Guidelines for Maximizing the Benefit of PixonVision Real-Time Video Processing

March 21, 2004

1. Executive Summary

- PixonVision is a real-time video-processing technology that improves three common sources of image degradation: blur, low light, and dynamic range.
- In order to take full advantage of PixonVision processing, it is important to understand the range of parameters under which each of the image degradations occurs and how it is attacked by PixonVision. Following is an executive summary for each one. This is followed by concrete examples, more technical details, and a glossary of terminology.
- **Blur:** PixonVision processing reduces the blur width by about a factor of 2, thereby increasing resolution by the same factor. Atmospheric blur usually dominates in optical applications. The angular width of atmospheric blur is typically tens of microradians, and under very good seeing conditions it can go down to ~10 microradians.[1]

  You need deblurring if the angular extent of the details that you want to resolve, as seen by the camera, is comparable to the blur width. If it is much larger, you do not need deblurring. If it is much smaller, an improvement by a factor of 2 will be insufficient.

  You also need to sample the blur width with 2–4 pixels. The angle subtended by a pixel, in microradians, is its width in microns divided by the focal length in meters. For example, for 5-micron pixels and a focal length of 1000 mm (1 m), a pixel subtends an angle of 5 microradians, perfect sampling for a blur width of 10–20 microradians.
- **Low light:** Noise suppression is required in order to discern finer gradations of intensity, or even to detect an object in the first place. PixonVision is unique in adapting the degree of noise suppression to the scene, in order not to lose resolution. It suppresses noise less for a small or narrow object, e.g., an antenna, and much more in the smooth background. The result is that narrow features stand out better above the background.

  The signal-to-noise ratio is smaller under low light conditions, either because the ambient light is low, or because the optics uses a high shutter speed or a large f-number. These are the conditions under which noise suppression by PixonVision processing is most important.
- **Dynamic range:** PixonVision uses adaptive contrast enhancement in order to overcome the human difficulty of simultaneously viewing shadow and glare or discerning detail through haze. This enhancement does not depend on camera characteristics such as pixel size, focal length, or f-number.

---

1 A microradian is 10^{-6} of a radian, i.e., the angle subtended at the camera by an object whose width is a millionth of its distance. For example, 10 microradians corresponds to 1 cm (about the width of a license-plate number) at a distance of 1 km.

2. Examples

- **Identification of people or license plate numbers:** The required resolution is about 1 cm in both cases. This corresponds approximately to the width of a license plate number on the one hand, and images a face with 20–30 resolution elements across on the other hand.

  The width of the atmospheric blur can change from day to day and from hour to hour. A typical value is around 20 microradians, and PixonVision processing can reduce it to 10 microradians. Under these conditions the license plate should be readable and the face identifiable at a distance of about 1 km. For a pixel size of 5 microns, the focal length should be 500–1000 mm (0.5–1 m), so the pixel subtends an angle of 5–10 microradians, providing adequate sampling of the atmospheric blur.

  Under worse atmospheric seeing, e.g., a blur width of 40 microradians, PixonVision processing can reduce it to 20 microradians. The license plate would then be readable and the face identifiable at a distance of about 500 m.

  The requirement for focal length is also reduced by a factor of two, but this should not lead you to use a camera with a smaller focal length. It is always better to use a camera with a large enough focal length, so you can take advantage of better atmospheric seeing. If conditions deteriorate, you can always reduce the focal length of the camera by zooming out. But if the maximum focal length of your camera is too small, and zooming in to the maximum zoom is insufficient, you are stuck.

  To test PixonVision resolution improvement, we recommend placing the target at a number of distances ranging from sufficiently near that the license plate is read easily, or a face identified with ease, say, 200 m, to far away, where either task is clearly impossible, say more than 2 km. The largest distance at which a license plate number can be read, or a face identified, will depend on the atmospheric blur, which does vary. But, if properly set, PixonVision processing should allow the detection to be made at about double the distance of the unprocessed video.

