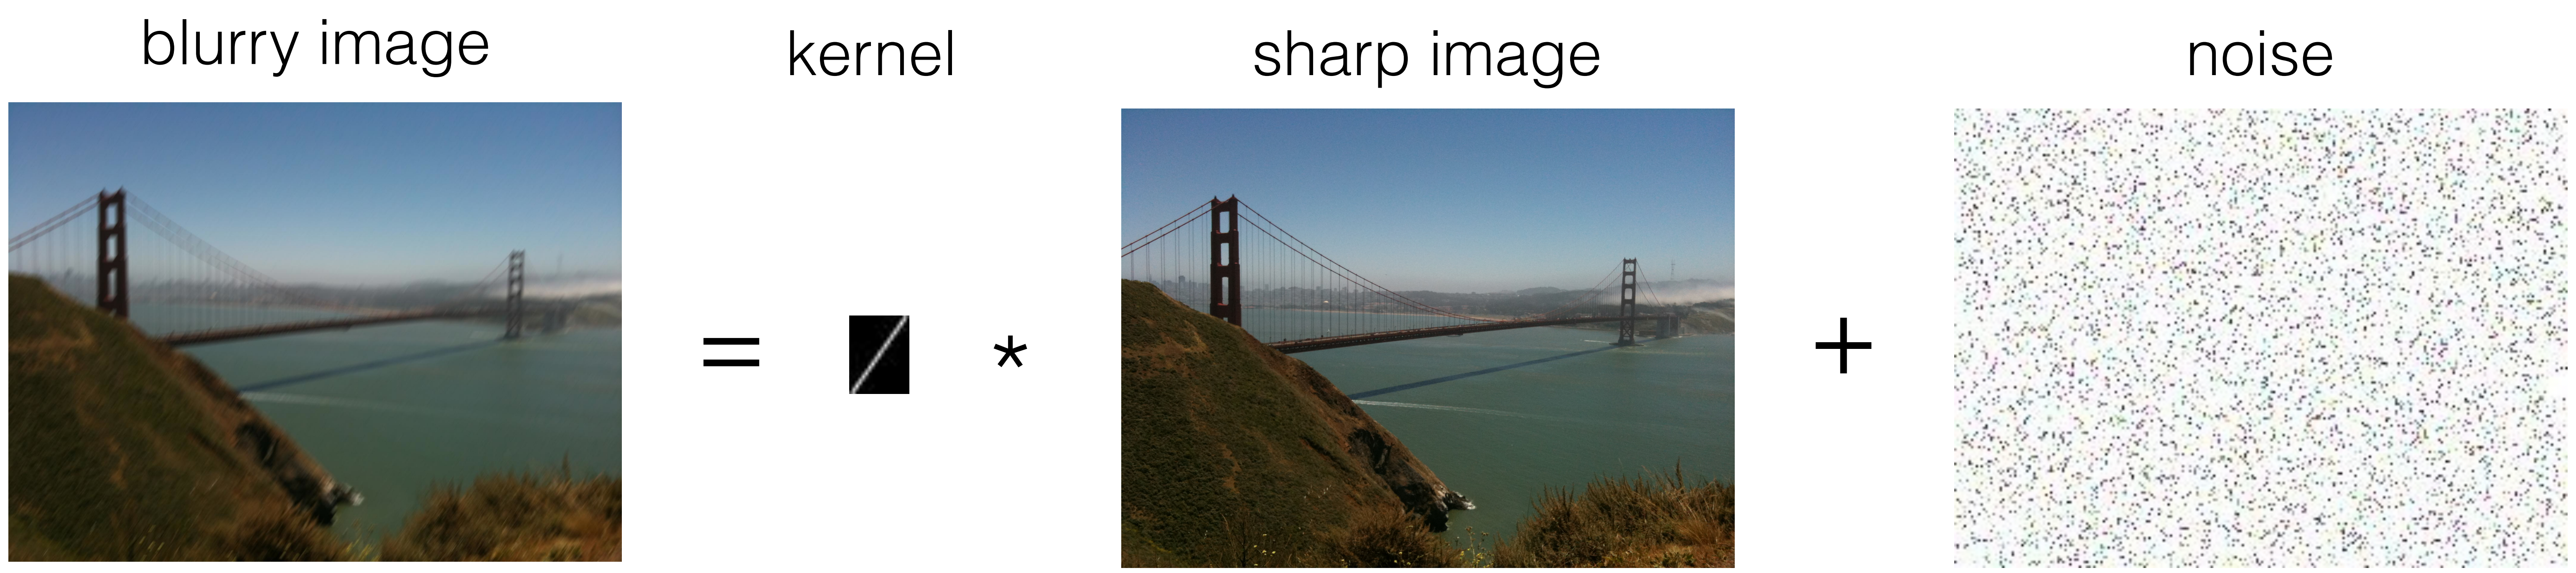

Model-based approaches

- If the camera translates along the X-Y axes and the scene is a fronto-parallel plane (or at infinity) a simple blur model is

