Photo Alignment and Sparse Cloud Optimization in Agisoft Metashape

Objective

The main goals of this lab exercise are to familiarize students with:

- Options for image alignment and keypoint generation in Metashape
- Tools to evaluate keypoints and the results of image matching
- Strategies for keypoint optimization to improve the SfM stereomodel
- How to interpret image alignment and stereo model results in the Metashape quality reports

Deliverables

Fill out and submit the questions at the end of this document.

Note: Please refer to the class BBLearn site for lab due dates. You may work together and help each other, but please make sure what you turn in is your own work.

Section 1: Alignment of Parker Farm Photos

The first thing we need to do is load the photos into Metashape and align them.

1. Download the Parker Farm RGB image dataset from the Drone Lab website (https://uidronelab.org/2024/08/22/photo-alignment-and-optimization-in-metashape/) or class Canvas page.

2. Open Metashape with a new/blank project. From the menu, choose **Workflow -> Add Folder** to add all of the images from the Parker Farm drone flight. Metashape will ask what kind of photos these are. Choose **Single Camera**. Once the photos are loaded into Metashape, you should see them in the photo tray at the bottom of the window.

3. From the main menu, choose **Tools -> Camera Calibration**. This dialog box contains information about the camera(s) that collected the photos from the image headers. If you have a camera where you know the calibration, you can enter that information here. Otherwise, Metashape will estimate the camera properties during the image alignment step. This is also where you enable the rolling shutter compensation if needed. The DJI Mavic 2 Pro drone has a rolling shutter, so **check Enable Rolling Shutter Compensation, and then click OK.**

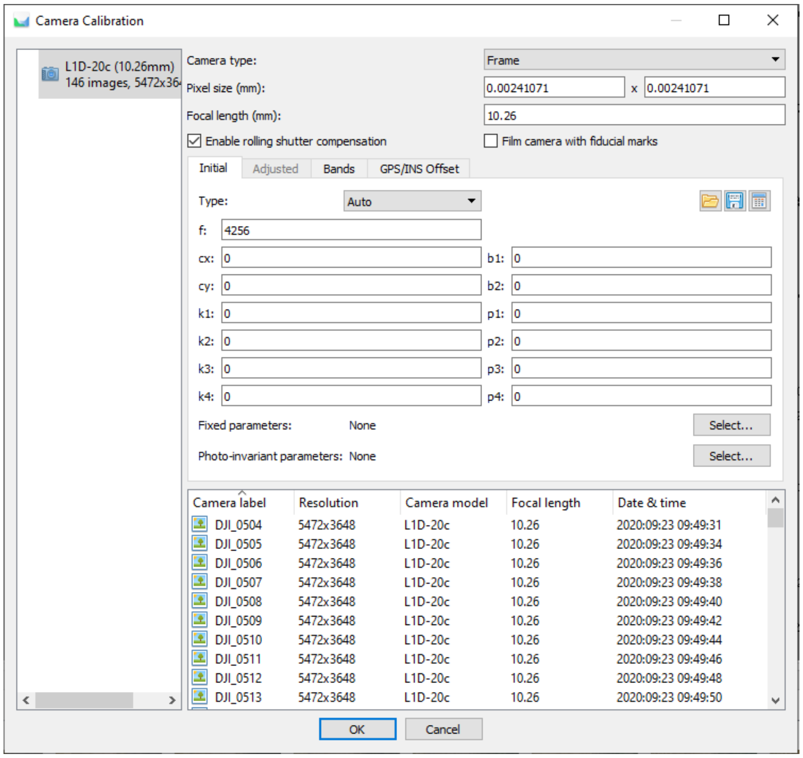

4. The next step is to align the photos. From the menu, choose **Workflow -> Align Photos**. This will open the Align Photos dialog box. The settings here control how the tie points are generated from the original images, how many key points per image are generated, and how many matches can be found between images. **See the figure below for the recommended settings at this stage and then click OK.** Some additional info on the options:

   Accuracy: Controls whether or not images are first resampled before key points are created per image. High (can take a really long time) uses the images at original resolution, and each step down resamples the previous step by a factor of four. For example, Medium resamples the images to a pixel size four times larger than the original photos. Generally, medium is adequate for most applications, and high does not produce results so much better that it justifies the extra processing time. Obviously, though, the more you resample the images, the poorer the quality of the resulting model will be.