- **Identification of vehicles or small boats or buildings:** The required resolution is about 10 cm in these cases, again with ~30 resolution elements across the object. The same calculation as above applies, except that the distances are 10 times larger, because the metric resolution is 10 times larger, but the angular requirements are the same.

- **Low light:** To test the effect of PixonVision noise suppression as a function of light level, film at dusk or dawn. Alternatively, use higher shutter speeds or increase the f-number of the optics (by narrowing the aperture or by zooming in) to restrict the light reaching the focal-plane array. To test the PixonVision contrast enhancement select scenes with variable levels of illumination and/or haze.

- **Dynamic Range:** To test the PixonVision contrast enhancement select scenes with variable levels of illumination (shadow and glare) and/or haze.

3. Blur

The blur in an image is measured by the point-response function (PRF), which is the image extended on the focal-plane array (FPA) by a point source at infinite distance. The characteristic figure of merit of the PRF is its full-width at half maximum (FWHM). That is the normal limit on resolution, since the images of two point sources at infinity, separated by a FWHM of the PRF, will appear to be blended into a single diffuse object.

PixonVision processing deblurs the image so that, after processing, the FWHM of the residual PRF is about half that of the original PRF, thus improving resolution by a factor of 2. Its advantage over all other deblurring techniques is that it does so cleanly, without amplifying noise, and without introducing false artifacts.

A major source of blur (except for observations from space) is due to the atmosphere. Under very good conditions, the angular FWHM of the atmospheric PRF is ~10 microradians. Under exceptional conditions, it may go down to ~5 microradians. Typically, especially for horizontal viewing, it is several times larger, i.e., tens of microradians. In any event, atmospheric blur is an issue only for strong telescopic applications in which resolution on the order of tens of microradians is important.

Two basic conditions need to be met for PixonVision processing to be effective. (This is true of any deblurring technique.) First, the resolution requirement must match the expected blur. PixonVision processing will give you an extra factor or 2 in resolution. Make sure that this factor of 2 is what you need. Second, let PixonVision do its job by using a camera with a sufficiently long focal length. Below are precise instructions on how to compute the required parameters for your application:

- Deblurring has to be meaningful for the objects of interest. PixonVision processing can reduce the FWHM of the PRF by about a factor of two. Hence, it is most meaningful if the required resolution is 0.5–1 x the FWHM of the PRF. There is little point in deblurring if the resolution of interest is much larger than the FWHM of the PRF, say, hundreds of microradians in the case of atmospheric blur. By the same token, PixonVision deblurring will be insufficient if the required resolution is much smaller than the FWHM of the PRF, say, less than a microradian.
- The FPA must sample the PRF sufficiently. A good rule of thumb is for the FWHM of the PRF to span 2–4 pixels on the FPA. The angular extent of a pixel in microradians can be written as

  $$\theta = a / F \quad ,$$

  where *a* is the size of the pixel in microns, and *F* is the focal length of the camera in meters. The angular extent of a pixel, θ, should then be compared with the FWHM of the PRF. For example, if the pixel width is a=5 micron, and the focal length of the camera is F=1000 mm (1 m), then θ=5 microradians. For atmospheric blur with FWHM of 10 microradians, there are therefore 2 pixels on the FPA per FWHM of the PRF, perfect for PixonVision processing. By contrast, if the focal length of the camera is F=100 mm (0.1 m), then θ is 50 microradians, which may be much larger than the atmospheric blur. Under such conditions, there is inadequate sampling.

  If the FWHM of the PRF spans more than 4 pixels on the FPA, it is usually better to reduce the focal length of the camera, which is easy to do in a typical camera with optical zoom. The limit of the optics is usually encountered in the other direction, when the focal length is insufficient at maximum zoom.

4. Low Light

Our ability to see in low light is limited by noise. The noise may result from the low light level itself (statistical intensity fluctuations due to the small number of photons detected) or from electronic or other sources. Whatever the origin of the statistical noise, it causes pixel-to-pixel variations that limit our ability to discern fine gradations of intensity.