$$f = k * u + n$$

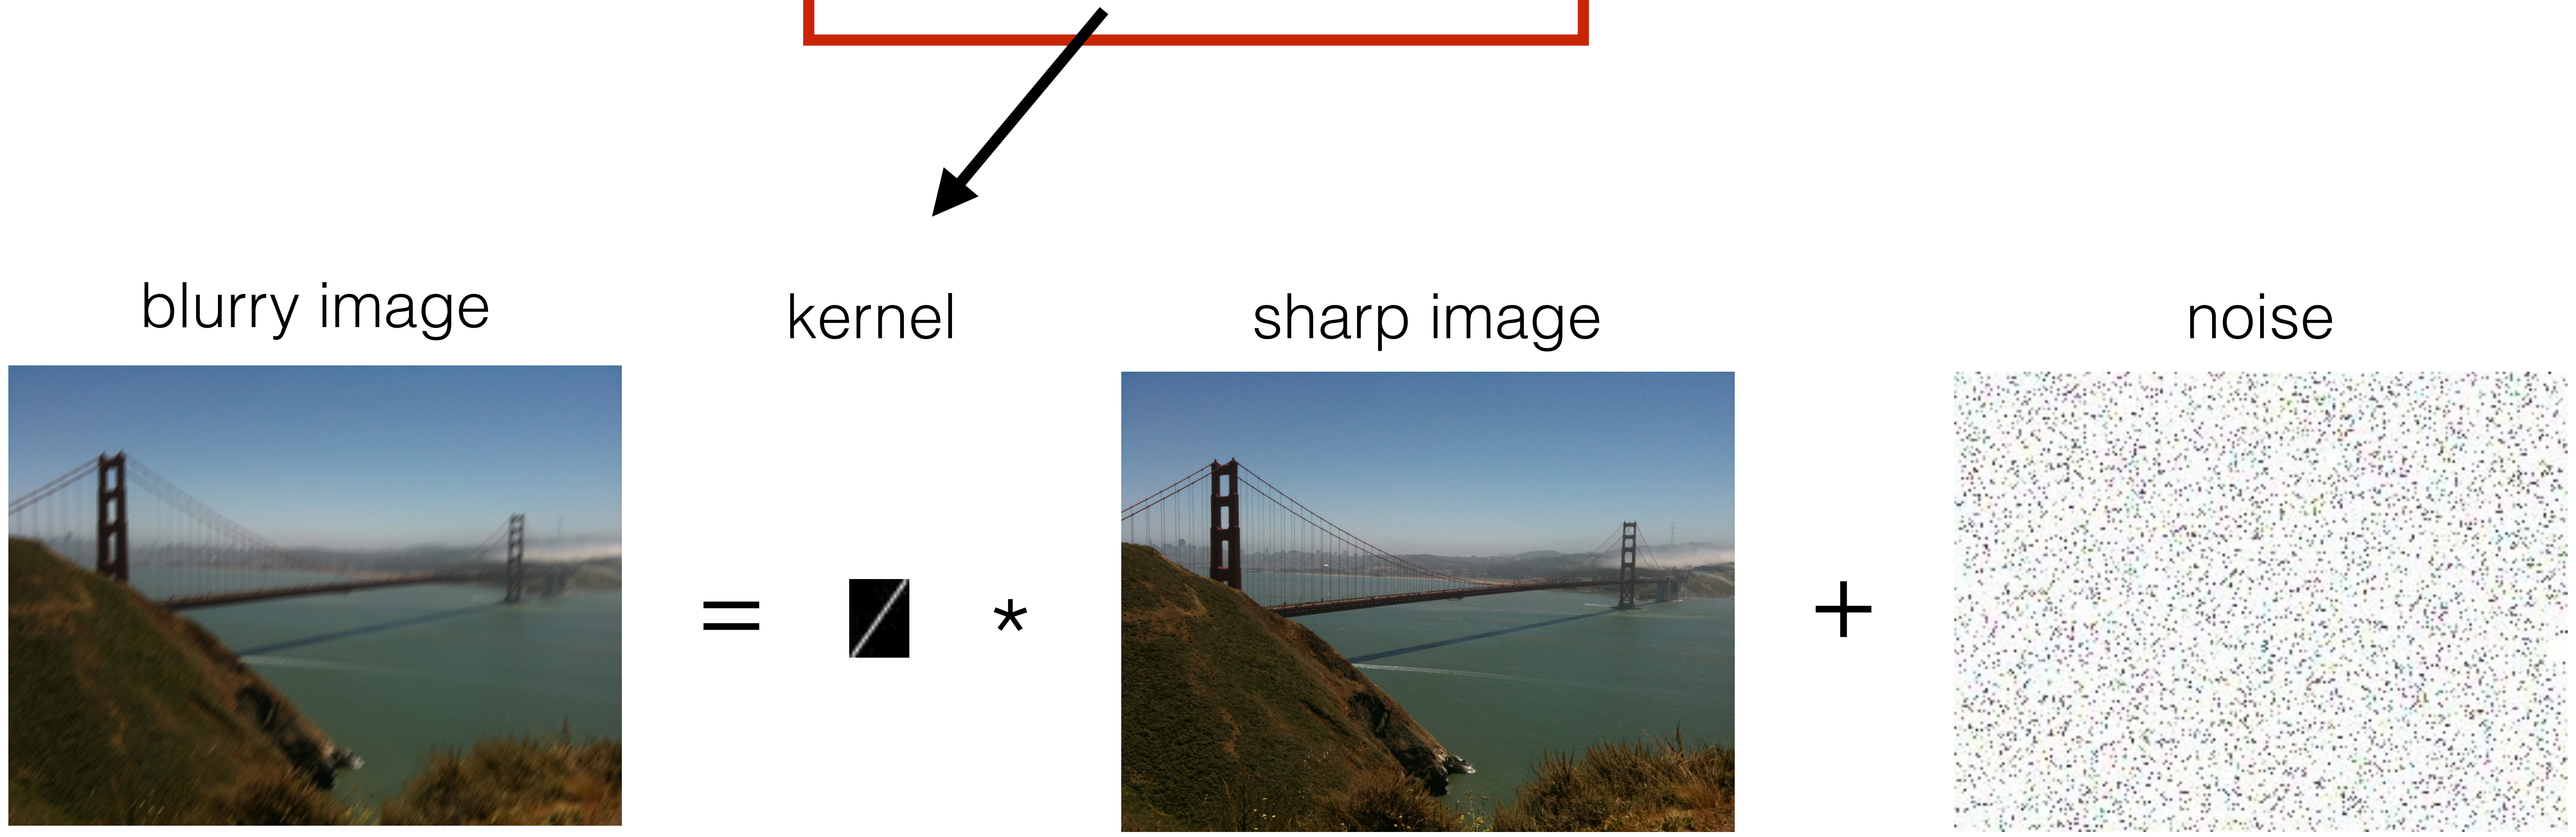

Model-based approaches

- If the camera translates along the X-Y axes and the scene is a fronto-parallel plane (or at infinity) a simple blur model is