   Generic preselection: Runs an initial photo alignment at low accuracy first to speed up key point generation at higher accuracy levels.

   Reference preselection: uses estimated camera locations from initial GPS coordinates or other reference information (e.g., order of photos) to help figure out where photos are in relation to each other and which photos are likely to have matching key points. ***This option can greatly speed up photo alignment if you have some sort of reference information for your photos.***

   Key point limit: The maximum number of key points that can be generated for any given photo. A value of zero will disable the maximum key point value per photo. The default value of 40,000 is generally sufficient. Increasing this number by a lot can generate many low-quality key points.

   Tie point limit: The maximum number of tie points between any two images. Tie points are filtered/kept according to their quality. A value of zero turns off the filtering and will return the maximum number of tie points between images, which can result in many low-quality tie points.

Adaptive camera model fitting: recalculates the camera interior orientation parameters using the image model. Generally, turn this option on unless you have good, known camera calibration information.

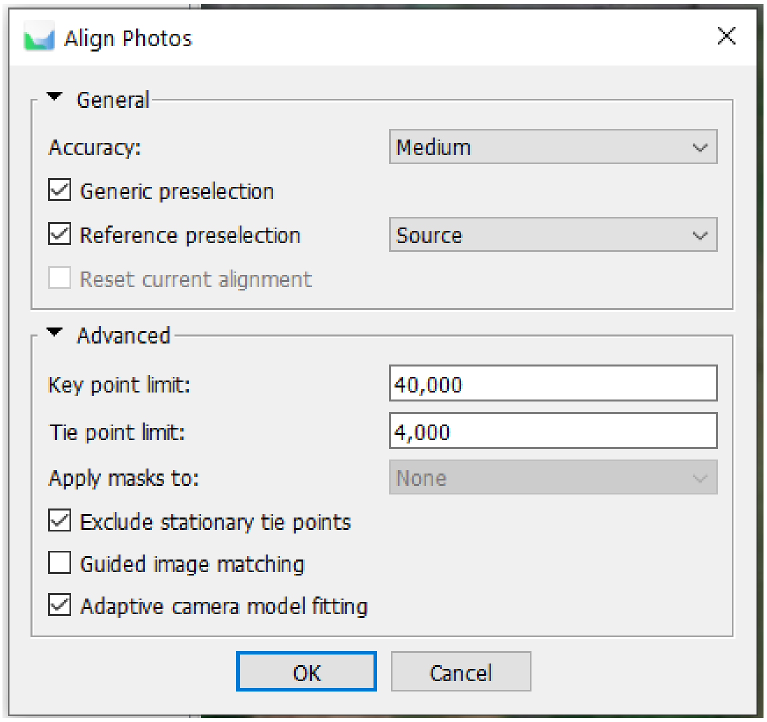

Note: The photo alignment process will take some time to complete. On my laptop, it took approximately 13 minutes to run.

5. If you imported the whole folder of images (and didn't look and remove the one weird photo), then you'll get the following warning once the photo alignment is complete.

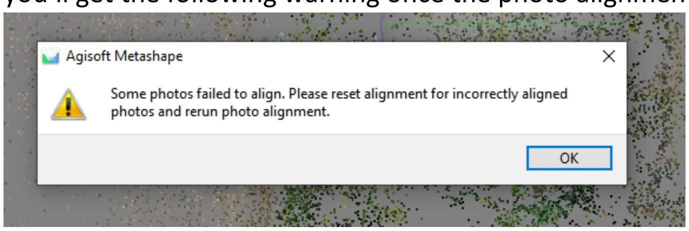

This is because one of the photos is an oblique photo of our group taken manually from the drone that didn't get removed before we brought the images into Metashape. You can see this photo as the blue dot that did not get transformed into an oriented photo representation in the model window.

If you have valid photos that do not align, it is usually a sign that something is wrong with the image quality or amount of overlap. You can use control points (i.e., markers) to create manual tie points in Metashape between images that won't align and overlapping images that do align and then redo the alignment procedure for individual images.

In this case, though, we'll just remove this photo by first figuring out which one it is (zooming in shows the photo number), right-clicking on that photo in the photo tray, and then choosing "Remove Cameras" from the context menu.

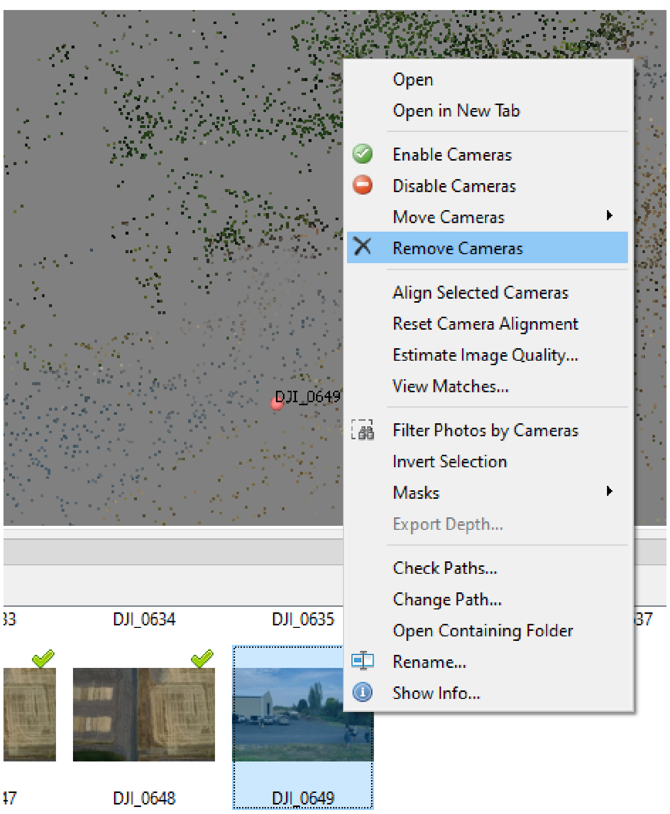

Section 2: Evaluating results of Image Alignment

Now that we have an initial 3D model, we can take a look at how well Metashape did at aligning our images. This image set is pretty good in terms of its geometric properties (e.g., overlap and number of overlapping images across the area) and the image quality, so there isn't much to really investigate for photo alignment issues. Still, it's good to be familiar with some of the basic tools.

1. First, notice in the lower left corner of the model window that it tells you how many total tie points there are across all the photos. This will be important as we start filtering the tie points to improve quality of the model.

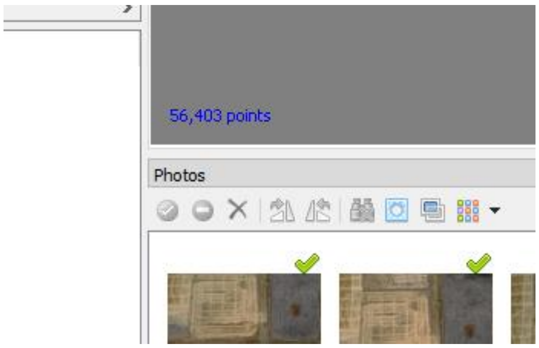

2. Right click on any photo in the photo tray and click "View Matches" from the context window to see how many tie point matches Metashape found between that photo and every other photo. This will show you "valid" tie points (i.e., ones that Metashape thought were good and kept) as well as the "invalid" ones it dropped. The drop-down box at the top selects a photo and the list shows all other photos that have tie points to it. Clicking on a photo in the list will show the matches on each photo. **Choose a photo and look at matches to another photo with a lot of tie points to it and a photo that does not have many tie points to it.** I have found this tool useful when I get situations where I have photos that align into two main blocks but the blocks themselves don't align well. This tool can help pinpoint images that may be messing things up in the overall alignment.

