

Altavian

Does Sharper Drone Imagery Mean Better Data?

Sharpness is arguably the most important factor in perceived image quality because it is one of the first things we see when we examine imagery. In their application to photogrammetric data, sharp images with high neighboring pixel contrasts are crucial to the detection and matching of features, yielding high informational content for the creation of 3D point clouds using dense matching methods. At the same time, having drone imagery that looks sharp and crisp can, on the surface, help sell your data product.

At Altavian, our drones' imagery sometimes appears 'softer', or less-sharp, than that collected from other drones. The reasons behind this are hard to explain without going into extensive detail. After all, comparing images with our own eyes is relatively simple; our eyes can focus easily on high-contrast objects and therefore, we see general sharpness as a beneficial trait. However, the way computer vision algorithms see images is very different than the way our eyes see them.

To shed some light on the topic and attempt to answer the question "why is your imagery soft?", we conducted a few experiments to show how sharpness affects images when processed into photogrammetric data products. What we found was that while artificial sharpening may make drone imagery appear crisp and clear, it will have hidden influence on the accuracy of your data.

UNDERSTANDING SHARPNESS IN DRONE IMAGERY

What we perceive as 'sharpness' is a combination of both resolution and acutance, or the subjective perception of sharpness that is related to the edge contrast of an image. Under the right conditions, high-quality, high-resolution camera systems can produce truly sharp imagery without any manipulation of the information that the sensor collects. On the other hand, lower-quality, lower-resolution cameras must make up for their softness by artificially boosting image edge contrast (sharpness) in their internal processing pipelines. This artificial sharpening is unfortunately what

some people have come to expect from digital imagery because it passes the ‘eye-test’ of looking sharp.

For any scientific data collection process, however, it is imperative that the information collected remains as faithful as possible to reality and that all perception biases are minimized. For example, cameras intended for general photography capture images whose raw information is first run through a number of processing steps to convert that information into an appealing picture. This processing generally includes image compression, artificial sharpening, tone curve mapping, exposure compensation, and other adjustments designed to make the image more aesthetically pleasing.

While this automated manipulation can be beneficial to amateur photographers, most professional photographers will take advantage of their camera’s RAW mode, in which more information is maintained in the image and any in-camera manipulation is minimized. The reason for this is so the photographer has more control over the processing in their photo editing software to suit their own visual style. This is the same methodology that applies to photogrammetric data; no matter what method is employed, processes which alter the raw data in any way nearly always involve either losing resolution (as in noise reduction) or increasing noise (as in sharpening).

The inherent problem with digital sharpening is that any “boost” given to an image is purely perceptual and cannot be said to truly recover information of the scene. Classic sharpening algorithms such as Unsharp Masking only boost acutance and have nearly nothing to do with resolution improvement. Therefore, image enhancements are a really bad idea for computational photography techniques, especially if the ultimate goal is high precision, reproducibility, and low uncertainty. In short, this is why we never artificially sharpen our drone imagery beyond what we cannot control within a consumer camera, yielding what we believe is drone imagery with the fewest processing artifacts and a more faithful representation of reality.

EXPERIMENT 1

To test our theory, we conducted two identical flights using an Altavian Nova F7200 aircraft with an MP22 sensor payload; the first flight using our standard, faithful camera setting, and the second flight adding +3 sharpening to the camera’s internal image profile. We wanted to see how the processed data from the control imagery compared to the imagery with the camera’s internal sharpening applied, as well as imagery sharpened using an Unsharp Masking algorithm.

The two flights were flown within the same hour, operated at 400 ft with a sidelap and overlap of 85%, both flights yielding datasets of 170 images. Using the control dataset (with default settings), a third and fourth dataset were created using RawTherapee to edit the sharpness of the photos with Unsharp Masking – one set with sharpness enabled at the program’s default setting and another with the sharpness amount increased to the maximum extent. This left us with four datasets to process:

1. Control: no image manipulation
2. In-Camera: +3 sharpening added in the camera's internal profile
3. Default: Using RawTherapee's default Unsharp Mask setting
4. Max: Using RawTherapee's maximum Unsharp Mask setting

We then took these datasets into Pix4DMapper Pro and processed them all using the same default '3D Maps' processing template with Geometrically Verified Matching enabled. According to the Pix4D processing reports, the results were as follows.

The results show that the control dataset outperformed both the default Unsharp Mask dataset and the in-camera sharpened dataset in the number of keypoints found, the number of keypoints matched, as well as total keypoint observations for the bundle adjustment. The reason for this is that, although the control images appear to be softer to the human eye, the more faithful features preserved in the untouched drone imagery allow the software to determine the relative positions of objects using a wider variety of contrast, outweighing the aesthetic enhancements yielded from the sharpening.

However, the software reports suggest that the RMSE was better with the default Unsharp Masking imagery, and the maximum sharpened images actually outperformed all other datasets in every category. As the sharpness increased, so did the number of features found in the images, and therefore, the number of features matched between images.