$$f = k * u + n$$

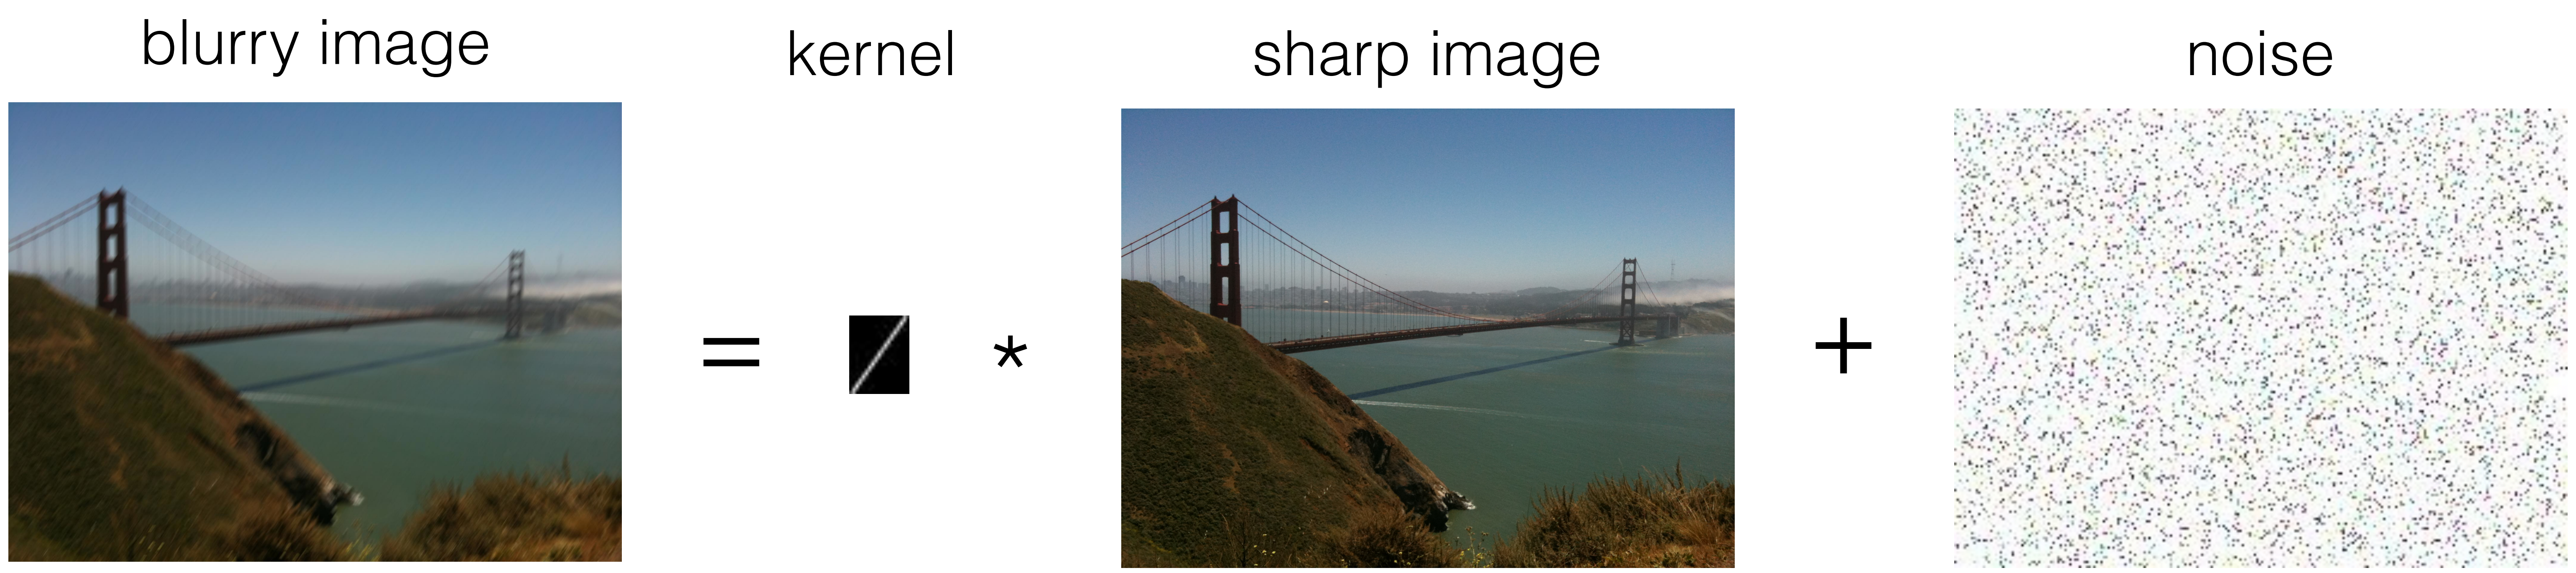

Model-based approaches

- If the camera translates along the X-Y axes and the scene is a fronto-parallel plane (or at infinity) a simple blur model is

$$f = k * u + n$$

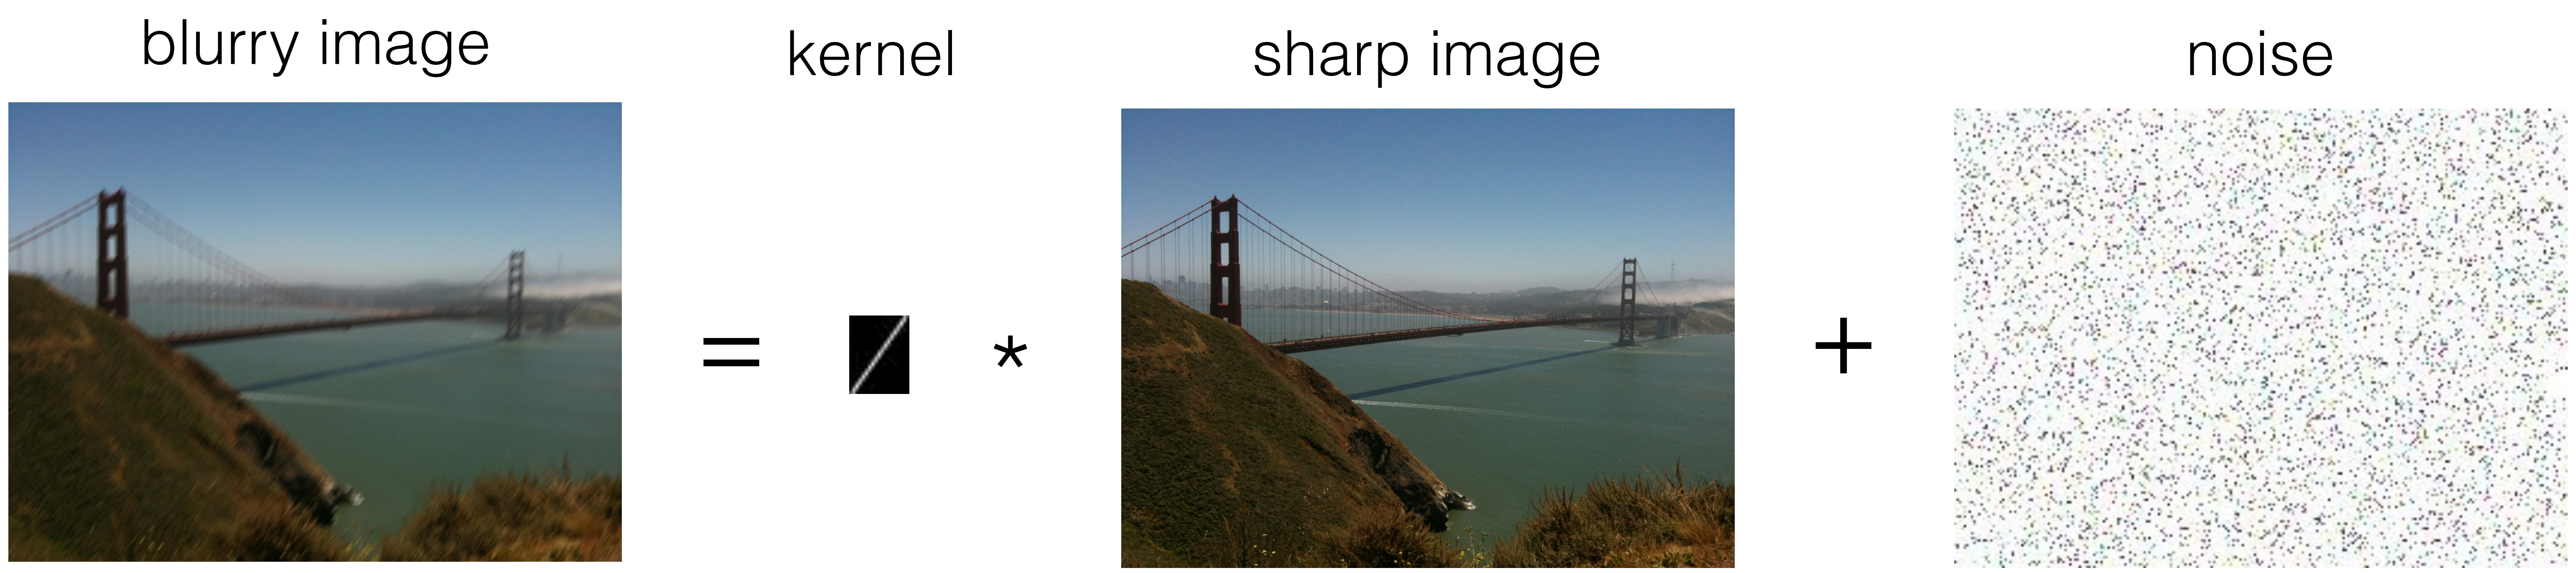

Blind deconvolution

- Recover **both** the blur kernel and the sharp image given the blurry image

$$f = k * u + n$$

- By using Maximum a Posteriori it can be posed as an optimization problem with some image prior (eg Total Variation*)

$$\min_{u,k} \lambda |\nabla u|_{2,1} + \frac{1}{2}|f - k * u|_2^2$$

*Chan and Wong *Total Variation Blind Deconvolution* TIP 1998 (also You and Kaveh 1996)