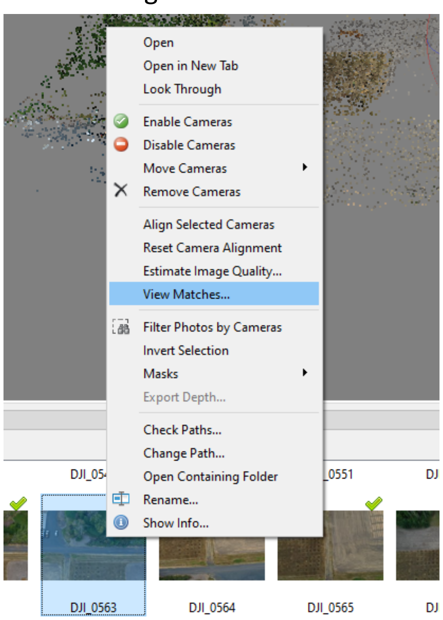

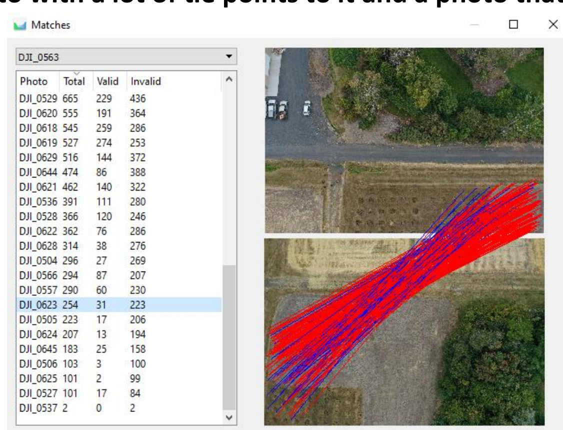

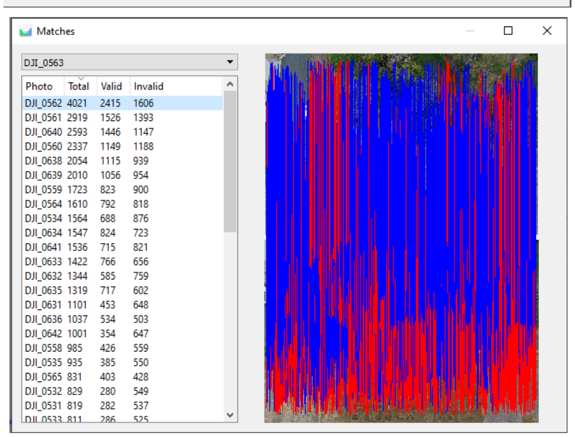

3. Creating a Processing Report from Metashape is also a good way to see how well the image alignment process worked. From the main menu, **click File, Export, Generate Report. Give this report a name and save it to a folder where you will save this project.** This will be the first of two reports we will create, so you will want to label it as such. Some useful pieces of information from the report are:

   Survey Data: Figure 1 shows the number of images that overlap for each of the tie points. This is generally useful information to make sure your flight paths and image collection were adequate. When a DEM is created for the project, the report will show a raster background here and not just the sparse point cloud.

   The flying altitude, ground resolution, and coverage area are estimated from the 3D model. These will change as we refine the model and as ground control information is added.

   Reprojection error is the RMSE of the positional error for all of the tie points. This is an initial indicator of how **precise** the model is, and this will also change as we improve the model.

   Camera Calibration: Figure 2 shows the results of the camera calibration model that Metashape computed – i.e., the average distortion produced by the camera lens (and that Metashape has accounted for). Table 2 gives the estimated camera interior orientation parameters and their error estimates.

   Camera Locations: Figure 3 shows how far (in X, Y, and Z directions) Metashape had to move the photos from their original estimated locations to get them to all align. This is a useful figure. Photos that were moved a large amount relative to the others may not have aligned properly or may be poor-quality images (and that you may want to remove). Table 3 gives the average positional error that Metashape has estimated for the photos based on the computed 3D model. Note that these values represent **precision**, and we can get better values to represent **accuracy** once we have ground control information.

   At this point, note the reprojection error and average positional error values, and make sure you save this report for comparison to later.

Section 3: Keypoint Optimization

Metashape filters out low quality key points and tie points during the image alignment process. However, if you look at the sparse point cloud from the side, you will usually see some spurious tie points.

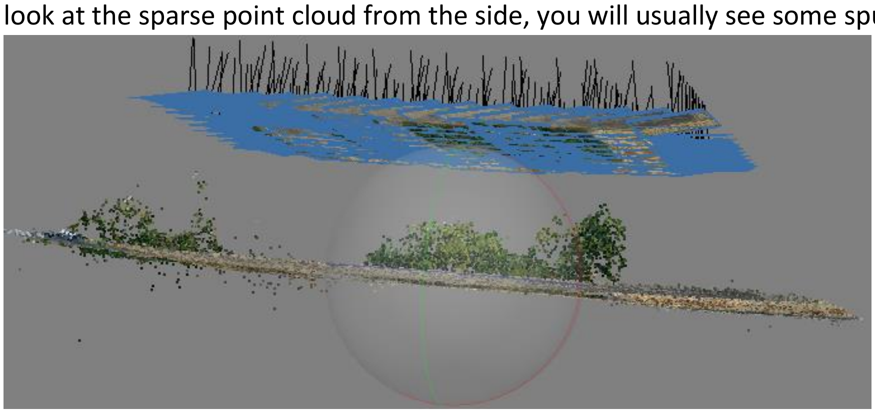

We can improve the model by identifying and removing low-quality tie points after the model has been created. To do this, we need to have information on the overall positional error of each image.

1. Note on the left side of the Metashape window the Workspace pane. At the bottom of this pane, click on the Reference tab. This opens the reference window which is where we can get information on positional error and also enter ground control information.

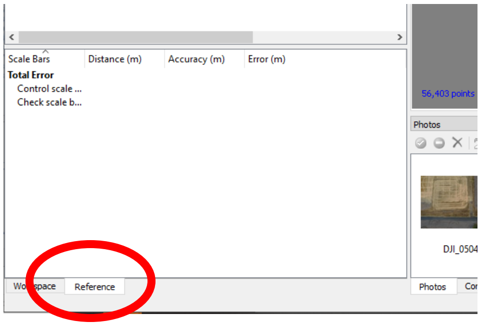

2. Click the "**View Errors**" button (see figure below) in the reference window to see the positional error for each image. Scroll to the bottom of the image list to see the total/average error in the X, Y, and Z direction. These values represent how far each image is off from the overall 3D model that was constructed during photo alignment. We will keep an eye on these values as we try to refine the model.

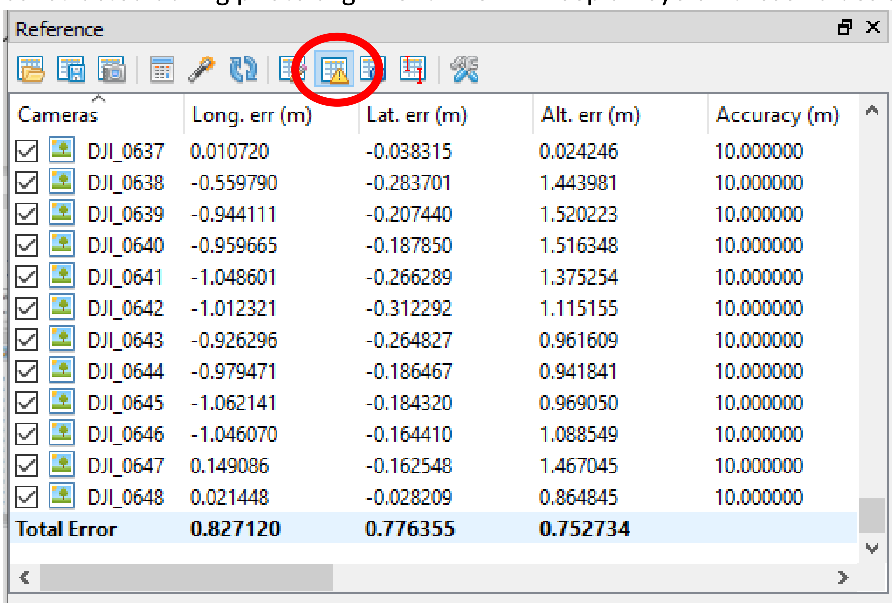

Reference

Cameras	Long. err (m)	Lat. err (m)	Alt. err (m)	Accuracy (m)
DJI_0637	0.010720	-0.038315	0.024246	10.000000
DJI_0638	-0.559790	-0.283701	1.443981	10.000000
DJI_0639	-0.944111	-0.207440	1.520223	10.000000
DJI_0640	-0.959665	-0.187850	1.516348	10.000000
DJI_0641	-1.048601	-0.266289	1.375254	10.000000
DJI_0642	-1.012321	-0.312292	1.115155	10.000000
DJI_0643	-0.926296	-0.264827	0.961609	10.000000
DJI_0644	-0.979471	-0.186467	0.941841	10.000000
DJI_0645	-1.062141	-0.184320	0.969050	10.000000
DJI_0646	-1.046070	-0.164410	1.088549	10.000000
DJI_0647	0.149086	-0.162548	1.467045	10.000000
DJI_0648	0.021448	-0.028209	0.864845	10.000000
Total Error	**0.827120**	**0.776355**	**0.752734**	

3. In the next step we'll start trying to refine the model. **From the menu, select Model, Gradual Selection.** This opens a dialog box where you can choose different criteria for evaluating the tie points.

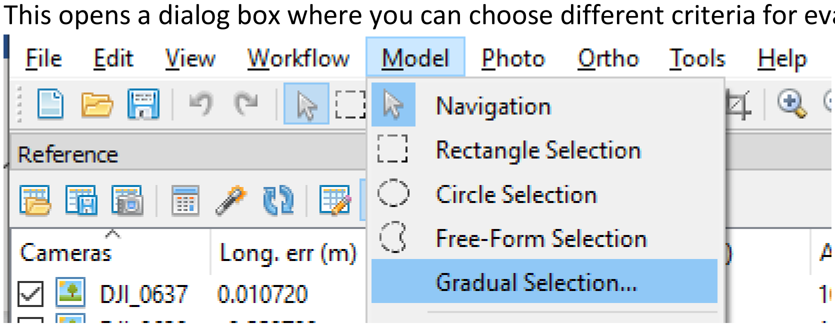

4. **Choose Reconstruction Uncertainty from the Criterion drop-down.** From the Metashape user's manual, "High reconstruction uncertainty is typical for points reconstructed from nearby photos with small baseline. Such points can noticeably deviate from the object surface, introducing noise in the point

cloud. While removal of such points should not affect the accuracy of optimization, it may be useful to remove them before building geometry in Point Cloud mode or for better visual appearance of the point cloud."

We will remove points with high reconstruction uncertainty in two steps. First, **use the slider or set the level to around 50.** You will notice that points in the sparse cloud turn pink as they are selected.

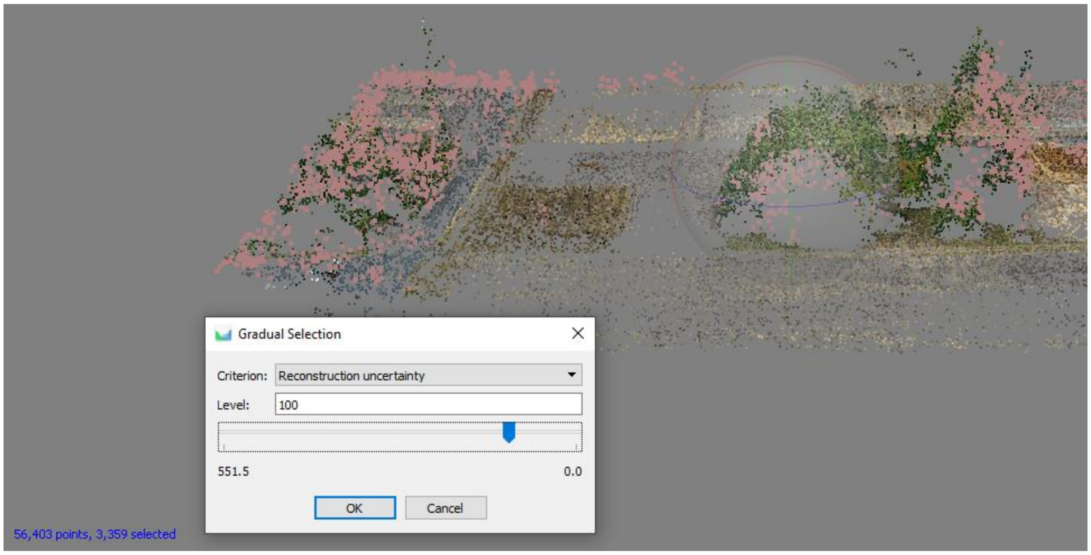

Once the points are selected. Click the "delete" button from the toolbar at the top of the Metashape window to remove these points.

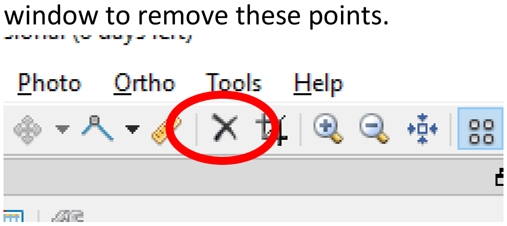

5. After you've deleted the points, we need to reoptimize the 3D model**. Click the Optimize Cameras magic wand button from the Reference window's toolbar.**

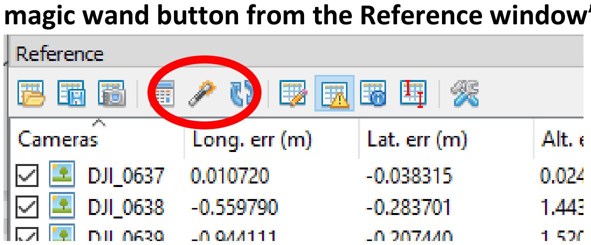

In the optimize camera alignment dialog box, **choose Adaptive camera model fitting and estimate tie point covariance, then click OK.**

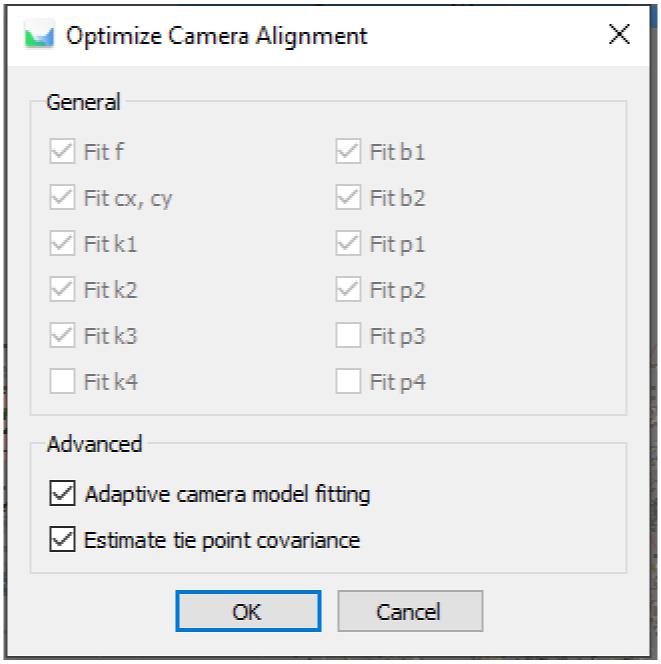

Note how the total error estimates in X, Y, and Z have decreased (probably very slightly, this is a good image set!).

6. **Repeat step 5 with selecting and removing tie points with a Reconstruction Uncertainty level more than 25. Don't forget to reoptimize after you delete them.** Notice how most of the spurious tie points are now gone from the model.

7. The next optimization step is to filter for Projection Accuracy. From the Metashape manual, "This criterion allows to filter out points which projections were relatively poorer localized due to their bigger size." This basically means tie points that are far out of alignment with the 3D model that has been computed.