Ignoring the in-camera sharpened data due to poor results in each category, we compared just the orthoimagery produced from each Unsharp Masked dataset, and the sharpened drone imagery also showed a noticeable enhancement in aesthetic quality. This was especially true when reviewing areas affected by common 'halo' artifacts found when point cloud densification is too low to properly model a feature. Control (left), default Unsharp Mask (center), maximum Unsharp Mask (right). Halo artifacts around center tree-line decrease as sharpness increases.

This enhancement in halo artifacts made more sense when we realized that, regardless of the number of keypoints matched, Pix4D consistently densified the point clouds to a greater extent relative to the sharpness of the images.

Control = 30,871,704 (116.58 ppm3)
Default = 30,951,714 (117.65 ppm3)
Max = 31,195,148 (118.07 ppm3)

This increased point cloud density in turn resulted in more detailed elevation models, further yielding proportionally larger orthomosaic files.

Control = 571 MB
Default = 608 MB
Max = 693 MB

APPEARANCES CAN BE DECEIVING

We know that photogrammetric data processing programs are designed to minimize mean pixel reprojection error, which accords with the least-squares bundle adjustment under the collinearity constraints. What we have found here is that the overly-defined features brought forth by artificial image sharpening leads to a higher degree of confidence in the software's bundle adjustment and therefore, a lower overall error. The problem, however, is that this confidence is entirely relative to the input data.

Well-matched keypoints at sufficiently large disparities are what lead to good parameter recovery in the bundle adjustment. Dense matching, on the other hand, relies on there being relatively small changes in features across fewer images. This makes depth recovery highly sensitive to the small inaccuracies in point estimates caused by the artificial sharpness affecting the same features slightly differently in each image.

It's important to understand that these low-level vision algorithms don't use color data independently, but instead fuse them into a single grayscale luminance channel, allowing them to most readily respond to contrast in the image. When artificial sharpening is introduced to the images, not only is contrast added to real, well-defined features, but also to image noise, and as noise becomes amplified, the software responds to it as an inseparable part of the true features. This means there's higher variability in the matches.

Due to the high overlap in data collection, the aforementioned weakness in the geometry needed for robust aerotriangulation increases the sensitivity to image artifacts. Counteracting this is the fact that many features are at least consistently matched within the nature of the overall model constraints. Gross outliers can be rejected, and biases in the data can be minimized, but only so far as the photogrammetric model's assumptions and parameterization can allow. Unmodeled errors can creep in without necessarily increasing RMSE when it is not computed using independent ground truth. Hence, lower RMSE values for highly sharpened images are explainable without necessarily meaning the accuracy of the product has been improved to the degree the numbers would suggest.

To confirm this understanding, we processed another dataset which included more well-known structures using the same Unsharp Masking methods. Although the software reported that overly-sharp drone imagery produced better results, we wanted to see exactly what effects this sharpening had on the ultimate model accuracy.

EXPERIMENT II

Processing the new drone imagery using the same Unsharp Masking method in RawTherapee, we again produced three datasets; a control (with no manipulation), a default (with RawTherapee's default Unsharp Mask setting), and a Max (with RawTherapee's Unsharp Mask amount set to the maximum extent). Again, Pix4D reported better results from the overly-sharpened drone imagery.

Drawing cross-sections of hard structures in the point clouds, we found some interesting results. Most notably, as the sharpness was increased, the model position consistently shifted and the elevation consistently dropped, by as much as 0.23m between the Control and Max datasets.

Cross sections show a consistent drop in elevation as sharpness increased.

Additionally, we saw that point cloud noise was further dispersed as the sharpness was increased, turning a noise deviation of 0.21m in the control data to 0.83m in the Max.

Feature noise is further dispersed as sharpness increased.

Vertical feature bleed, a common photogrammetric artifact in which the software fails to detect the sudden difference between the edge of a tall surface and the beginning of a low surface, was also negatively impacted from a dispersity of 0.6m in the control data to 1.45m in max.

To determine the absolute extent of these anomalies and more clearly define the effects of sharpness, we introduced ground control into the solution and compared the deviation across all three datasets.

The results show that, as more image sharpness is artificially added, the software reports increasingly better bundle adjustment precision as the absolute accuracy equally decreases. This is again caused by the software's optimization goal of minimizing pixel mean reprojection error. While the bundle adjustment accuracy technically is better as it uses the widest-separated points, the aerotriangulation, which uses narrower shifts, is being degraded, minimizing the data's absolute accuracy.

CONCLUSION

While sharpness is an overall appealing trait of drone imagery, it is important to understand how that sharpness is achieved and how it impacts your bottom line. Images collected with small point-and-shoot cameras with smaller sensors are often over-processed to compensate for their noise and the limitations of the optics to which they are mated. The artificial sharpness imposed by their internal image profiles provides a false appearance of clarity, which may ultimately have significant impacts on your data's file sizes and overall accuracy.

So, if you're reviewing drone imagery that seems soft, it may not be blur, but rather a sign of a clear end-to-end process, without artificial adjustments, that is, as a result, a more accurate representation of reality.