A little problem

- The TV prior has a little flaw

$$\min_{u,k} \lambda |\nabla u|_{2,1} + \frac{1}{2}|f - k * u|_2^2$$

*Levin et al Understanding and evaluating blind deconvolution algorithms CVPR 2009

A little problem

- The TV prior has a little flaw

$$\min_{u,k} \lambda |\nabla u|_{2,1} + \frac{1}{2}|f - k * u|_2^2$$

- Compare the true solution (u, k) with the no-blur solution (f, δ)

$$\boxed{f \equiv \delta * f \equiv k * u}$$

*Levin et al Understanding and evaluating blind deconvolution algorithms CVPR 2009

A little problem

- The TV prior has a little flaw

$$\min_{u,k} \lambda|\nabla u|_{2,1} + \frac{1}{2}|f - k * u|_2^2$$

- Compare the true solution (u, k) with the no-blur solution (f, δ)

$$\boxed{f \equiv \delta * f \equiv k * u}$$

*Levin et al Understanding and evaluating blind deconvolution algorithms CVPR 2009

A little problem

- The TV prior has a little flaw

$$\min_{u,k} \lambda |\nabla u|_{2,1} + \frac{1}{2}|f - k * u|_2^2$$

- Compare the true solution (u, k) with the no-blur solution (f, δ)

$$\boxed{f \equiv \delta * f \equiv k * u}$$

- Only the image prior is left in the cost, but the prior favors the no-blur solution!

$$\boxed{|\nabla f|_{2,1} \leq |\nabla u|_{2,1}}$$

*Levin et al Understanding and evaluating blind deconvolution algorithms CVPR 2009

Revisiting total variation BD

- The complete problem statement is

$$\min_{u,k} \lambda |\nabla u|_{2,1} + \frac{1}{2} \left| f - k * u \right|_2^2$$

$$\text{s.t. } k \succeq 0, \qquad |k|_1 = 1$$

  where the constraints on the blur kernel ensure that the blur is non negative and adds up to 1 (or, equivalently, its L_1 norm is 1)

- The L_1 norm constraint fixes the scale ambiguity between u and k; without it the minimization would make the scale of u tend to 0 and make the image prior irrelevant

Fixing the scale ambiguity

- The complete problem statement is

$$\min_{u,k} \lambda |\nabla u|_{2,1} + \frac{1}{2} \left|f - k * u\right|_2^2$$
$$\text{s.t. } k \geq 0, \qquad |k|_1 = 1$$

- If all we need is to fix the scale of k, then L_p norms could be used too

- Would $p \neq 1$ make a difference?

L_p normalization

- The new problem statement is

$$\min_{z,w} \lambda |\nabla z|_{2,1} + \frac{1}{2} \left| f - w * z \right|_2^2$$

$$\text{s.t. } w \succeq 0, \qquad |w|_p = 1$$

L_p normalization

- The new problem statement is

$$\min_{z,w} \lambda |\nabla z|_{2,1} + \frac{1}{2} \left| f - w * z \right|_2^2$$

$$\text{s.t. } w \succeq 0, \qquad |w|_p = 1$$

- Now substitute $k = w / |w|_1$ and $u = |w|_1 z$

L_p normalization

- The new problem statement is

$$\min_{z,w} \lambda |\nabla z|_{2,1} + \frac{1}{2} \left| f - w * z \right|_2^2$$
$$\text{s.t. } w \geq 0, \qquad |w|_p = 1$$

- Now substitute $k = w / |w|_1$ and $u = |w|_1 z$

- Obtain the equivalent formulation

$$\min_{u,k} \lambda |k|_p |\nabla u|_{2,1} + \frac{1}{2} \left| f - k * u \right|_2^2$$
$$\text{s.t. } k \geq 0, \qquad |k|_1 = 1$$

L_p normalization

- The new problem statement is

$$\min_{z,w} \lambda |\nabla z|_{2,1} + \frac{1}{2} \left| f - w * z \right|_2^2$$
$$\text{s.t. } w \geq 0, \qquad |w|_p = 1$$

- Now substitute $k = w / |w|_1$ and $u = |w|_1 z$

- Obtain the equivalent formulation

$$\min_{u,k} \lambda |k|_p |\nabla u|_{2,1} + \frac{1}{2} \left| f - k * u \right|_2^2$$
$$\text{s.t. } k \geq 0, \qquad |k|_1 = 1$$

which has a regularization parameter that depends on the blur L_p norm

L_p normalization

- The equivalent formulation is almost like the previous total variation form

$$\min_{u,k} \lambda |k|_p |\nabla u|_{2,1} + \frac{1}{2} \left| f - k * u \right|_2^2$$

$$\text{s.t. } k \succeq 0, \qquad |k|_1 = 1$$

L_p normalization

- The equivalent formulation is almost like the previous total variation form

$$\min_{u,k} \lambda |k|_p |\nabla u|_{2,1} + \frac{1}{2} \left| f - k * u \right|_2^2$$
$$\text{s.t. } k \succeq 0, \qquad |k|_1 = 1$$

- Let us compare now the true solution (u, k) with the no-blur solution (f, δ)

$$|k|_p |\nabla u|_{2,1} \qquad \text{vs} \qquad |\nabla f|_{2,1}$$

L_p normalization

- The equivalent formulation is almost like the previous total variation form

$$\min_{u,k} \lambda |k|_p |\nabla u|_{2,1} + \frac{1}{2} \left| f - k * u \right|_2^2$$

$$\text{s.t. } k \succeq 0, \qquad |k|_1 = 1$$

- Let us compare now the true solution (u, k) with the no-blur solution (f, δ)

$$|k|_p |\nabla u|_{2,1} \qquad \text{vs} \qquad |\nabla f|_{2,1}$$

- When $p = 2$ the term $|k|_p < 1$ if $k \neq \delta$ and this makes the LHS term small

Rescuing the TV prior

- **Theorem** Assume the gradients of the true sharp image u to be i.i.d. zero-mean Gaussian and the true blur kernel k to have finite support. Given a blurry image $f = k * u$, the new formulation favors with high probability the true blur/image pair (u, k) over the trivial no-blur pair (f, δ) for $p \geq 2$.

Jin, Roth, Favaro Normalized blind deconvolution ECCV 2018

Optimization

- We use the Frank-Wolfe algorithm and alternate between blur and image
- Advantages

1. For the first time it is possible to optimize the cost function exactly
2. Coarse to fine scheme is not needed
3. Careful initialization is not necessary (can start with $k = \delta$)
4. Regularization parameter is not changed during the iteration time
5. The formulation is convex separately in each variable

Quantitative evaluation

Table 1: Quantitative comparison on the entire SUN dataset [39] (640 blurry images).

Method	mean error ratio	maximum error ratio	failure cases
Cho & Lee [7]	9.198	113.491	224
Krishnan *et al.* [20]	12.015	142.668	475
Levin *et al.* [23]	6.695	44.171	357
Sun *et al.* [39]	2.581	35.765	44
Xu & Jia [44]	3.817	75.036	98
Perrone & Favaro [31]	2.114	8.517	**7**
Chakrabarti [4]	3.062	11.576	64
Michaeli & Irani [24]	2.617	9.185	30
Pan *et al.* [29]	**1.914**	23.279	11
`PN`	2.299	**6.764**	8
`FW`	2.195	**6.213**	8

Sum of squared differences ratio $\frac{|u - \hat{u}|^2}{|u - \hat{u}_*|^2}$

$\hat{u}_*$ estimated with GT kernel
$\hat{u}$ estimated with estimated kernel

Levin, A., Weiss, Y., Durand, F., Freeman, W.T.: Efficient marginal likelihood optimization in blind deconvolution. In: CVPR (2011)

Quantitative evaluation

Table 2: Quantitative comparison on the small BSDS dataset [1] (72 blurry images).

Method	mean error ratio	maximum error ratio	failure cases
Sun *et al.* [39]	2.648	15.152	12
Xu & Jia [44]	3.645	22.272	13
Perrone & Favaro [31]	2.093	7.493	4
Chakrabarti [4]	3.768	11.809	9
Michaeli & Irani [24]	3.458	23.001	14
Pan *et al.* [29]	2.058	13.516	3
Yan *et al.* [46]	2.022	12.237	3
L^1 normalization	2.211	7.821	3
weight decay (heuristic)	2.591	8.762	2
L^2 blur prior (classic)	2.487	7.953	4
`PN`	2.011	4.676	**0**
`FW`	**1.983**	**4.387**	**0**

Sum of squared differences ratio $\frac{|u - \hat{u}|^2}{|u - \hat{u}_*|^2}$

$\hat{u}_*$ estimated with GT kernel

$\hat{u}$ estimated with estimated kernel

Levin, A., Weiss, Y., Durand, F., Freeman, W.T.: Efficient marginal likelihood optimization in blind deconvolution. In: CVPR (2011)

Qualitative comparisons

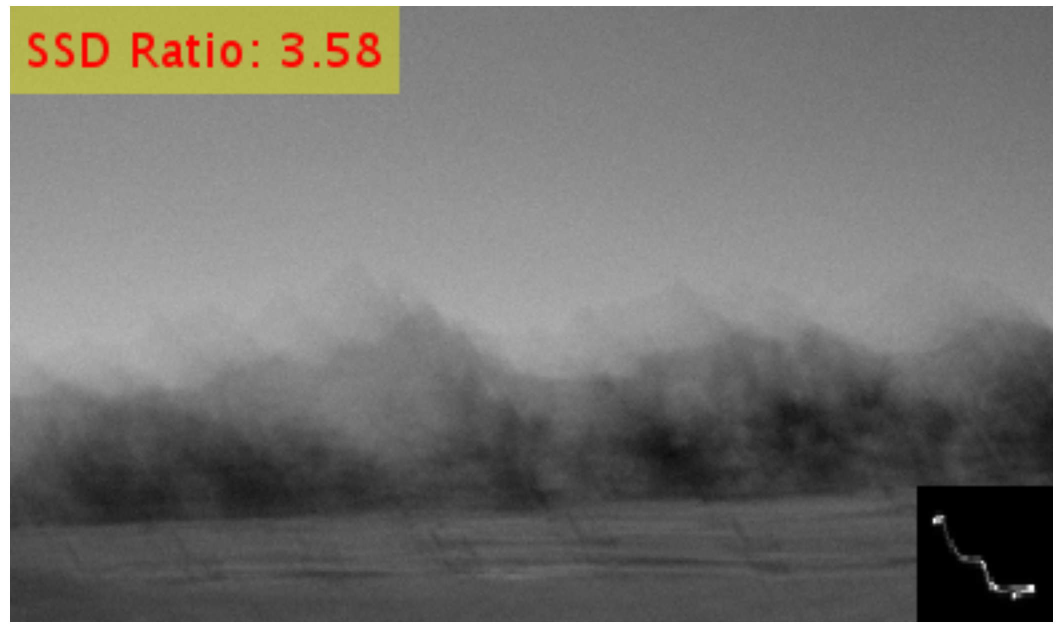

input

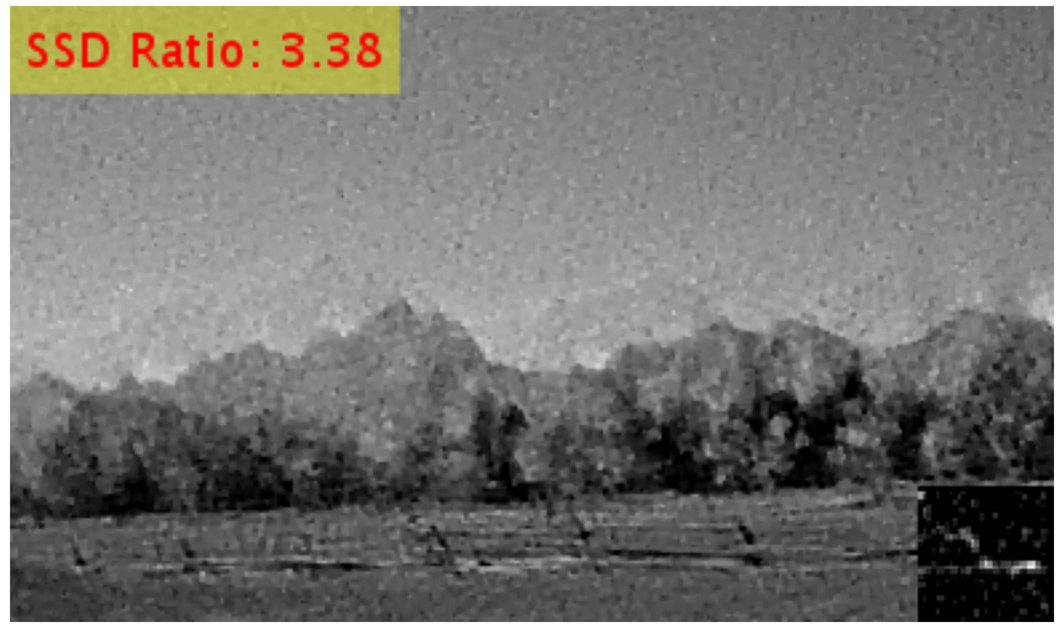

Xu and Jia 2010

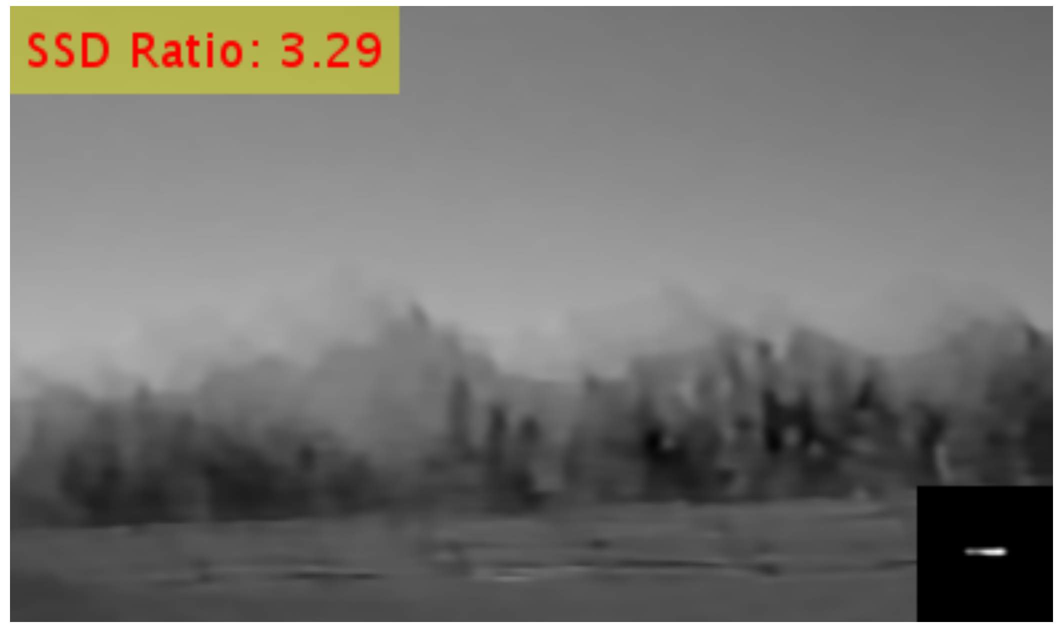

Chakrabarti 2016

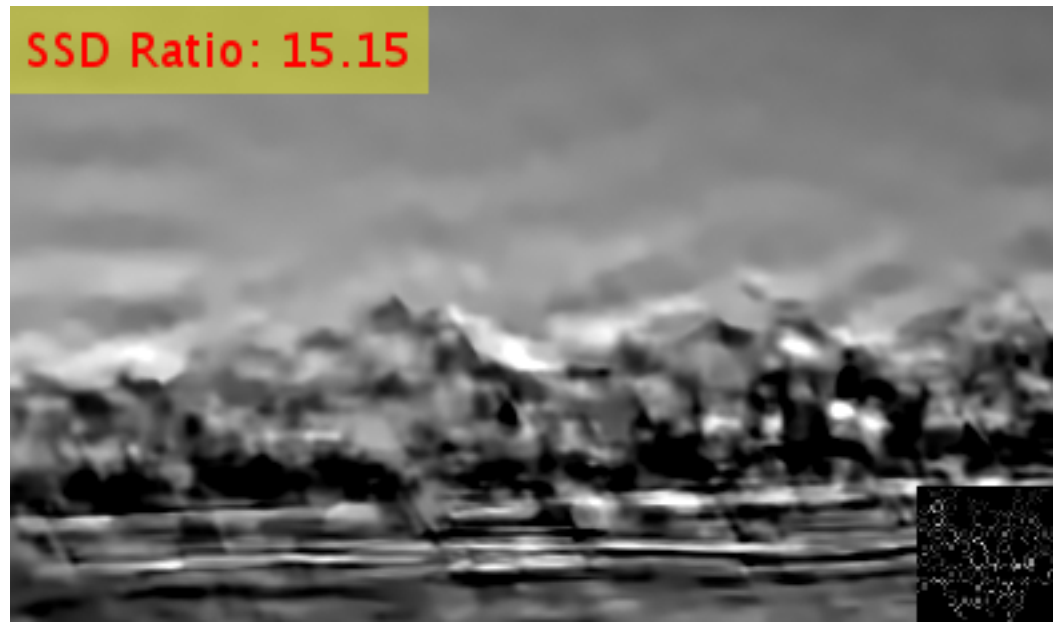

Sun et al 2013

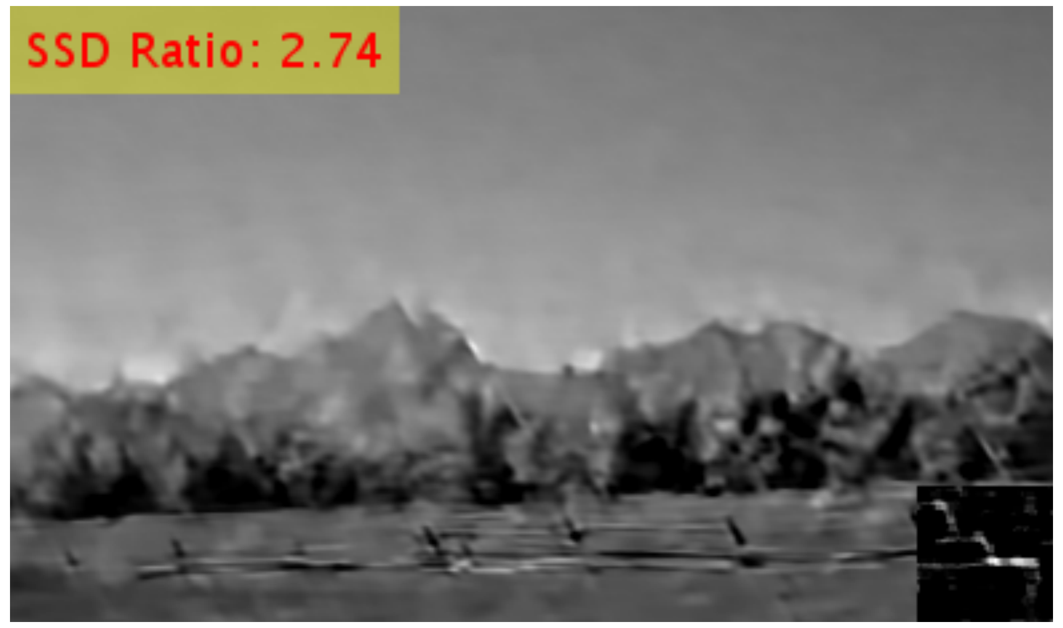

Michaeli and Irani 2014

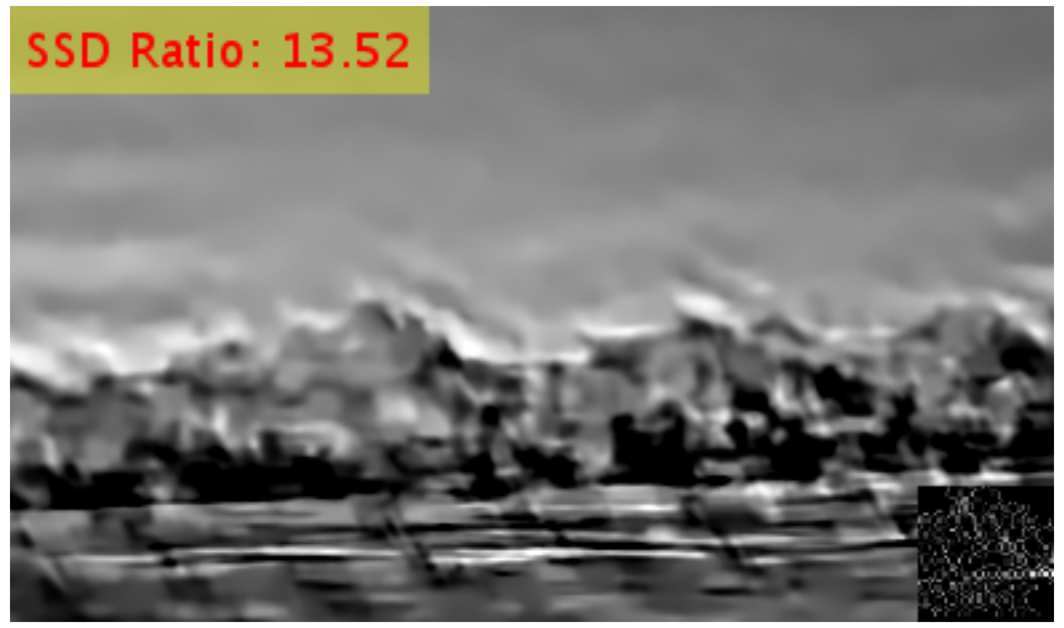

Pan et al 2016

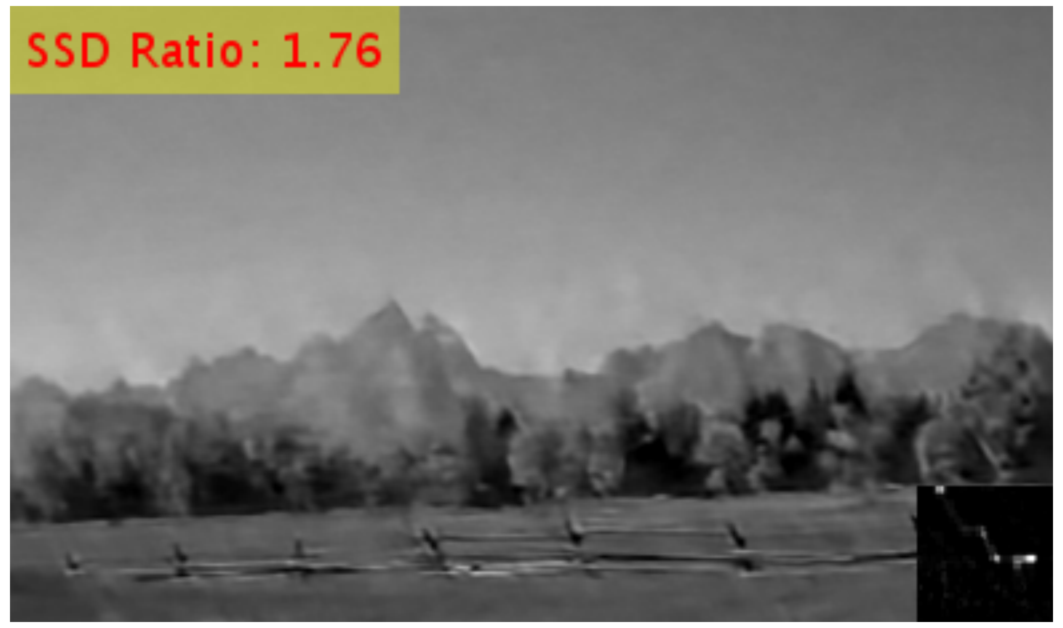

Perrone and Favaro 2016

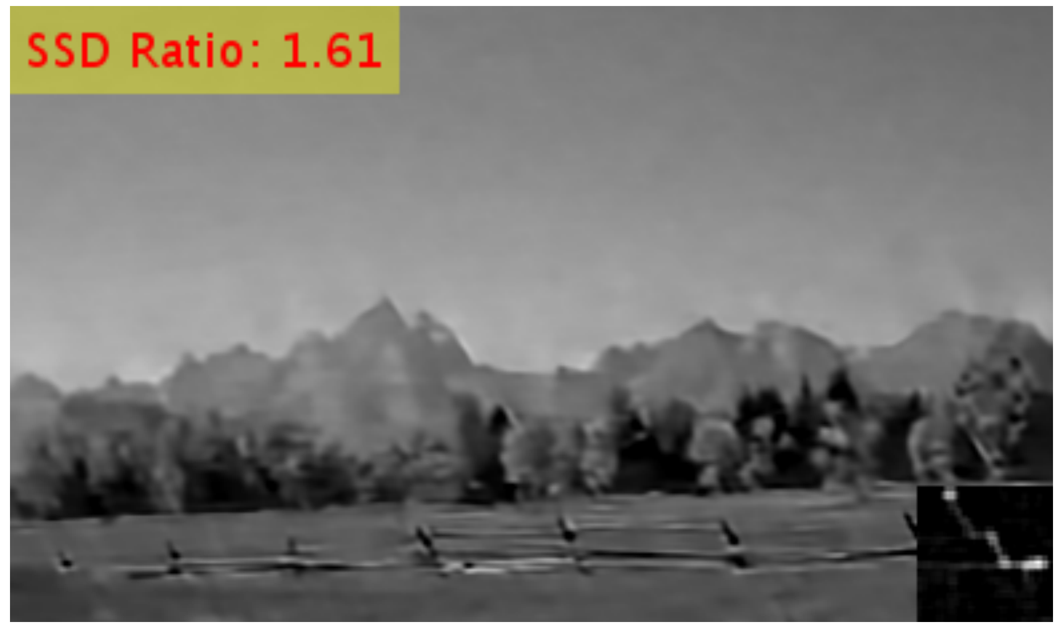

FW

Worst cases in real images

input

Pan et al 2016

PN

Conclusions

- Deep learning methods will probably prevail in the end
- There are some limitations that might take time to address
- Can we trust that the reconstruction is not a hallucination of the data?

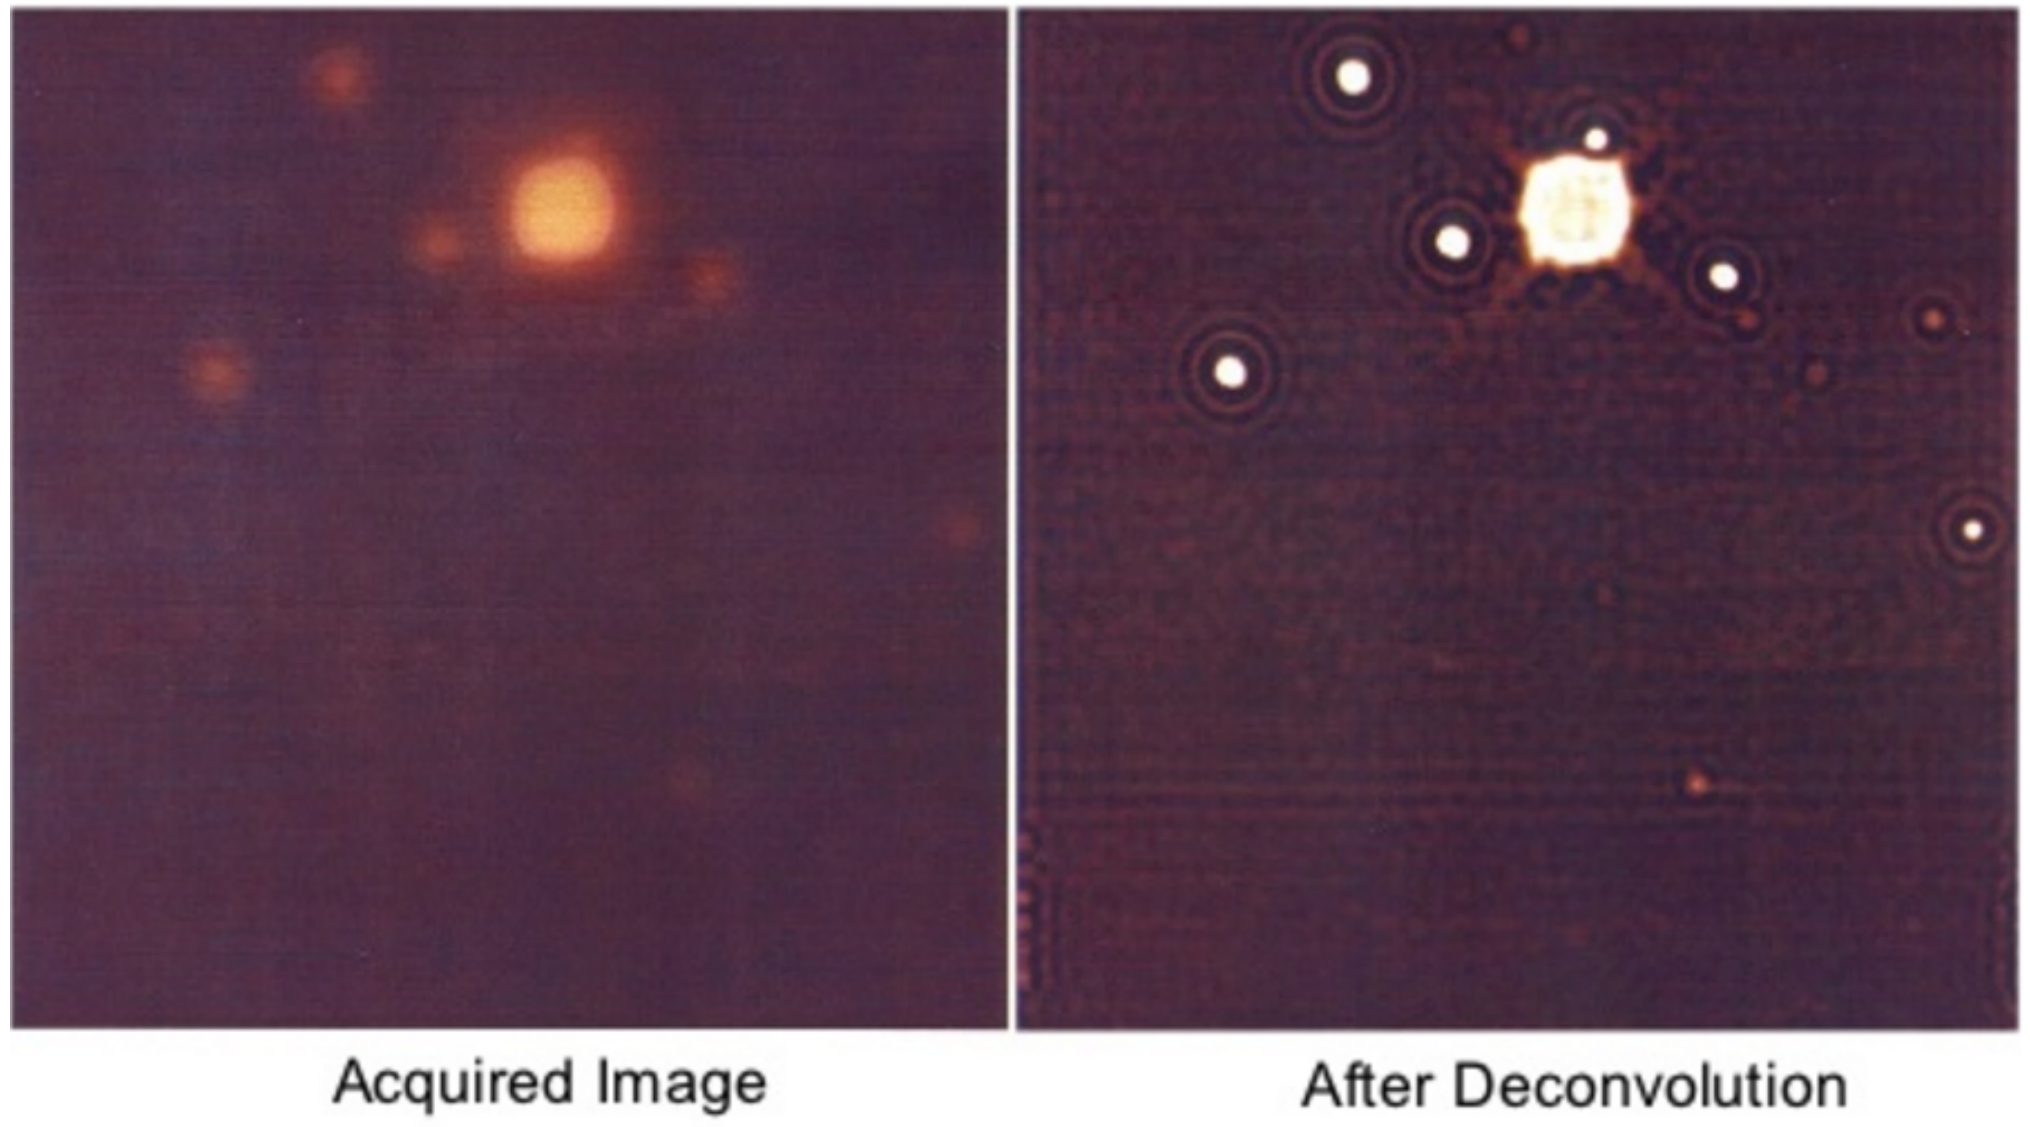

Conclusions

- In contrast model-based methods are easily interpretable

- There is still quite a bit to do even with simple formulations

- It pays to pay attention to the details

Blurry Input

7 Frame Estimates

Blurry Input

7 Frame Estimates