USGS
science for a changing world

Processing Coastal Imagery With Agisoft Metashape Professional Edition, Version 1.6—Structure From Motion Workflow Documentation

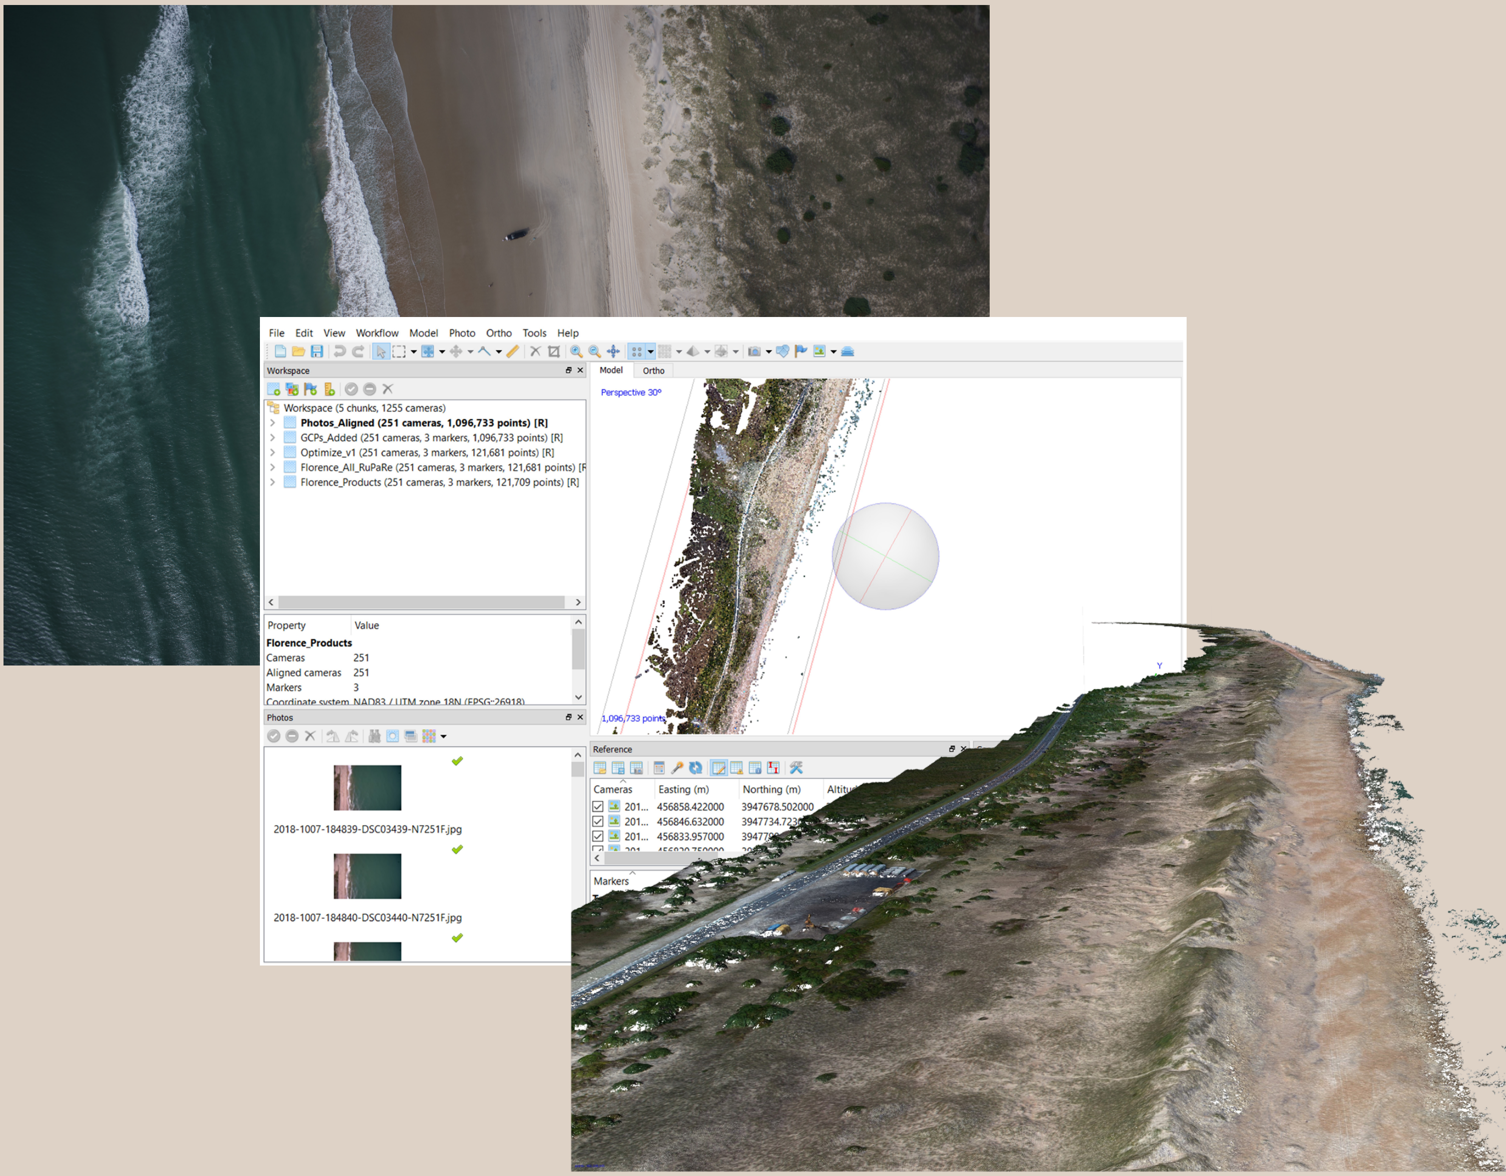

Open-File Report 2021–1039

U.S. Department of the Interior
U.S. Geological Survey

Cover. Collage illustrating the process of structure from motion on coastline imagery. Coastline imagery (top left) is processed by Agisoft Metashape software (middle), resulting in a three-dimensional model (bottom right). Images by the U.S. Geological Survey.

Processing Coastal Imagery With Agisoft Metashape Professional Edition, Version 1.6—Structure From Motion Workflow Documentation

By Jin-Si R. Over, Andrew C. Ritchie, Christine J. Kranenburg, Jenna A. Brown, Daniel Buscombe, Tom Noble, Christopher R. Sherwood, Jonathan A. Warrick, and Phillipe A. Wernette

Open-File Report 2021–1039

U.S. Department of the Interior
U.S. Geological Survey

U.S. Geological Survey, Reston, Virginia: 2021

For more information on the USGS—the Federal source for science about the Earth, its natural and living resources, natural hazards, and the environment—visit https://www.usgs.gov or call 1–888–ASK–USGS.

For an overview of USGS information products, including maps, imagery, and publications, visit https://store.usgs.gov/.

Suggested citation:
Over, J.R., Ritchie, A.C., Kranenburg, C.J., Brown, J.A., Buscombe, D., Noble, T., Sherwood, C.R., Warrick, J.A., and Wernette, P.A., 2021, Processing coastal imagery with Agisoft Metashape Professional Edition, version 1.6—Structure from motion workflow documentation: U.S. Geological Survey Open-File Report 2021–1039, 46 p., https://doi.org/10.3133/ofr20211039.

ISSN 2331-1258 (online)

Acknowledgments

This document would not have been possible without the expertise and dedication of Neffra Matthews of the Bureau of Land Management. This work was supported by the U.S. Geological Survey response and recovery efforts under congressional appropriations through the Additional Supplemental Appropriations for Disaster Relief Act, 2019 (Public Law 116–20; 133 Stat. 871).

Contents

Figures

Table

Conversion Factors

International System of Units to U.S. customary units

Multiply	By	To obtain
centimeter (cm)	0.3937	inch (in.)
meter (m)	3.281	foot (ft)

Datum

Vertical coordinate information is referenced to the North American Vertical Datum of 1988 (NAVD 88).

Horizontal coordinate information is referenced to the North American Datum of 1983 (NAD 83).

Elevation, as used in this manual, is a height referenced to a vertical datum.

Abbreviations

3D	three-dimensional
4D	four-dimensional
API	application program interface
CCG	camera calibration group
CPU	central processing unit
DEM	digital elevation model
DSM	digital surface model
DTM	digital terrain model
EXIF	exchangeable image file
GCP	ground control point
GNSS	global navigation satellite system
GPS	global positioning system
GPU	graphics processing unit
GUI	graphical user interface
INS	inertial navigation system
PCIe	peripheral component interconnect express
PPK	postprocessing kinematic
RGB	red, green, blue
RMS	root mean square
SEUW	standard error of unit weight
SFM	structure from motion
UAS	unoccupied aerial system
USGS	U.S. Geological Survey
UTM	Universal Transverse Mercator
VRAM	video random access memory

Processing Coastal Imagery With Agisoft Metashape Professional Edition, Version 1.6—Structure From Motion Workflow Documentation

By Jin-Si R. Over,[1] Andrew C. Ritchie,[1] Christine J. Kranenburg,[1] Jenna A. Brown,[1] Daniel Buscombe,[1] Tom Noble,[2] Christopher R. Sherwood,[1] Jonathan A. Warrick,[1] and Phillipe A. Wernette[1]

Introduction

Structure from motion (SFM) has become an integral technique in coastal change assessment; the U.S. Geological Survey (USGS) used Agisoft Metashape Professional Edition photogrammetry software to develop a workflow that processes coastline aerial imagery collected in response to storms since Hurricane Florence in 2018. This report details step-by-step instructions to create three-dimensional (3D) spatial products from both singular and repeated collections of shoreline aerial imagery. The products can be used for real-time hazard guidance and future forecasting and recovery endeavors.

This work was supported by the Remote Sensing Coastal Change project of the USGS Coastal/Marine Hazards and Resources Program and by congressional appropriations through the Additional Supplemental Appropriations for Disaster Relief Act, 2019 (Public Law 116–20; 133 Stat. 871). The funds from the Disaster Relief Act have been used by this project and program to assist with recovery efforts from Hurricane Florence and Michael following damage to the coastlines. The scope of the Remote Sensing Coastal Change project has been extended from documentation of Hurricane Florence (2018) impacts in North Carolina to documentation of storm impacts and recovery from Hurricane Dorian (2019) and Isaias (2020). The tasks covered by the storm documentation include producing data for hurricane impact models and assessments of future coastal vulnerability.

The workflow described in this report takes advantage of the flexibility, quality retention, and low-cost nature of SFM, as demonstrated by earlier users (James and Robson, 2012; Westoby and others, 2012; Fonstad and others, 2013); the benefits and caveats of using SFM in the earth sciences are summarized by James and others (2019). The methodological contributions set out in this report are fully transparent, based on sound photogrammetric principles, and the result of rigorous testing. These best practices have been developed for singular and repeat collections of shoreline imagery that use a combination of camera locations and ground control points to produce digital elevation models and orthorectified imagery products. This documentation details step by step instructions for using Metashape Professional Edition software, but is ultimately intended to provide sufficient background on photogrammetric principles and the workflow as used in the USGS hurricane relief and recovery efforts to allow consistent replication and application of these techniques to novel datasets.

This report is not intended to replicate the Metashape user manual (Agisoft LLC, 2020b). The steps, values, and suggestions presented in this report are not necessarily generic and may not be suitable for every SFM project. Where possible, choices are explained where they are based on the software or the data type analyzed. The workflow detailed in this report addresses processing nadir and near-nadir aerial imagery collected from approximately parallel quasilinear flight paths along the coast to measure geomorphological change over long and narrow continuous natural features, such as barrier islands. The software to produce SFM products is undergoing continual evolution, but this snapshot documentation can serve as a go-to guide when there is a need to process large landscape imagery datasets with little or no formal training in SFM or photogrammetric principles. Although this workflow is specific to Metashape Professional Edition, those working with similar software can still find relevant guidance in this document. Photogrammetric theory has not changed substantially over time, and SFM practices with aerial imagery taken of natural environments typically proceed along similar steps (Shervais, 2015). Because the use of SFM is changing, the best practices in this resource may require updates.

The photogrammetry field is diverse, with the fields of archeology and geomorphology especially having benefited from SFM workflows where it augments (Nouwakpo and others, 2016) and potentially replaces more expensive topographic survey techniques (for example, aerial light detection and ranging [lidar] and terrestrial laser scanning). When supported by a Global Positioning System (GPS) of less than 10-centimeter (cm) accuracies and pixel sizes, SFM can measure landscapes at similar qualities over multikilometer scales and detect landscape change with decimeter

[1]U.S. Geological Survey.

[2]TN Photogrammetry, LLC.

accuracy. SFM workflows can also be run rapidly, enabling near-real-time delivery of products that are valuable to a wide audience.

With this in mind, SFM workflows have been developed with the use of unoccupied aerial systems (UASs) to map coastal cliffs and dunes, river banks and floodplains, grasslands, and river ice (Woodget and others, 2015; Warrick and others, 2017; Alfredsen and others, 2018; Forsmoo and others, 2018; Carrivick and Smith, 2019); underwater SFM applications employ specialized cameras to map submerged settings, such as coral reefs (Storlazzi and others, 2016; Hatcher and others, 2020). Paleontology, archaeology, and architectural historical preservation are also benefitting from the growing technical applications of SFM with use of UASs and other camera platforms to map dinosaur trackways, ruins, and historical sites (Matthews and others, 2006; Yilmaz and others, 2007; Jones and Church, 2020). The workflow detailed in this report specifically targets mapping coastal systems from airborne flights (Sherwood and others, 2018), but could be applied in other scenarios and landscapes with similar image quality and collection methods.

Overview

Much of this workflow is software specific. Although there are open-source SFM programs, the standalone photogrammetric processing commercial software Agisoft Metashape Professional Edition facilitates grouping images by time and space, encompasses the entirety of the SFM workflow described in this report from raw imagery to digital elevation models (DEMs) and orthomosaics (Turner and others, 2014), allows the user to maintain control over model parameterization, and is scalable and scriptable using Python or JavaScript. The Metashape Professional Edition version 1.6 user manual (Agisoft LLC, 2020b) instructs on the basics of lens coefficients, image captures, and provides a generic SFM workflow; for the purposes of this report, the manual will be referred to as the “Metashape user manual”. Many of the processing steps are applicable to earlier versions of Agisoft Metashape, Agisoft Metashape Professional, and its predecessor Agisoft Photoscan (U.S. Geological Survey, 2017), but earlier versions may not include all the features discussed in this report. This guide specifically describes the requirements, procedures, and inputs necessary to create high-quality SFM coastal products from aerial imagery, including georeferenced orthomosaics and DEMs, using Metashape Professional Edition version 1.6.5 (Agisoft LLC, 2020a).

Quick Start Guide

A quick-reference, condensed version of the workflow detailed in this report is detailed in table 1. The table is not a replacement for the expanded instructions in this report; rather, it serves as an access point for the detailed instructions and complements the visual workflow (fig. 1; see also “Starting a New Project” section of this report). The most subjective part of the workflow is when to stop removing tie points and stop optimizing. There are multiple indicators (see “Error Reduction—Optimization and Camera Calibration” section of this report) that may help with this decision on a case-by-case basis.

Metashape Professional Edition Settings

To keep projects consistent and comparable, applicable settings in Metashape Professional Edition (referred to as Metashape for this report) are provided in this section. Some of the settings may require a deeper understanding of photogrammetry, computer science, or Metashape terms. Note that the term “camera” in Metashape is used to describe a specific image at a specific time and location, and the term is also the tool used to take photographs. Metashape uses “image”, “photograph”, and “camera” interchangeably. This creates an unwanted complexity in the software; in this text, the usage of “camera” or “sensor” will be used to describe the physical instrument and “image” will be used wherever possible to describe a photograph at a specific time and location. Settings and terminology may be further explained or become clearer later in this report.

The Metashape workspace has eight display panels, all of which can be toggled on and off under the “View” menu. For example, a workspace may use the Workspace, Photos, Reference, and Console panels (fig. 2). The general steps to process a project are listed under the Workflow menu and are detailed in this section. Follow Tools >> Preferences to access the general and advanced settings dialog.

External Inputs and Processes

Metashape processes imagery into spatially georeferenced products. The shoreline imagery inputs are either aerial nadir, near-nadir, or high-angle oblique images. Metashape does not support camera-specific (.raw) formats, so the images must be converted to a compatible format (such as .dng, .tiff, or .jpg). The quality of the imagery will also affect or limit the quality of the final products; projects with low resolution or blurry image inputs may not be fully supported. Image quality can be assessed based on sharpness in Metashape by right-clicking on a photo >> Estimate Image Quality >> “Apply to” >> “All cameras”. A quality assessment of 1 (unitless) is considered the best possible; consider removing or disabling images with a quality less than 0.5 from the project. The images must also have position information, which may be derived from postprocessing kinematic (PPK) techniques, or at least three ground control points (GCPs) in the

Table 1. Quick start guide for the structure from motion workflow used to aid hurricane relief and recovery efforts by the U.S. Geological Survey.

[GPS, Global Positioning System; INS, inertial navigation system; m, meter; RMSE, root mean square reprojection error value (reported in pixels); NAD83(2011), North American Datum of 1983 with 2011 adjustment; UTM, Universal Transverse Mercator; NAVD88, North American Vertical Datum of 1988; GCP, ground control point; DEM, digital elevation model; 3D, three-dimensional; 4D, four-dimensional; >, greater than; <, less than; ≤, less than or equal to; %, percent; —, no specification]

Step	Menu	Function	Action	As used in this workflow
			Image setup and alignment	
1	Main Menu >> Workflow	Add imagery	Navigate to directory with imagery and select all imagery associated with the project	.jpg file format
2	Workspace pane >> Cameras	Add camera group	Put images from same collection into a new camera group	For example, separate images from 2018-10-06 and 2018-10-07 into camera groups
3	Main Menu >> Tools	Camera calibration	Import precalibration (optional)	—
			Set GPS or INS offset for lever arm	0.1/−0.15/1.25 (0.05) (in meters)
			Create camera calibration groups	Make for each 4D or 3D collection
	Reference panel >> Import	Add camera positions (optional)	Import file of coordinates and altitude for each image	.txt file of latitude, longitude, and height above ellipsoid
4	Reference panel >> Settings	Set camera accuracy	Under Measurement Accuracy, set Camera accuracy	0.1/0.15 (in meters)
		Set tie point accuracy	Under Image Coordinates Accuracy, set “Tie point accuracy”	1 (default; in pixels)
		Set camera coordinate reference system	Select from “Coordinate System” dropdown	NAD83(2011)
5	Main menu >> Workflow	Align imagery	Change settings to: high, generic, reference-source, 60000, 0	—
			Add ground control points (optional)	
6	Reference panel >> Import	Add GCPs	Import or detect	Load from .csv file
7	Reference panel >> Settings	Set GCP accuracy	Under “Measurement Accuracy”, set “Marker accuracy”	0.02/0.03 (in meters)
			Under “Image Coordinates Accuracy”, set “Marker accuracy”	0.5 (in pixels)
		Set GCP coordinate reference system	Make sure dropdown “Coordinate System” matches cameras in step 5	NAD83(2011)
8	Reference panel >> Markers	Filter by marker	Manually adjust GCPs on individual images	—
			Error reduction and bundle adjustment	
9	Workspace panel >> Chunk	Duplicate	Before optimization, duplicate chunk and rename	Uncheck “Keep key points”
10	Reference panel >> Optimize	Initial optimization	Check on: f, cx, cy, k1, k2, k3, p1, p2, tie point covariance	—

Table 1. Quick start guide for the structure from motion workflow used to aid hurricane relief and recovery efforts by the U.S. Geological Survey.—Continued

[GPS, Global Positioning System; INS, inertial navigation system; m, meter; RMSE, root mean square reprojection error value (reported in pixels); NAD83(2011), North American Datum of 1983 with 2011 adjustment; UTM, Universal Transverse Mercator; NAVD88, North American Vertical Datum of 1988; GCP, ground control point; DEM, digital elevation model; 3D, three-dimensional; 4D, four-dimensional; >, greater than; <, less than; ≤, less than or equal to; %, percent; —, no specification]

Step	Menu	Function	Action	As used in this workflow
		Error reduction and bundle adjustment—Continued		
11	Main Menu >> Model >> Gradual Selection	Reconstruction uncertainty (geometry)	Set level: 10 (if >50% points are selected, increase until <50% points are selected)	—
			Delete points and optimize (optional)	—
			Check on: f, cx, cy, k1, k2, k3, p1, p2, tie point covariance	—
			Monitor: projections goal >100, RMSE (in pixels) ≤0.3	—
12	Main Menu >> Model >> Gradual Selection	Projection accuracy (pixel matching errors)	Set level: 2 (if >50% points are selected, increase until <50% points are selected)	Projection accuracy set to level 3
			Delete points and optimize	—
			Check on: f, cx, cy, k1, k2, k3, p1, p2, tie point covariance	—
			Monitor: projections goal >100, RMSE (in pixels) ≤0.3	—
13	Main Menu >> Model >> Gradual Selection	Reprojection error (pixel residual errors)	Set level: 0.3 or the level when about 10% of points are selected	—
			Delete points and optimize	—
			Check on: f, cx, cy, k1, k2, k3, p1, p2, tie point covariance	—
			Repeat until level 0.3 selects no points or <10% of original tie points remain	—
			Monitor: projections goal >100, RMSE (in pixels) ≤0.3	Stop if <10% original number of tie points remain, the Errors column exceeds Accuracy column in the Reference panel, or the Errors column starts to increase
14	Reference panel >> Optimize	Optimization (optional)	Check on: f, cx, cy, k1, k2, k3, p1, p2, additional coefficients, tie point covariance, and fit additional corrections	Coefficients b1, b2, and k4 are inappropriate to model or enable with most consumer grade cameras.

Table 1. Quick start guide for the structure from motion workflow used to aid hurricane relief and recovery efforts by the U.S. Geological Survey.—Continued

[GPS, Global Positioning System; INS, inertial navigation system; m, meter; RMSE, root mean square reprojection error value (reported in pixels); NAD83(2011), North American Datum of 1983 with 2011 adjustment; UTM, Universal Transverse Mercator; NAVD88, North American Vertical Datum of 1988; GCP, ground control point; DEM, digital elevation model; 3D, three-dimensional; 4D, four-dimensional; >, greater than; <, less than; ≤, less than or equal to; %, percent; —, no specification]

Step	Menu	Function	Action	As used in this workflow
		Error reduction and bundle adjustment—Continued		
15	Reference panel >> Settings	Tighten tie point accuracy and continue Reprojection error gradual selection (optional)	Change tie point accuracy (in pixels) from 1 to 0.1–0.3	—
			Optimize	—
			Check on f, cx, cy, k1, k2, k3, p1, p2, additional coefficients, tie point covariance, fit additional corrections	—
			Monitor: projections goal >100, RMSE (in pixels) ≤0.18	—
			Stop if RMSE ≤0.18, <10% original number of tie points remain, or the Errors column exceeds the Accuracy column in the Reference panel or RMSE begins to increase	
			Repeat step 14 using 10% of points per gradual selection. Stop if RMSE (in pixels) ≤0.18 or begins to increase	—
		Build dense point clouds, DEMs, and orthomosaics		
16	Main Menu >> Region	Resize output region	Use resize and rotate region tools to redefine output area	—
17	Workspace panel >> Chunk	Duplicate	Create duplicate chunk for each 4D flight chunk and rename	Not necessary in 3D
		Disable or delete	Disable or delete cameras not related to product being produced	Repeat for each flight chunk
18	Main Menu >> Workflow	Build dense cloud	Set Quality: high; depth filtering: mild; check all	—
	Main Menu >> Tools or Main Menu >> Model	Edit dense cloud	Use filter by confidence or manually select flyers and sinkers	—
19	Main Menu >> Workflow	Build mesh (optional)	Height field, dense cloud, high, enabled	If built from depth maps, the setting is arbitrary
		Build DEM	Select from "Coordinate System" dropdown, dense cloud, interpolation disabled	Build in NAD83(2011)
		Build orthomosaic	Build from DEM, mosaic	Build from a separate interpolated DEM to avoid holes in the orthomosaic
		Export and file naming		
20	Main Menu >> File	Export	Change coordinate reference system to UTM projection and vertical datum (may need to download a geoid)	NAD83(2011) UTM Zone 18N + NAVD 88 (meters)
			Set resolution size	0.25 m
			Set region boundaries to round numbers	Nearest 10s of meters

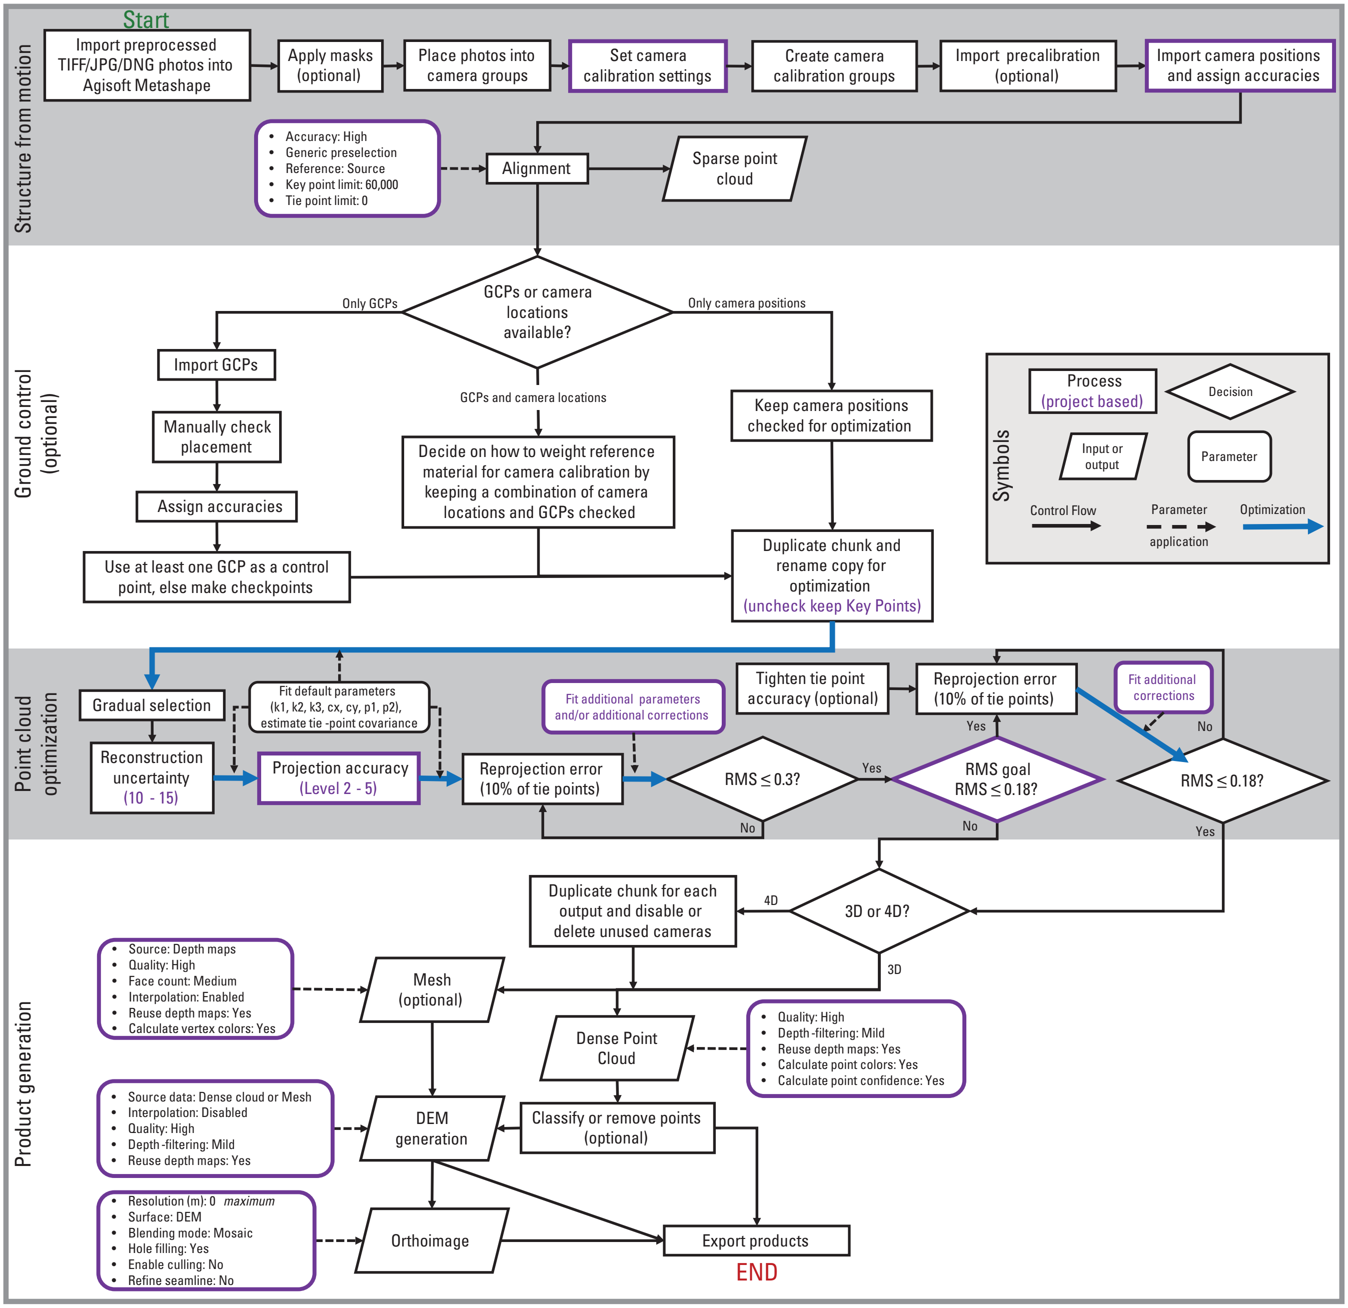

Figure 1. Schematic diagram summarizing a workflow using Agisoft Metashape Professional Edition. The workflow conceptualizes four main stages: structure from motion, ground control, point cloud optimization, and product generation. Each stage contains a series of processes, decisions, inputs, and parameters important for consideration in the software. The workflow can be applied to projects with optional inputs. Steps or inputs that may change based on the requirements of a project are indicated by purple text or a purple outline. GPS, Global Positioning System; INS, inertial navigation system; GCP, ground control point; RMS, root mean square; %, percent; ≤, less than or equal to 3D, three-dimensional; 4D, four-dimensional; DEM, digital elevation model; m, meter.

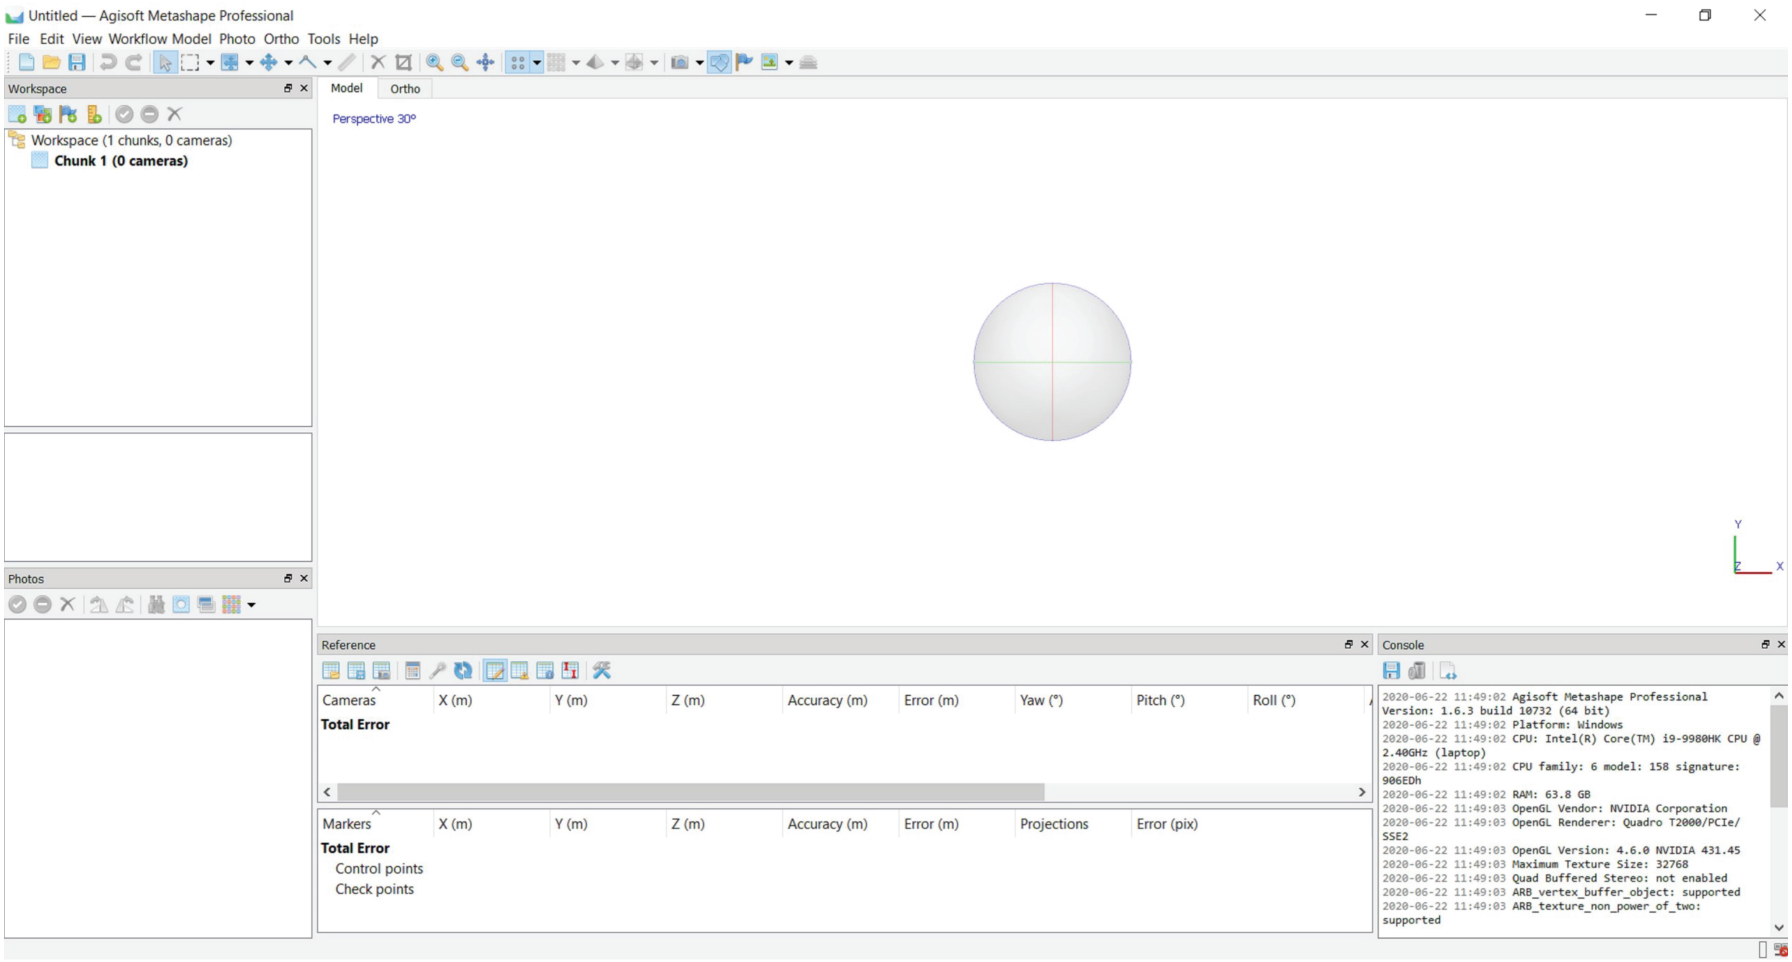

Figure 2. Screenshot of a new, blank project workspace in Agisoft Metashape Professional Edition, showing a standard arrangement of working panels (Workspace, Photos, Reference, and Console) for a single-monitor workstation.

reconstructed scene to produce georeferenced photogrammetry products (Smith and others, 2016; Sanz-Ablanedo and others, 2018); Agisoft suggests at least 10 evenly spread GCPs for high quality results (Agisoft LLC, 2020b). Well-constrained camera positions (at decimeter accuracy) can be used to create products with horizontal and vertical uncertainties of less than 1 meter (m; Turner and others, 2014), but the addition of even one accurately-located GCP to the project is enough to increase overall vertical accuracies (Forlani and others, 2018). Although SFM products can be created without GCPs, their addition to the workflow is encouraged, if only to check the accuracy of the camera alignment and the final products.

The Metashape workflow used for this report uses a lossless .jpg 8-bit file format, determined to be the best compromise between file size and quality for the scale of project. Images in the workflow used for this report were captured in true color with about 75 to 80 percent forward overlap and 60 to 70 percent side overlap from an altitude of 300 m from a piloted, fixed-wing aircraft. For reproducibility context, SFM products have been made with the same workflow using imagery obtained from drones with set flight paths, greater overlap, and lower altitude than for imagery in this study, as well as imagery obtained from the National Oceanic and Atmospheric Administration (2020) that have less overlap and are captured from a higher altitude. Note that insufficient overlap may result in poor Metashape reconstructions and a distinct “train-track” pattern where there are consistent gaps in the point cloud. The data on the camera position for the workflow used for this report were collected with a global navigation satellite system (GNSS), stored in an external file, and corrected using PPK techniques.

The GNSS signals at the antenna mounted atop the aircraft were collected at a sample rate of 1 hertz by a dual-frequency survey-grade GNSS receiver. These data, along with reference data from no fewer than four National Geodetic Survey continuously operating reference stations and precise satellite ephemerides published by the International GNSS Service, were combined in NovAtel GrafNav (version 8.80.2720) GNSS postprocessing software (NovAtel, 2020) to produce trajectories with low (1-hertz) GPS sampling frequency with an estimated horizontal accuracy of 10 cm and a vertical accuracy of 15 cm (Kranenburg and others, 2020). This trajectory, in combination with precisely recorded shutter activation event times, generates antenna positions at the moment of each image capture. Metashape treats the camera position as the nodal point of the lens, the place where all light rays converge from the scene before being projected onto the sensor. A more accurate camera position is obtained when the lever arm offset from the nodal point of the lens to the antenna reference point of the GNSS antenna (GPS or inertial navigation system [INS] offset) is measured within the frame of reference of the camera.

File Types

Metashape supports two project file types. The .psx file type is uncompressed and has companion files that must be copied together with the Metashape working file. Project .psz files are compressed and will have longer loading and saving times but do not have companion folders. The .psx file format supports DEMs, orthomosaics, and network processing, whereas the .psz file format does not.

Within each project file, Metashape parses imagery into “chunks”, which is what groups of images, camera models, and products listed and edited in the Workspace panel are called in Metashape. Chunks can be duplicated during the workflow process as an easy way to save changes because there is no way to undo certain processes. Only one chunk can be active at any time. Duplicated chunks are, by default, exact copies and include key points and products such as depth maps unless specifically deselected, which can save space (fig. 3*B*). Chunks should be duplicated before major changes or after a time-consuming step. For example, duplicating a chunk after alignment creates an active copy, which can be used in further processing, and the now inactive original chunk can be returned to or duplicated again if a mistake in subsequent steps requires a need to reprocess from the previous state. In the example, without a duplicated chunk, the steps up to and including the alignment would need to be redone. To work with a chunk, right-click on the chunk to open a context menu; actions include adding images and folders and exporting products. The “Show Info…” option from the context menu can be selected to observe associated errors, number of tie points, and all other relevant parameters of the chunk as it has been processed (fig. 3*A*).

Metashape Preferences Dialog

The Metashape Preferences dialog allows users to adjust general settings for the software. This section describes settings that may affect the workflow, but not all settings will be addressed (for example, the “Navigation” tab); for settings not mentioned, the default setting is assumed; a comprehensive description of all settings is available in the Metashape user manual. Metashape Preferences may be reset to the manufacturer’s default in the “Advanced” tab.

General Tab

Changes to language, units, keyboard shortcuts, and miscellaneous preferences are detailed in this section (fig. 4*A*).

User Interface >> Default view

If the “Default view” is set to anything other than Point Cloud, (such as Dense Cloud or Model), there is a chance of the program crashing if the dense cloud or model file size is unexpectedly large. This can be a serious consideration with

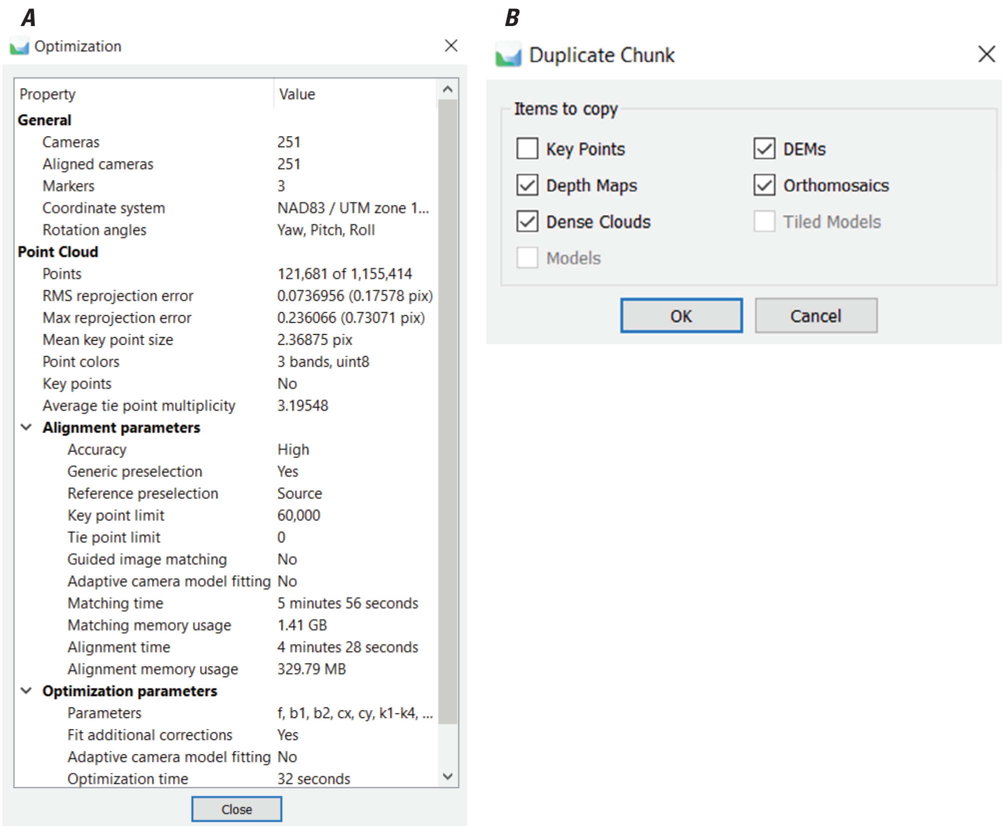

Figure 3. Screenshots of dialogs in Agisoft Metashape Professional Edition used with chunks. *A*, The Show Info dialog for a chunk named “Optimization” that displays the various properties and values of a chunk; and *B*, The Duplicate Chunk dialog with various products in the “Items to copy” list selected.

large datasets, such as the complete dataset generated for the workflow used for this report, which has over 35,000 images (Kranenburg and others, 2020). The “Point Cloud” setting will also help shorten loading time when opening a project.

Miscellaneous >> Write log to file

Every action in Metashape is timestamped and viewable in the Console panel and is also logged to a file if “Write log to file” is selected. It is good practice to enable “Write log to file” and keep track of the log file size and available space. Very large log files may be difficult to open with many text viewers or editors. New projects warrant a new log file.

GPU Tab

Discrete (not part of the motherboard chipset) graphics processing units (GPUs) can speed up processing substantially. The “GPU” tab is used to display all the GPUs available for use and enable and select the GPUs to be used for processing. Performance gains using GPUs are a function of the number of NVIDIA Compute Unified Device Architecture (CUDA) cores or AMD brand shader processing cores, the amount of video random access memory (VRAM), the speed and bandwidth of the peripheral component interconnect express (PCIe) bus, the clock speed of the GPU and central processing unit (CPU), and the availability of CPU cycles for data transfer to and from the GPUs. Performance improvements are variable based not only on the type of GPU installed but on the overall system architecture. Puget Systems regularly posts benchmarking tests using Metashape on different GPUs (Puget Systems, 2020); discussions that provide helpful information are found on the Agisoft forum (Agisoft LLC, 2021a). The “Use CPU when performing GPU accelerated processing” option is a feature that enables CPUs to work on tasks, such as selecting and matching tie points, that otherwise would be performed much faster by the GPU (fig 4*B*). The option, if enabled, can slow down processing or cause a workstation to become unresponsive, but in systems with many or fast CPU cores, it may speed up processing. Before enabling, benchmark a project large enough to keep the GPUs at full capacity for at least several tens of minutes because there is substantial CPU overhead required to send or receive data to and from the GPUs.

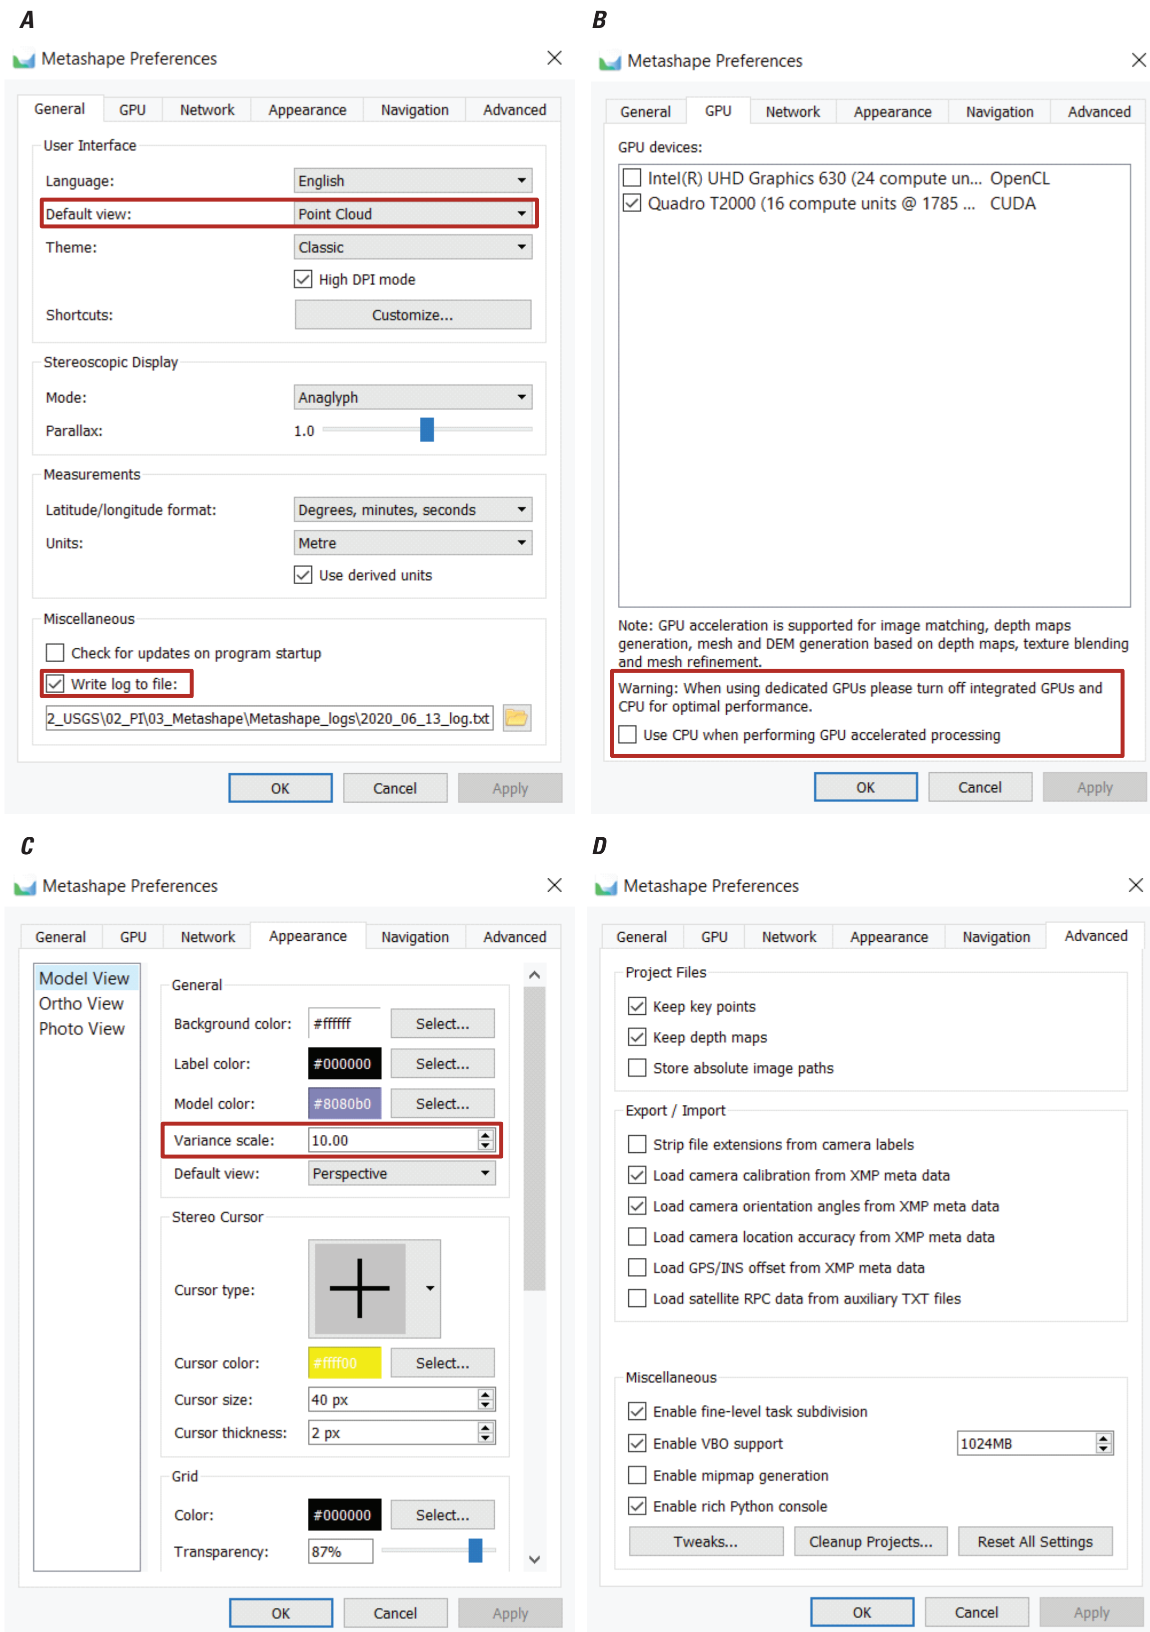

Figure 4. Screenshot of Metashape Preferences dialog in Agisoft Metashape Professional Edition. *A,* The "General" tab; *B,* the "GPU" tab; *C,* The "Appearance" tab; and *D,* The "Advanced" tab. The settings used in the structure from motion workflow for U.S. Geological Survey hurricane relief and recovery efforts are selected in each tab. Red outlines highlight option boxes changed from the default or that are discussed in detail in the text.

Appearance Tab

This tab contains options for changing color schemes and properties for the three basic views: Model, Ortho, and Photo, and tools within the views, such as masks, cursors, and hillshading (fig. 4*C*). In the Model view, the only setting we suggest changing is the “Variance scale”. This option can be found at Model view >> General >> “Variance scale”.

The “Variance scale” is a visualization of estimated position error of tie points displayed as vectors whose length and color correlate to the magnitude of the error. A higher value will make inspection of error distribution easier in subsequent stages of the workflow. The suggested value is 10. A discussion on variance and where it can be used can be found in the “Error Reduction—Optimization and Camera Calibration” section of this report.

Advanced Tab

The following advanced settings were used to process the data for the workflow detailed in this report (fig. 4). Miscellaneous >> “Reset All Settings” will revert the project to default settings, even those not under the Metashape Preferences dialog. Select the following options:

- Project Files >> “Keep key points” is important because Metashape supports incremental image alignment, which is the addition of new images without resetting the alignment (see the “Starting a New Project” section of this report for more information on key points). However, this option adds substantially to the file size; keeping key points can easily double the project size with multiple chunks, but copied chunks can have key points removed once the image alignment is set (see “Other Settings and Features” section of this report).

- Project Files >> “Keep depth maps” allows multiple dense cloud, depth map, or mesh products to be created without recomputing the depth maps each time. This option can save time. Keeping depth maps also adds to file size, though the increase is relatively small compared with the increase from key points or images; to reduce file size increases, depth maps can be turned off when copying chunks.

- Miscellaneous >> “Enable fine-level task subdivision” enables out-of-core processing and reduces memory requirements. This allows processes to be split between cores and can improve performance when working with large datasets. In tests with a few thousand images, fine-level task subdivision improved both speed and alignment results. This option must be enabled for network processing.

- Miscellaneous >> “Enable VBO support” generally improves performance of dense clouds and meshes with high-end video cards by storing model details (vertices or points) in GPU memory. In some cases, a driver issue or using an integrated (onboard) graphics cards can make Metashape graphically unstable, leading to software crashing. If this occurs, uncheck this option.

- Miscellaneous >> “Enable rich Python console” can be selected if running scripts in the Console panel because the console provides more convenient autocomplete functionality with the Tab key, adds color coding for syntax, and allows the user to input and run multiline script snippets rather than line by line.

Deselect the following options:

- Project Files >> “Store absolute image paths”, which, when enabled, makes it harder to move project (PSX) and image files stored on the same volume to a different volume without breaking projects. If images are stored on a separate drive, this setting is irrelevant. Metashape also has built-in tools to repair broken image links; right-click on an image in the Workspace panel and select “Check Paths…”.

- Export/Import >> “Strip file extensions and camera labels” should be deselected, because it can cause image loading problems when using scripts. Including the file extension in the file name prompt is usually useful.

Other Settings and Features

Other settings regularly accessed in Metashape by the user include the Reference Settings and Camera Calibration dialogs. The Reference Settings dialog is the interface used to change the coordinate reference systems and accuracies of cameras and markers. The dialog is accessed through the crossed hammer and crescent wrench symbol on the far right of the Reference panel menu (or button; see the “Adding Camera Positions”, “Adding Control Points and Check Points”, and “Error Reduction—Optimization and Camera Calibration” sections of this report). Follow Tools >> Camera Calibration to change settings specific to the camera, such as the frame type, focal length, and options related to lens coefficients. The Camera Calibration dialog is covered in more detail in the “Adding Images, Camera Groups, and Camera Calibration Groups” and “Error Reduction—Optimization and Camera Calibration” sections of this report.

Starting a New Project

This section provides an overview of the SFM workflow (fig. 1) and defines common terms. For a more comprehensive introduction to the theory of photogrammetry and SFM with diagrams see, for example, Schenk (2005) and Carrivick and others (2016). A typical SFM project in Metashape processes imagery in two sequential steps: (1) image alignment and (2) generation of a spatial product. Alignment is the process of searching for unique feature points, known as “key points”, in each individual image and matching them across images pairwise into what are known as “tie points”. During alignment, the geometry of the scene is also reconstructed; the program finds the position of tie points relative to the cameras in a process commonly called aerial triangulation. This process calculates the camera position (x,y,z) and attitude (ω,φ,κ), which is the exterior orientation, for each image. Alignment, or more specifically the aerial triangulation, is aided by the addition of estimated or observed exterior orientations; however, a remarkable aspect of SFM is that these orientations are not required. In Metashape, alignment parameters are computed in a local spherical coordinate system and translated to a geographic coordinate system only when a coordinate system is defined. Alignment also refines the camera calibration model through the interior orientations, which refers to camera lens and sensor variables, recorded in Metashape as the camera calibration model, and consists of a set of coefficients (f, cx, cy, b1, b2, k1, k2, k3, k4, p1, p2; discussed in the “Adding Images, Camera Groups, and Camera Calibration Groups” section of this report) that describe camera lens and sensor geometry.

The final products of the alignment step are a calibrated camera model, estimated locations and orientations of cameras, a set of estimated camera calibration coefficients, a geolocated set of tie points with their red-green-blue (RGB) values (the sparse point cloud), and the estimated location of GCPs, if any. The GCPs, named “markers” in Metashape, can be added to improve the quality of the alignment (see the “Adding Control Points and Check Points” section of this report). The camera calibration can also be refined by the selective removal of weak tie points (see the “Error Reduction—Optimization and Camera Calibration” section of this report). The removal, by the process of targeted, iterative (gradual) selection of reconstruction uncertainty, projection accuracy, and reprojection error metrics, is critical for high-quality reconstructions. The elimination of tie points is balanced by the requirement of maintaining enough tie points to allow image alignment over the entire scene.

At each step, after the addition of GCPs or removal of tie points, the tie point positions, camera position, and lens model are adjusted in a least-squares sense. Metashape calls this “optimizing”; it is commonly known in photogrammetry as a “bundle adjustment” or “bundle block adjustment” and has been a component of photogrammetry since the 1950s (for example, Brown, 1958, 1976; Triggs and others, 2000). Incorporating iterative adjustments to the camera position (exterior orientation) and camera calibration model (internal orientation) into the bundle block adjustment, as is done in Metashape, is a modern optimization practice. The bundle adjustment is essentially a series of equations that tries to recreate, or solve, the camera sensor’s exterior orientation as it was exactly when it took each image; a minimization of the difference between the observed image locations and predicted image points. The bundle adjustment is, in this sense, the solution. In addition, specific metrics generated by Metashape during alignment and optimization can be used to evaluate the overall quality of the camera model. These include the standard error of unit weight (SEUW), the root mean square (RMS) reprojection error (both weighted and unweighted), mean key point size, and RMS errors in marker and camera locations.

After the geometric model and lens calibration are complete and the sparse point cloud has been edited, processing in Metashape can proceed to the second step of generating a spatial product (see “Creating Products” section of this report). Products may include the following:

- depth maps based on a multicamera dense matching algorithm using previously estimated camera positions;
- a dense point cloud built from the depth maps;
- a polygonal model (mesh) of textured or untextured vertices and faces;
- a DEM, sometimes known as a digital surface model, to distinguish from ground-only elevation models, derived from either the sparse cloud, depth maps, dense cloud, or mesh; or
- an orthomosaic or RGB averaged product, made by projecting orthorectified images onto a surface (DEM or mesh).

Although the two broader steps above describe a general workflow common to most SFM projects, there are many details involved in each step (fig. 1). For postprocessing steps, such as point cloud classification, removing shadows, and calculating vegetation indices, see the Metashape user manual.

The approach to photogrammetric reconstruction of landscapes from aerial imagery may be considered three-dimensional (3D) or four-dimensional (4D). A 3D process, which is discussed in this report, involves the reconstruction of a scene at one instance, using images taken as close to simultaneously as possible. The comparison of 3D projects of the same area can produce change maps (Chirico and others, 2020), but a 4D workflow may be more appropriate for repeat surveys.

Processing in 4D refers to reconstruction of a scene using images from multiple collection efforts with the intent of detecting a change in the same scene. Two or more sets of images are aligned together with the assumption that there are substantial areas of unchanged or stable features between the collections. These unchanged features can strengthen the reconstruction and alignment compared with separately processing the different sets of images in 3D. Warrick and others (2017) and Sherwood and others (2018) demonstrated the use of 4D processing. This report includes details where the workflow differs in 3D and 4D processing. The appropriate method depends on the number of image datasets available, the time between the datasets, and amount of landscape change between data collection efforts. For example, recurring aerial surveys of a beach after storms within a year with the same coverage would benefit from a 4D processing method. On the other hand, an attempt to align a nadir flight with an oblique flight of the same beach years apart may not benefit from a 4D approach.

The SFM workflow described in this report assumes use of the Metashape graphical user interface (GUI). However, the majority of the workflow can be done either in batch processing (see the “Batch Processing” section of this report) or with scripting through the Agisoft application programming interface (API; Agisoft LLC, 2020c). Example source code for Metashape is available from Logan and Ritchie (2020). The examples and screenshots in this report are taken from a project of about 250 near-nadir aerial images, part of a larger imagery data release (Kranenburg and others, 2020), taken on different dates at Pea Island National Wildlife Refuge in the Outer Banks of North Carolina after Hurricane Florence; the project was processed using the 4D technique described by Sherwood and others (2018). Camera locations derived from both postprocessed kinematic GNSS and GCPs were available for this project (Kranenburg and others, 2020; Brown and others, 2021).

Adding Images, Camera Groups, and Camera Calibration Groups

A project in Metashape starts with the addition of images. To add images, use Workflow >> “Add Photos…”. This will create a “Cameras” folder under the active chunk in the Workspace panel and display as “Cameras (0/n aligned)”. Once a “Cameras” folder is created, camera group subfolders can be created in the Workspace panel. Camera groups are useful for project organization and can be used to separate images by collection, location, or camera type, and commonly have an associated camera calibration group (CCG). It is important to note that camera groups in the Workspace panel are treated completely independently from the camera groups in the Camera Calibration dialog by Metashape.

To create a new group, in the Workspace panel, right-click on any “Cameras” folder and select “Add Camera Group”. To add images directly into a group from a “Cameras” folder, select a series of images and drag into the group or right-click “Move Cameras” and select “New Camera Group” or right-click on a group and select “Add Photos…” (fig. 5). Rename groups to distinguish among separate groups. Actions affecting Camera groups such as select, enable, or disable are done in subsequent product creation steps.

Next, open a dialog from Tools >> Camera Calibration, to display information about any automatically or manually created CCGs. This is the space to view or modify CCG properties and export or import camera calibration models. If the images added have different dimensions, focal length, or other parameters of their associated exchangeable image file (EXIF) format information, multiple CCGs may automatically be created in the left panel. If EXIF data are corrupted or missing, make a new CCG for each group of images that have different collection dates, altitudes, or resolutions, or for each group of images that were taken with different camera settings, such as focus or aperture settings. Depending on the camera model, a simple powering down and back up can change the lens geometry (powering down on some camera models retracts the lens, changing the focal length or resetting the aperture) enough to possibly warrant a new CCG. Understanding when to assign different CCG requires a careful consideration of the stability of the relation between the lens and sensor, particularly when consumer-grade cameras are used.

Photogrammetric calibration is performed on each CCG to determine the lens calibration coefficients. These coefficients make up the internal orientation, the parameter that links the pixel coordinates of an image point with the corresponding coordinates in the camera reference frame. The exterior orientation refers to the position and orientation of the camera with respect to a world reference frame. Collectively, the interior (camera model) and exterior (camera pose) orientations specify the transformation from point coordinates in the local camera coordinate system to the pixel coordinates in the image. The camera model is calibrated with the goal of approximating the camera lens sensor geometry that produced a given image. This transformation or approximation is modeled using equations that differ with the camera model type and is one of the more complex configurations in Metashape; frame cameras are not solved the same as fisheye or spherical cameras. The equations differ, but corrections of lens sensor geometry of all camera models in Metashape are solved with the same interior orientation variables. The following coefficients are tracked in the Camera Calibration dialog and can be solved for in the optimization steps (see the “Error Reduction—Optimization and Camera Calibration” section of this report):

f	focal length;
cx, cy	principal point offset;
k1, k2, k3, k4	radial distortion coefficients;
p1, p2	tangential distortion coefficients; and
b1, b2	affinity and nonorthogonality (skew) coefficients.

The coefficients k4, b1, and b2 are not default Metashape coefficients and must be explicitly selected to be used in the alignment (see the “Aligning Images” section of this report) and manually added in the camera calibration optimization procedure. Set up the “Camera Calibration” dialog before alignment (fig. 6), following these steps (data input into steps for the workflow in this report are listed in brackets at the end of the appropriate step):

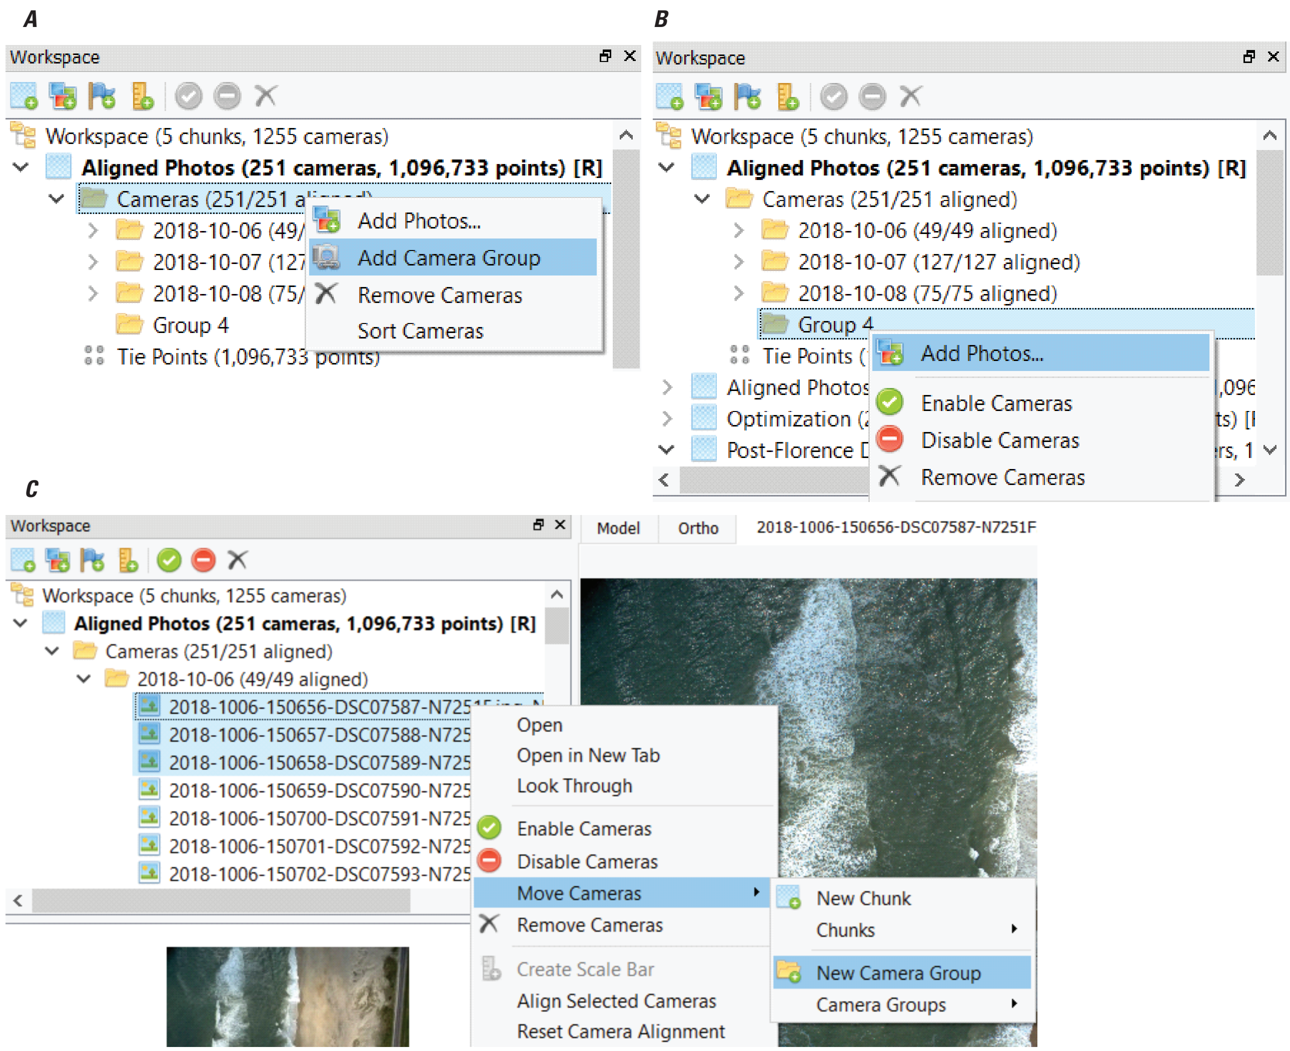

Figure 5. Screenshots of the Workspace panel in Agisoft Metashape Professional Edition highlighting options to “Add Camera Group” and “Add Photos…” to camera groups. *A*, Create a camera group by right-clicking on a folder and selecting “Add Camera Group”; *B*, Add images to a camera group by right-clicking a camera group folder and selecting “Add Photos…”; and *C*, Move one or more images to a new camera group by right-clicking the selected images, clicking “Move Cameras” and selecting “New Camera Group”.

1. Set “Camera type” [Frame]. Choose based on equipment used to collect imagery.
2. “Pixel size (mm)”: and “Focal length (mm)”: Fill in the boxes if known.
3. Go to the “GPS/INS Offset” tab and check the “Enable reference” box to enter the applicable information [0.1/−0.15/1.25 (0.05)].
4. Select the images from the list at the bottom of the dialog to be added to a CCG, distinguished by date, time, name, and other important information, then right-click and select “Create Group”. Note that preselecting the images by camera group in the Workspace panel will carry through the selection into the dialog.
5. Right-click the new CCG in the left panel of the dialog and select “Rename Group”.
6. (Optional; see the “Precalibration” section of this report) Load or import a precalibration settings file from the yellow folder icon in the “Initial” tab. The “Type” will switch from “Auto” to “Precalibrated”, and a pencil symbol will appear on the selected CCG camera icon.
7. Repeat steps 4–6 as necessary for each CCG. Note that for steps 1–3 to be applied to all CCGs, they must be performed prior to steps 4–6; otherwise, each CCG must be individually selected, and the information reentered.
8. Click “OK”.

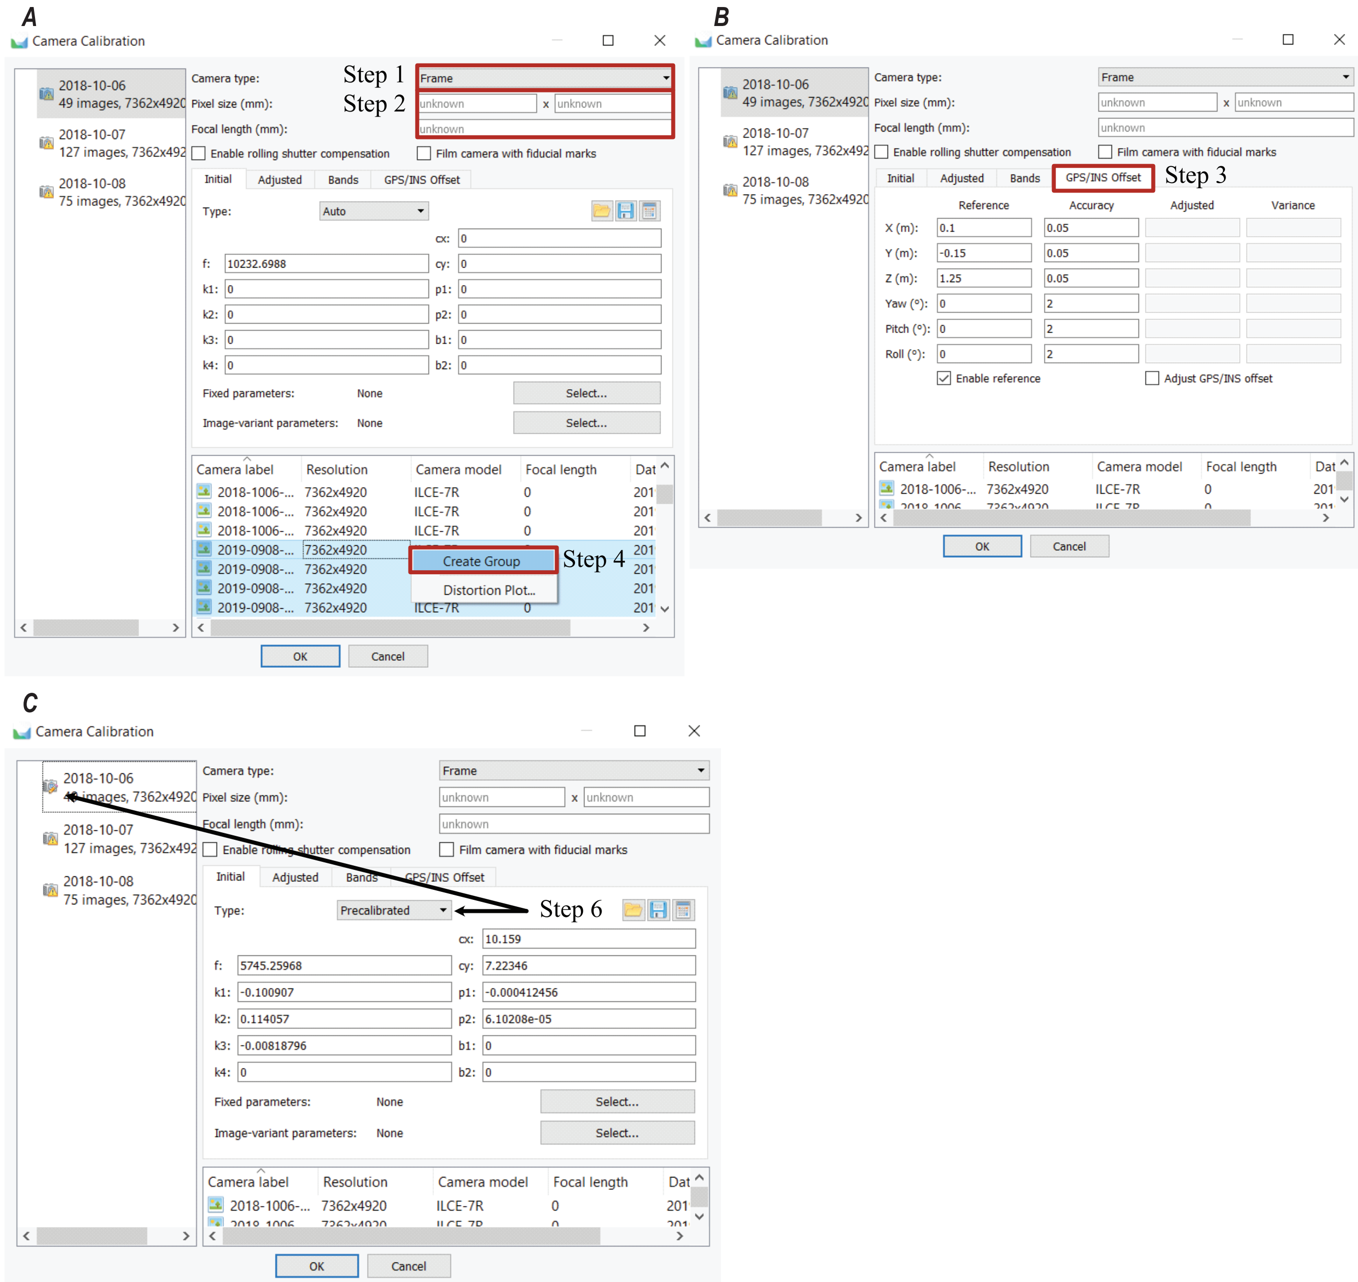

Figure 6. Steps used to set up appropriate prealignment settings in the Camera Calibration dialog in Agisoft Metashape Professional Edition, with certain tabs and buttons indicated with red boxes and black arrows for different steps. *A*, View of the "Initial" tab for camera calibration group (CCG) 2018-10-06 before a precalibration with steps 1, 2, and 4; *B*, View of the "GPS/INS Offset" tab completed with step 3; and *C*, View of the "Initial" tab showing that CCG 2018-10-06 has been precalibrated (step 6).

Precalibration

The optimization process can be sped up by preinitializing lens coefficients in the camera model from a previously solved model. A precalibration file of the best or most recent camera model with lens coefficients solutions can be loaded from the Camera Calibration dialog (fig. 6). Metashape supports the loading of precalibration formats commonly used by other organizations and software including Pix4D (.cam) and USGS camera calibration files (.txt); the "Tie points and camera calibration, orientation data export" section of the Metashape user manual includes a complete list of available precalibration formats. Within Metashape, a camera calibration will be refined based on initial values set by a previous dataset so that it learns from previous iterations (transfer learning), a functionality that is designed to speed up model convergence time and accuracy for similar datasets. If there are multiple camera groups that use the same precalibration, loading a precalibration model before all the groups are created will help avoid repeating steps 1–3 in the "Adding Images, Camera Groups, and Camera Calibration Groups" section of this report. At the end of the optimization process (see "Camera Calibration Check" section of this report), the adjusted coefficients can be exported in the same Camera Calibration dialog as an Agisoft .xml file or one of the other supported formats. Coefficients imported during the precalibration step can also be fixed in the Camera Calibration dialog using the "Fixed parameters" option in the "Initial" tab, which forces the software to use the camera model without further solving for any of the parameters. This should only be used for very stable camera models, but this can be useful if trying to align very few images, or images with low overlap.

Masks

A second, optional step after adding images is the application of masks, which can be used to limit the part of the original images analyzed by Metashape. This can be done at any step before alignment. Masks are binary images consisting of zeros and ones, where the zeros represent parts of the corresponding image to be masked. The masked part of the image is the part that the user does not wish to use for 3D (or 4D) reconstruction. This is particularly useful for aerial oblique imagery with large parts of water and projects focusing on shorelines. The masks are applied before the alignment process to remove propagating objects in the image set, such as moving vehicles and ocean waves, which can create false tie points and therefore false topography and high error. Masking could also be used to ignore whole classes of a scene, for example, all water or all vegetation if there was a need and if masking could be reliably achieved by an external image analysis procedure, such as automated image segmentation (Buscombe and Ritchie, 2018). Masking can also speed up alignment, partly because the collective number of pixels to be aligned and the area used in key point feature detection in the masked image set is greatly reduced from the original dataset. Regions of no interest or with high visual noise, if not masked out during alignment, can be masked from the dense point cloud or can be removed manually from the final product by editing the dense point cloud or final DEM.

Masks can be loaded from external sources or generated automatically in Metashape if background image data are available. