Creating a Low-Light, Motion-Blurred Image Processing Pipeline

Alexander Evelson
aevelson@stanford.edu

Andrea Baldioceda
abaldio@stanford.edu

Motivation

- When processing images, there is a consistent trade-off between blur and noise. Given that noise is high-frequency content, denoising by nature blurs an image. By contrast, deblurring introduces noise into the image.
- In low-light conditions, the situation is compounded by the need to brighten an image. The process of brightening the image itself introduces additional noise.
- Taking images while in motion, especially of wildlife, though not necessarily of course, introduces additional blur.
- Past projects have considered one issue or the other, but few have attempted to tackle the basket of issues in an end-to-end manner.
- We thus set out to create a pipeline that takes these dark, blurry images and processes them into sharp, bright images.

Image Processing Pipeline

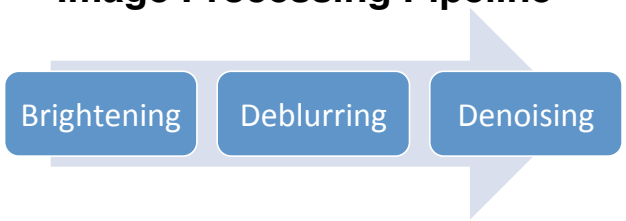

Image Brightening

- To brighten the image, we used a standard gain approach. At the end of this step, the image was still blurry and now additionally noisy.

Dark, Blurry Image

Brightened

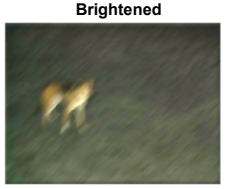

References

[1]B. Choi, D. Kim, W. Kyung and Y. Ha, "Multi-spectral flash imaging under low-light condition using optimization with weight map."
[2]K. De and V. Masilamani, "Image Sharpness Measure for Blurred Images in Frequency Domain."
[3]Z. Hu, S. Cho, J. Wang and M. Yang, "Deblurring Low-Light Images with Light Streaks."
[4]F. Li, Y. Ji and J. Yu, "A Hybrid Camera System for Low-Light Imaging."
[5]W. Lai et. al, "A Comparative Study for Single Image Blind Deblurring."
[6]S. McCloskey, "Temporally coded flash illumination for motion deblurring."
[7]J. Wang, T. Xue, J. T. Barron and J. Chen, "Stereoscopic Dark Flash for Low-light Photography."

Image Deblurring

Known Blur Kernel

- We first operated under the assumption that we knew the details of the blur kernel to test how efficiently the pipeline would work under simple conditions. We used two approaches to deblur the image in this case.

Wiener Deconvolution

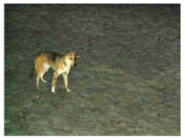

ADMM

- In this instance, we noticed that ADMM had fewer artifacts but was not as sharp as Wiener Deconvolution.

Unknown Blur Kernel

- Our algorithm, using a fixed length, selected the theta for which the output had the most frequencies below a given threshold (low-pass filter). Then, for that optimal theta, we selected the length with the most frequencies in an optimal band (mid-pass filter).

Wrong Theta

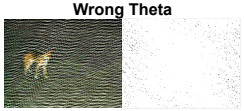

Right Theta, Short Length

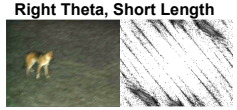

Right Theta, Long Length

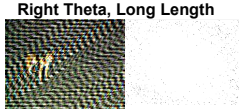

Proper Output

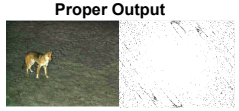

MATLAB Blind Deconvolution

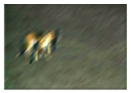

PSNR compared to Wiener
MATLAB Blind 26.8795 (integer len & theta) 23.7992 (non-integer len & theta)
Our Algorithm Inf (integer len & theta) 42.6849 (non-integer len & theta)

- Our approach outperformed MATLAB's and approximated the same output as the Wiener Deconvolution with the actual parameters of the motion blur.

Image Denoising

- Once we deblurred the images, we implemented three denoising methods to see if we could remove the noise from the brightening and deblurring parts of the process without simultaneously overblurring the image and setting us back. For our baseline, we used a simple Gaussian filter. We hoped to improve on that with a bilateral filter and using non-local means.

Gaussian Filter

Bilateral Filter

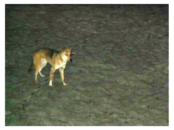

Non-Local Means

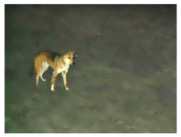

- Of the approaches, non-local means reduced the noise most significantly. However, it also took the longest time and overly blurred most of the image. For most cases, it thus appears that a bilateral filter strikes the proper balance. However, to generate an image where the subject, such as an animal, is more focused than the background, NLM is appropriate.

Discussion and Concluding Thoughts

- Our experience demonstrated that even a dark, blurry image can be processed into an appealing one with appropriate brightness and sharpness, but without too much noise.
- Even our best outputs had some artifacts, and the tradeoff between noise and blur was evident throughout, but the tradeoff was not so drastic as to prevent the results from being significantly better than the starting image.
- Our basic approach to motion deblurring outperformed MATLAB's function, demonstrating the opportunities for blind deconvolution. Many of the state-of-the-art approaches rely on deep learning, and that should allow for a relaxation of constraints (such as the one we used – motion blur in a single direction with a magnitude of less than 30 pixels).
- Our approach considered the scenario where the motion blur is uniform across the entire image. We would also be interested in a scenario where the subject of the image is in motion but the rest is still. This would be interesting to consider in the context of low-light wildlife photography.