Bayesian Deblurring with Integrated Noise Estimation

Uwe Schmidt · Kevin Schelten · Stefan Roth   Department of Computer Science, TU Darmstadt, Germany

TECHNISCHE UNIVERSITÄT DARMSTADT

Image Blur & Our Deblurring Approach

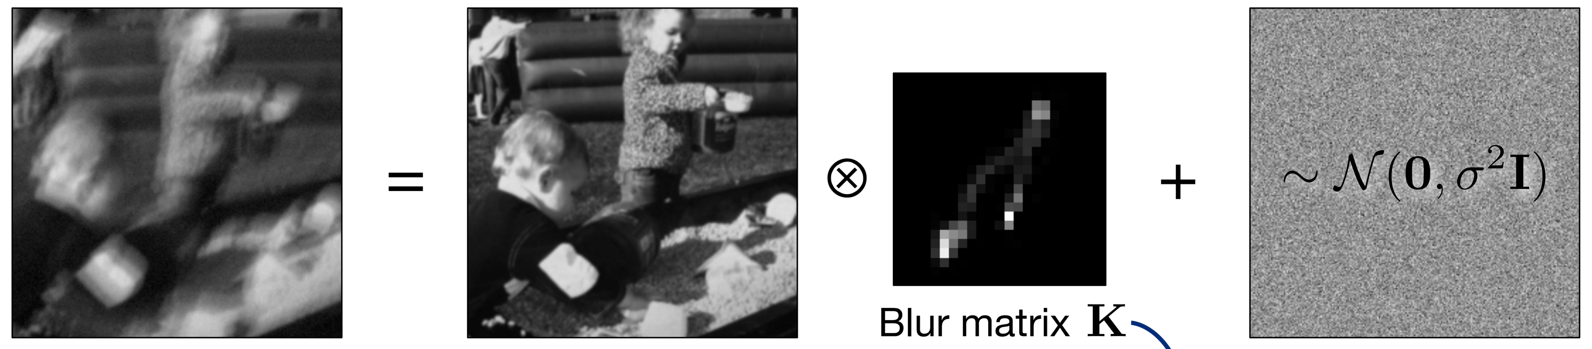

Input: Blurred image $\mathbf{y}$ — Desired: Original image $\mathbf{x}$ — Blur matrix $\mathbf{K}$ or kernel $\mathbf{k}$ with $\mathbf{Kx} \equiv \mathbf{k} \otimes \mathbf{x}$ Assumed known — Gaussian noise with unknown noise level σ

Output: Deblurred image $\hat{\mathbf{x}}$

$$\hat{\mathbf{x}} = \arg\min_{\tilde{\mathbf{x}}} \int_{\mathbf{x}} \int_{\sigma} \|\tilde{\mathbf{x}} - \mathbf{x}\|^2 \, p(\mathbf{x}|\mathbf{y}, \mathbf{K}, \sigma) \, p(\sigma|\mathbf{y}, \mathbf{K}) \, \mathrm{d}\mathbf{x} \, \mathrm{d}\sigma$$

Noise level σ is integrated out — Posterior (image restoration) — Noise prior (here uniform)

Bayesian minimum mean squared error (*MMSE*) / least squares (*BLS*) estimate → integrate over $\mathbf{x}$

Problem

- **Noise level** of corrupted image **not known in practice**
  → mostly assumed to be known anyway
- **Accurate noise level important** for image restoration
  → noise **estimation** from corrupted (i.e. blurred) image **difficult**
- Here: non-blind deblurring
  → image blur is known or estimated beforehand
  → important component of blind deblurring [3]

