

Problem description

Multi-frame denoising of still images

MM Solution AD.

A sequence of 8 images are captured quickly one-by-one while camera is held in hand and unintentional hand shake is present. When capturing in low light conditions ($<20\text{Lux}$), the images appear blurred and noisy because the camera uses long exposure (1/15..1/5 sec) and high gain (ISO). The target is to combine multiple frames to produce an image with less noise and less (ideally no) blur. The processing algorithm should not be too computationally expensive. As usual in the image processing, the target is to get a good looking image, rather than removing the noise and the blur to achieve the ground-truth latent image. So any brain-hacking tricks are acceptable.

Good-looking image means:

- no unnatural artifacts. Better little less sharpness than having artifacts (e.g. ringing).
- little high-frequency noise (up to STD about 1-2 for 8-bit image) creates an impression of sharper image
- the brain is extremely sensitive to the quality of straight edges and less sensitive to noise, present in image areas with fine details (grass, leaves).

A cropped part of an image captured at 10 Lux with the strongest blur within the series is presented below.



Next image presents the image with least blur in the same sequence.



Noise characteristics

STD = 6.5 (from 8bit pixel values)

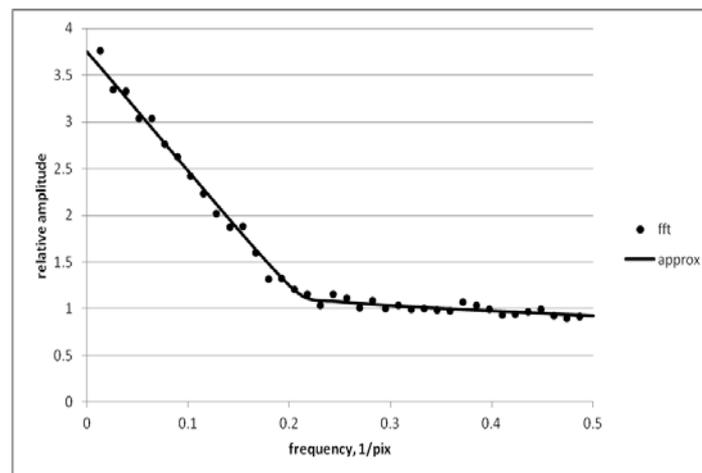
STD changes at different brightness about 2 times. The presented value is the maximum STD, observed at pixel values about 64..128

The noise is Gaussian, but not white - Low frequencies amplitudes are higher.

The next figure show a noise spectrum, measured on a flat (without details) noisy part of one of the images in the same sequence. This is a "pink noise", typical for electronic circuits like the analog amplifiers embedded in the image sensor.

More noise or less brightness can be expected at lower light, so there is no upper bound of the noise STD - the more noise the algorithm can handle, the better.

For sure, the noise in each individual image is more than the edges we would like to enhance and see in the output image.



noise spectrum

Blur characteristics

Blur kernel changes smoothly within the image - it is not same for the whole image, as opposed to the common assumption in most de-blur articles. This is caused by 2 facts:

- camera shake is mainly rotational
- contemporary image sensors use rolling shutter, i.e. different rows of the image are exposed in different moments, for same amount of time. Hence, the camera motion during exposures of the rows varies.

In typical hand-held capture case, the blur is up to about 20pix at 4000x3000 image resolution, at 1/15 sec exposure time, close to line.

Alignment

Appart from blur, hand shake causes the images not to be aligned to each other. Alignment is NOT part of that project. We have an algorithm, which aligns the images with up to 2pix error at 4000x3000 image resolution.

Preprocessing

The images are already pre-processed - color conversion and gamma correction is applied, so they are not in linear space (lightness-vs-pixel value is not linear). Finding the gamma applied is possible, but hard, so better to avoid it. The processing parameters are same for all the images in the sequence, including exposure time and gain.

Applying different exposure time and gain to the images is possible, but it would lead to different lightness-vs-pixel response of each image and differences in color. Thus, worse alignment and even worse moving-objects detection.

Moving-objects detection and handling is not part of this project, but yet, it happens somewhere, so we have to keep in mind its presence. Basically, it leaves the moving objects only from one of the images. Another approach is to leave the moving objects to create ghosts - if the exposure periods of the subsequent images are contiguous, this would create a moving object motion blur, as if captured with a very long exposure period, i.e. the image will look consistent.

Approaches

To decrease the blur, exposure time can be reduced, but this leads to darker image, i.e. less SNR

To decrease the noise, either the gain could be decreased (thus less brightness, less overall SNR) or exposure could be increased while decreasing the gain (keep brightness, but increase blur).

Average all images - decreases the noise up to $\sqrt{\text{number-of-images}}$. Decreases the blur also since different images are blurred in different directions. Yet, the residual blur may be more than the one of the least blurred image.

Deblurring - there are many algos, simultaneous deblurring of multiple frames shows good results (in the articles), but is very computationally expensive. Yet, some simple approach is computationally acceptable. Note, deblurring increases the noise, especially the high frequencies. Also, deblurring is very prone to creating "ringing" artifacts near strong edges and overexposed areas.