A figure of merit for the effect of noise is the ratio of signal to noise, called the signal-to-noise ratio (SNR). The SNR varies across an image because the image intensity changes, and sometimes also because the noise level varies.

Low SNR can also indirectly cause loss of resolution, even when blur is not an issue, i.e., on scales much larger than the FWHM of the PRF. Finer details are lost simply because the range of signal intensities over which they vary is comparable to the noise level or smaller and they are not noticed. A criterion for the detection of a signal difference is that it is at least 3–5 times larger than the noise.

Noise reduction is achieved by averaging signal over neighboring pixels. Since the statistical fluctuations on top of the signal can be either positive or negative, the averaging reduces the statistical fluctuations due to noise. PixonVision noise reduction is unique in maximizing noise reduction without compromising resolution by adapting the number of pixels over which the signal is averaged to the local image conditions around it.

PixonVision noise suppression is therefore of greatest benefit for images with variable imaging conditions. For example, if the scene contains antennae and cables that need to be detected, PixonVision noise suppression will be small at the location of the antennae and cables, in order not to lose resolution, and larger in the background, where larger noise suppression does not cause loss of information. In this way, the desired features (antennae and cables in this case) stand out against the background.

Noise suppression is more important, of course, under conditions of low light, when the SNR is smaller. The light level depends on three factors: (i) the ambient light, (ii) the shutter speed, and (iii) the f-number of the optics (the ratio between the diameter of the aperture and the focal length of the camera).

5. Dynamic Range

The human visual system can see only a limited range of intensities. It is particularly difficult simultaneously to see details in side-by-side dark and bright features, e.g., neighboring shadow and glare. For example, in an 8-bit digital display the intensity ranges from 0 to 255 data numbers (DN). An area of shadow, however, might use only DN=0–40, while a neighboring area in bright sunlight might use DN=215–255, so the brightness and contrast are not optimal in either area. Each area is therefore making inefficient use of the available intensity range, making it more difficult for the user to see.

Displays usually allow the user to adjust the global brightness and contrast of the image but not to adjust it differentially across the image. A shadow range of DN=0–40 can be expanded up to DN=0–255, but the sunlit area would be then be saturated at maximum intensity, and no details would be seen. Conversely, a sunlit are of DN=215–255 can be expanded down to DN=0–255, but the shadow would then be zeroed at minimum intensity, and details would again be lost.

PixonVision provides adaptive contrast enhancement to overcome this difficulty, automatically adjusting the brightness and contrast controls in each location in the image to fit the conditions there. This allows the shadow range of DN=0–40 to be expanded up, while the sunlit range of DN=215–255 is simultaneously expanded down.

Another important application is reduction of haze. The area of interest in the haze may span only a limited intermediate range of intensities, while uninteresting foreground objects are darker and the sky is brighter. The result is a narrow range of intensities for the area of interest, e.g., DN=100–140. PixonVision processing will again adjust the contrast adaptively, expanding the range of intensities in the region of haze without compromising darker and brighter areas.

6. Glossary

DN	Data number: the digital reading from a pixel in the focal-plane array. A typical optical detector outputs 8-bit data (DN=0–255).
F-NUMBER	The ratio between the diameter of the aperture and the focal length of the optics.
FPA	Focal-plane array: a chip of pixel detectors placed in the focal plane of the camera.
FWHM	Full-width at half maximum: the width of the effective point-response function, and hence a measure of resolution.
MICRORADIAN	10^{-6} of a radian: the angle subtended at the camera by an object whose width is a millionth of its distance. For example, 10 microradians corresponds to 1 cm (about the width of a license-plate number) at a distance of 1 km.
PRF	Point-response function: the image extended on the focal-plane array by a point source at infinite distance.
SNR	Signal-to-noise ratio: the ratio between signal (image intensity) and noise. It varies across an image because the image intensity changes, and sometimes also because the noise level varies.