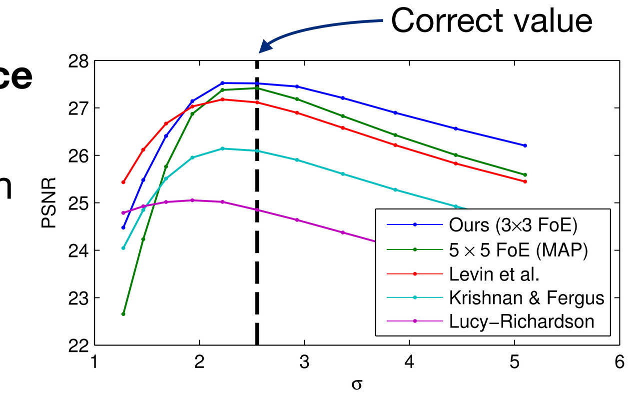

Deblurring performance for various assumed noise levels
(all methods suffer from incorrect value)

Approach

- **Sampling-based Bayesian inference with learned image prior**
  → allows to treat the **noise level as nuisance parameter** and integrate it out, no preceding noise estimation necessary
  → enables MMSE/BLS estimation for improved restoration performance with learned high-order MRF prior [8]
  → **additional benefit**: integrating over images also allows to **obtain noise estimate**
  → noise estimation specialized to deblurring and only based on natural image prior
- **No parameters to tune**
  (e.g., no noise-dependent regularization weights)
- Special case: blind denoising with noise estimation
- Applicable to non-uniform blur

Code available at **goo.gl/tqqsk**

Results

- **State-of-the-art performance** (deblurring *and* noise estimation)
- Performance difference on average negligible to the case where noise level is known

Bayesian Inference with Integrated Noise Estimation

0. Bayesian approach

$$p(\mathbf{x}|\mathbf{y}, \mathbf{K}, \sigma) \propto p(\mathbf{y}|\mathbf{x}, \mathbf{K}, \sigma) \cdot p(\mathbf{x})$$

Gaussian likelihood $\mathcal{N}(\mathbf{y}; \mathbf{Kx}, \sigma^2\mathbf{I})$; Natural image prior $p(\mathbf{x})$

1. High-order Markov Random Field (MRF) prior

- Fields of Experts (FoE) [6] model $p(\mathbf{x}) = \frac{1}{Z} e^{-\epsilon\|\mathbf{x}\|^2/2} \prod_{c \in \mathcal{C}} \prod_{i=1}^{N} \phi(\mathbf{J}_i^{\mathrm{T}}\mathbf{x}_{(c)}; \boldsymbol{\alpha}_i)$ with
  $\phi(\mathbf{J}_i^{\mathrm{T}}\mathbf{x}_{(c)}; \boldsymbol{\alpha}_i) = \sum_{z_{ic}=1}^{J} \alpha_{iz_{ic}} \mathcal{N}(\mathbf{J}_i^{\mathrm{T}}\mathbf{x}_{(c)}; 0, \eta^2_{iz_{ic}})$ (Gaussian scale mixture experts [8])
  Discrete indicator variables
- Latent variable representation $p(\mathbf{x}, \mathbf{z})$ with $p(\mathbf{x}) = \sum_{\mathbf{z}} p(\mathbf{x}, \mathbf{z})$, useful for inference

2. MMSE estimation

- Use MMSE instead of prevalent MAP estimate and analyze its aptitude for deblurring
- Extend sampling-based MMSE estimation from denoising [8] $(\mathbf{K} = \mathbf{I})$ to deblurring:

$$\hat{\mathbf{x}} = \arg\min_{\tilde{\mathbf{x}}} \int \|\tilde{\mathbf{x}} - \mathbf{x}\|^2 p(\mathbf{x}|\mathbf{y}, \mathbf{K}, \sigma) \, \mathrm{d}\mathbf{x} = E[\mathbf{x}|\mathbf{y}, \mathbf{K}, \sigma]$$

- In practice use $p(\mathbf{x}, \mathbf{z}|\mathbf{y}, \mathbf{K}, \sigma) \propto p(\mathbf{y}|\mathbf{x}, \mathbf{K}, \sigma) \cdot p(\mathbf{x}, \mathbf{z})$ ($\mathbf{z}$ eventually marginalized out)

3. Integrating noise estimation

- In contrast to previous work, assume noise level to be *unknown in the spatial domain:*

$$p(\mathbf{x}, \mathbf{z}, \sigma|\mathbf{y}, \mathbf{K}) \propto p(\mathbf{y}|\mathbf{x}, \mathbf{K}, \sigma) \cdot p(\mathbf{x}, \mathbf{z}) \cdot p(\sigma|\mathbf{y}, \mathbf{K})$$

- Obtain MMSE estimate by integrating (marginalizing) σ out:

$$\hat{\mathbf{x}} = \arg\min_{\tilde{\mathbf{x}}} \iint \|\tilde{\mathbf{x}} - \mathbf{x}\|^2 p(\mathbf{x}, \sigma|\mathbf{y}, \mathbf{K}) \, \mathrm{d}\mathbf{x} \, \mathrm{d}\sigma = E[\mathbf{x}|\mathbf{y}, \mathbf{K}]$$

Benefit:

- Can also obtain MMSE estimate of the noise level σ (integrate $\mathbf{x}$ out):

$$\hat{\sigma} = \arg\min_{\tilde{\sigma}} \iint \|\tilde{\sigma} - \sigma\|^2 p(\mathbf{x}, \sigma|\mathbf{y}, \mathbf{K}) \, \mathrm{d}\mathbf{x} \, \mathrm{d}\sigma = E[\sigma|\mathbf{y}, \mathbf{K}]$$

No further assumptions necessary

4. Sampling-based inference

- Approximate MMSE estimate through sequence of samples from the joint posterior:

$$\left\{\{\mathbf{x}^{(1)}, \mathbf{z}^{(1)}, \sigma^{(1)}\}, \ldots, \{\mathbf{x}^{(T)}, \mathbf{z}^{(T)}, \sigma^{(T)}\}\right\} \sim p(\mathbf{x}, \mathbf{z}, \sigma|\mathbf{y}, \mathbf{K})$$

- Obtain samples of posterior through Gibbs sampling from conditional distributions
  - $p(\mathbf{z}|\mathbf{x}, \mathbf{y}, \mathbf{K}, \sigma)$ decomposes into univariate discrete distributions [8]
  - $p(\mathbf{x}|\mathbf{z}, \mathbf{y}, \mathbf{K}, \sigma)$ is a multivariate Gaussian
  - $p(\sigma|\mathbf{x}, \mathbf{z}, \mathbf{y}, \mathbf{K})$ is a Gamma distribution: $\mathcal{G}\left(\frac{1}{\sigma^2}; \frac{n}{2} + 1, \frac{2}{\|\mathbf{y} - \mathbf{Kx}\|^2}\right)$
- MMSE estimate can be approximated by averaging samples after B burn-in iterations
- Alternatively, use Rao-Blackwellized MMSE estimation [4] (more efficient) by averaging conditional expectations:

$$\hat{\sigma}_{\mathrm{RB}} \approx \frac{1}{T - B} \sum_{t=B+1}^{T} \sqrt{\frac{\|\mathbf{y} - \mathbf{Kx}^{(t-1)}\|^2}{n + 2}} \qquad \hat{\mathbf{x}}_{\mathrm{RB}} \approx \frac{1}{T - B} \sum_{t=B+1}^{T} \mathbf{Q}^{-1}_{\mathbf{z}^{(t)}} \mathbf{K}^{\mathrm{T}} \frac{\mathbf{y}}{(\sigma^{(t)})^2}$$

#pixels of y; Precision of Gaussian $p(\mathbf{x}|\mathbf{z}, \mathbf{y}, \mathbf{K}, \sigma)$

Experiments

Qualitative – Preserves textured *and* smooth regions

Original — Blurred, PSNR = 23.89dB — **Ours** (3×3 FoE), PSNR = 29.05dB — **Ours** (pairwise MRF), PSNR = 28.84dB
5×5 FoE (MAP) [6], PSNR = 28.81dB — Levin et al. [2], PSNR = 28.54dB — Krishnan & Fergus [1], PSNR = 28.36dB — Lucy-Richardson [5], PSNR = 27.01dB

Quantitative – State-of-the-art results for non-blind deblurring and blind denoising, including noise estimation in both applications

Deblurring	PSNR (dB) GT	PSNR (dB) NE	Estimate $\hat{\sigma}$ avg.	Estimate $\hat{\sigma}$ $\langle\epsilon\rangle$
Lucy-Richardson [5]	25.38	25.34	NA	
Krishnan & Fergus [1]	26.97	26.86		
Levin et al. [2]	28.03	27.96		
5×5 FoE (MAP) [6]	28.44	28.33		
Zoran & Weiss [9]	NA		2.52	8.21%
Ours (pairwise MRF)	28.24	28.17	2.55	**4.02%**
Ours (3×3 FoE)	**28.66**	**28.61**	2.64	4.28%

Deblurring 64 test images with noise level $\sigma = 2.55$

Denoising	PSNR (dB) GT	PSNR (dB) NE	Estimate $\hat{\sigma}$ avg.	Estimate $\hat{\sigma}$ $\langle\epsilon\rangle$
5×5 FoE (MAP) [6]	27.44	—	NA	
5×5 FoE (MAP) [7]	27.86	—		
Pairwise MRF (MMSE) [8]	27.54	—		
3×3 FoE (MMSE) [8]	**27.95**	—		
Zoran & Weiss [9]	NA		23.16	8.8%
Ours (pairwise MRF)	—	27.16	22.81	10.1%
Ours (3×3 FoE)	—	**27.88**	24.20	**5.8%**

Denoising 68 test images with noise level $\sigma = 25$

Comparable relative performance also for *SSIM*

Average image restoration quality (PSNR) and average relative noise estimation error ($\langle\epsilon\rangle$)

GT: Using ground truth noise level   **NE**: Noise level integrated out (ours) or estimated beforehand (with Zoran & Weiss [9])

Sorted PSNR differences for the 64 deblurred images ($\sigma = 2.55$) between our method (3x3 FoE) and all others. *Below dashed line → ours better*

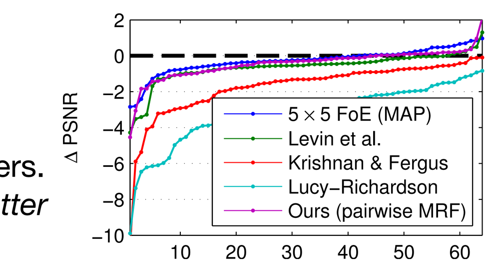

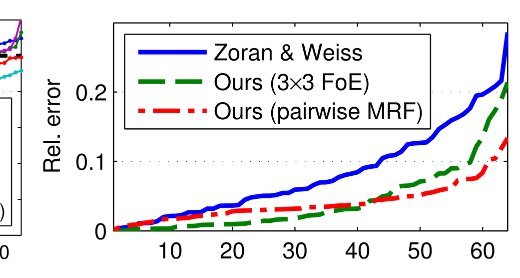

Noise estimation results for the 64 deblurred images ($\sigma = 2.55$) in terms of average relative estimation error. Each curve sorted separately.

References

[1] D. Krishnan and R. Fergus. Fast image deconvolution using hyper-Laplacian priors. In NIPS*2009.
[2] A. Levin, R. Fergus, F. Durand, and W. T. Freeman. Image and depth from a conventional camera with a coded aperture. ACM T. Graphics, 26(3), July 2007.
[3] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman. Understanding and evaluating blind deconvolution algorithms. In CVPR 2009.
[4] G. Papandreou and A. Yuille. Gaussian sampling by local perturbations. In NIPS*2010
[5] W. Richardson. Bayesian-based iterative method of image restoration. J. Opt. Soc. America, 62(1):55–59, 1972.
[6] S. Roth and M. J. Black. Fields of experts. IJCV, 82(2):205–229, Apr. 2009.
[7] K. G. G. Samuel and M. F. Tappen. Learning optimized MAP estimates in continuously-valued MRF models. In CVPR 2009.
[8] U. Schmidt, Q. Gao, and S. Roth. A generative perspective on MRFs in low-level vision. In CVPR 2010.
[9] D. Zoran and Y. Weiss. Scale invariance and noise in natural images. In ICCV 2009.