Total Variation Blind Deconvolution: The Devil is in the Details

UNIVERSITÄT BERN

Daniele Perrone and Paolo Favaro

{perrone, favaro}@iam.unibe.ch

IEEE 2014 Conference on Computer Vision and Pattern Recognition

CVPR 2014

Blind Deconvolution

blurry image | PSF | sharp image | noise

Input | Outputs

Task: Estimate a sharp image u and the point spread function (PSF) k from a single blurry image f.

Clearing the fog: Several approaches, which one is right?

Method 1:
Sharp image marginalization.

Method 2:
Energy reweighting.

Method 3:
Use of filtered images.

Method 4:
Edge enhancement.

Levin et. al. [1] show that these algorithms are not supposed to work. Yet they do.

We address the following question: Why do these algorithms work despite theoretical results showing that they cannot?

Summary of our findings

1) The findings of Levin et. al. [1] are correct: the **exact** minimization of a large class of energies with texture priors leads to a no-blur solution.

2) Many algorithms still work because they do **not** minimize the claimed cost.

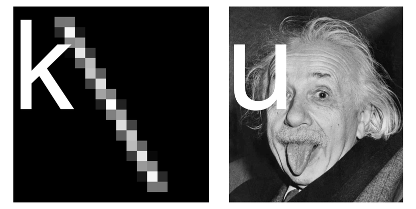

3) Delayed normalization (**scaling**) of the blur is key.

Total Variation Blind Deconvolution

Blind deconvolution is typically solved by minimizing a variation of the following cost function.

$$\underbrace{\arg\min_{u,k\geq 0,\mathbf{1}^T k=1} \frac{1}{2}\|u*k-f\|_2^2}_{\text{Data fitting term}} + \overbrace{\lambda}^{\text{Regularization parameter}}\underbrace{\|u\|_{BV}}_{\text{Total variation regularization}} \quad (1)$$

Theorem 1: A large class of regularization terms, such as the total variation, favor the blurry image and not the sharp one (extension of the results in Levin et. al. [1]).

$$\left\| \right.$$  $$\left. \right\|_{BV} \leq \left\| \right.$$ 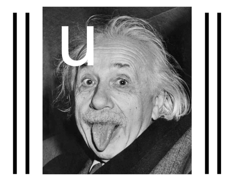 $$\left. \right\|_{BV}$$

Projected Alternating Minimization

A common approach to minimize (1) is to alternate between the following steps.

1 $$u^t \leftarrow \min_u \frac{1}{2}\|u*k^{t-1}-f\|_2^2 + \lambda\|u\|_{BV}$$

2.1 $$\hat{k}^t \leftarrow \min_k \frac{1}{2}\|u^t*k-f\|_2^2$$

2.2 $$k^t \leftarrow \frac{\max\{\hat{k}^t,0\}}{\|\max\{\hat{k}^t,0\}\|_1}$$

Proposition: The energy in (1) is not minimized by the Projected Alternating Minimization (PAM) algorithm.

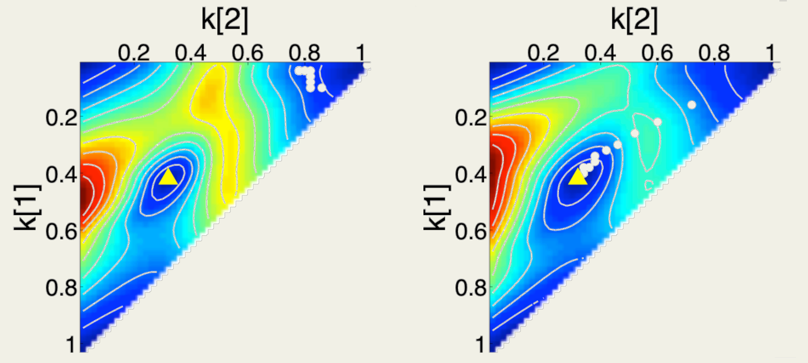

$\lambda = 0.001$ $\lambda = 0.01$

Theorem 2: for a 1D step function blurred with a blur of support equal to 3 pixels and for $\lambda > \lambda_0$, the PAM algorithm estimates the true blur in two steps.

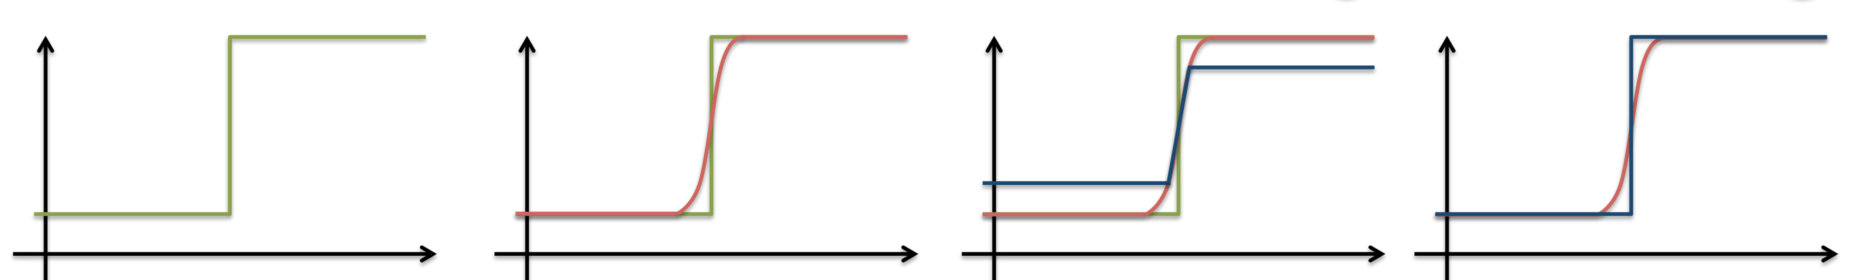

Color legend: Desired u Blurred image f Current estimate u^t

Exact Alternating Minimization

1 $$u^t \leftarrow \min_u \frac{1}{2}\|u*k^{t-1}-f\|_2^2 + \lambda\|u\|_{BV}$$

2 $$k^t \leftarrow \arg\min_{k\geq 0,\mathbf{1}^T k=1} \frac{1}{2}\|u^t*k-f\|_2^2$$

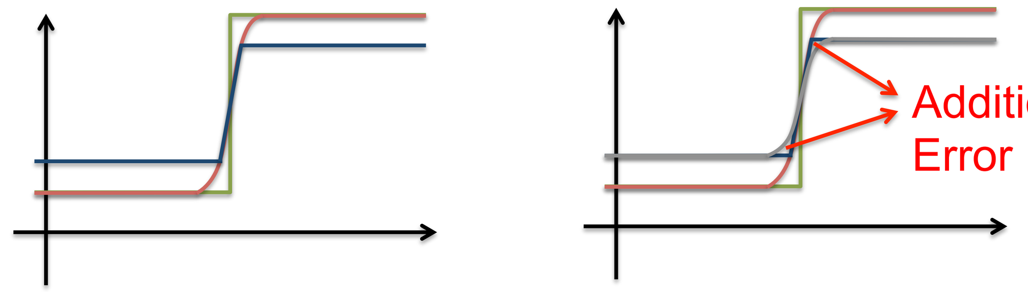

Color legend: Desired u Blurred image f Current estimate u^t Blurred u^t

Theorem 3: for the same signal and value of λ of Theorem 2 the Alternating Minimization (AM) algorithm either converges to the no-blur solution or becomes unstable.

Experimental Validation

The PAM algorithm without any additional heuristics achieves state-of-the-art results.

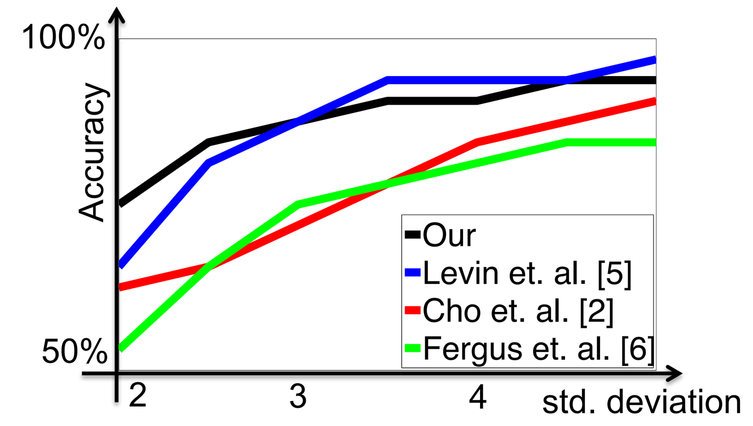

Blurred

Our result

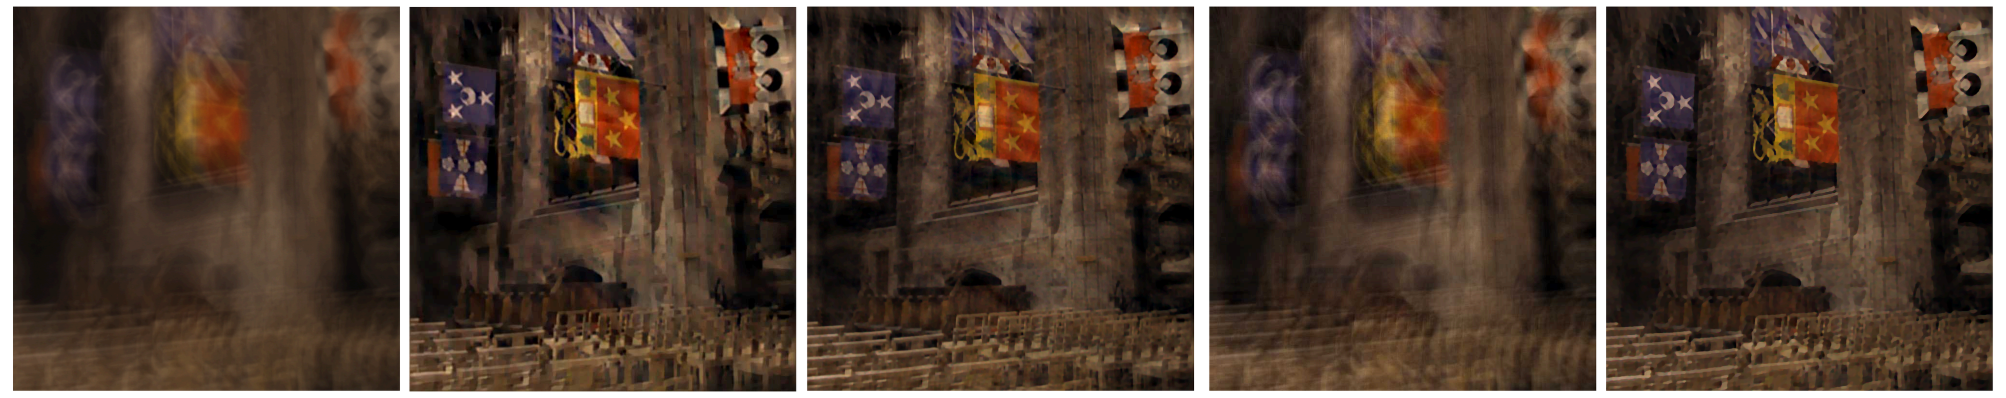

Blurred | Cho et. Al. [2] | Xu et. Al. [3] | Whyte et. Al. [4] | Our result

Code available at:
http://www.cvg.unibe.ch/dperrone/tvdb/

References:

[1] Levin, A., Weiss, Y., Durand, F. and Freeman, W.T., "Understanding Blind Deconvolution Algorithms," *PAMI* 2011.

[2] Cho, S., and Lee, S., “Fast motion deblurring”, In *ACM SIGGRAPH Asia 2009.*

[3] Xu, L., and Jia, J., “Two-phase kernel estimation for robust motion deblurring”. *ECCV 2010*.

[4] Whyte, O., Sivic, J., Zisserman, A., and Ponce, J., “Non-uniform Deblurring for Shaken Images”, *IJCV 2012.*

[5] Levin, A., Weiss, Y., Durand, F. and Freeman, W.T., "Efficient marginal likelihood optimization in blind deconvolution”, *CVPR 2011*.

[6] Fergus, R., Singh, B., Hertzmann,, A., Roweis, S. T. and Freeman, W.T., “Removing camera shake from a single photograph”, In *ACM SIGGRAPH 2006.*