   We will do this step in two parts as well. First, remove the tie points with a projection accuracy level greater than 20. Then remove points with a projection accuracy level greater than 10. Don't forget to optimize between each step. Notice if your total error values are decreasing.

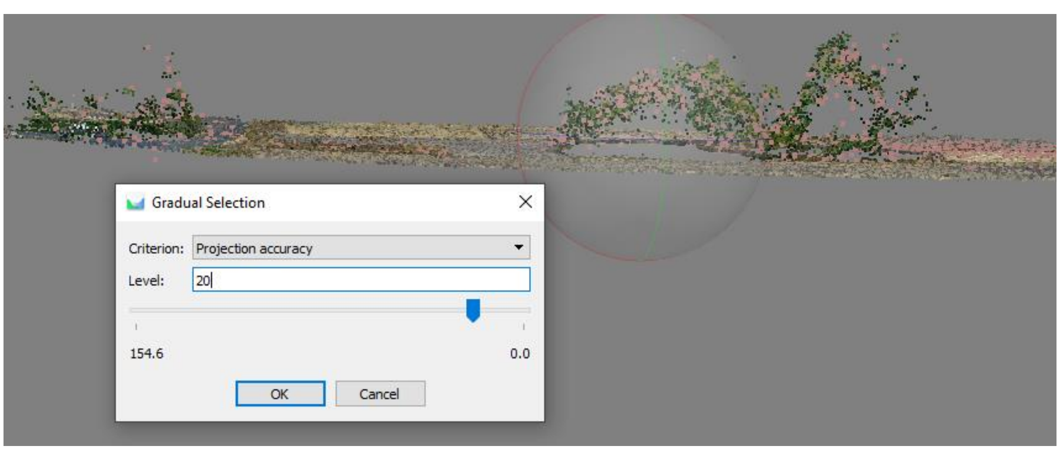

8. In the final step, we'll remove points with high reprojection error. From the manual, "High reprojection error usually indicates poor localization accuracy of the corresponding point projections at the point matching step. It is also typical for false matches. Removing such points can improve accuracy of the subsequent optimization step."

   The goal for this step is to remove the 10% of the tie points from the sparse cloud with the worst reprojection error. To find these, **select reprojection error and then move the slider until you have about 10% of the points selected (watch the point count in the lower-left corner of the model window). Delete these points and reoptimize.**

   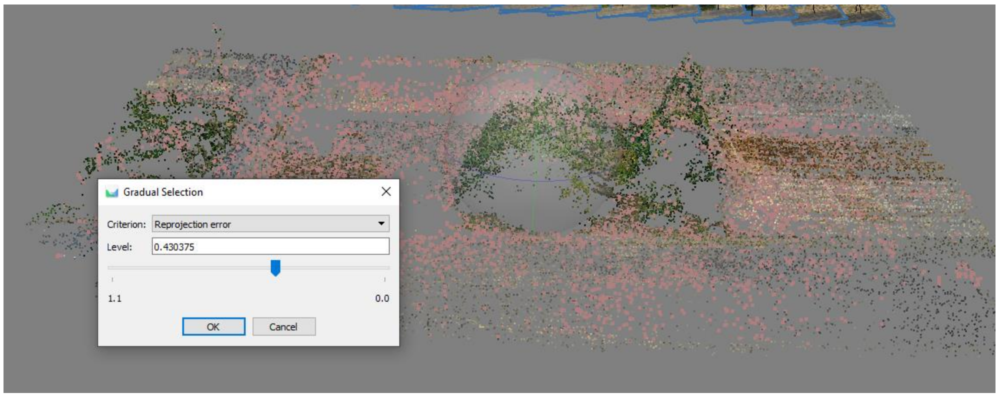

9. **Create a second report after all the optimization steps.**

10. **Make sure to save your Metashape project! We'll use it again next week.**

The result of all these steps should be a model that's dialed in and ready to generate some photogrammetric products from. In case you're wondering, the order of these refinement steps is the product of a lot of trial and error. The exact values for each step will vary by project, but when we add ground control information in the next lab, these values will have more meaning.

When you're finished, fill out and submit the question sheet for this lab.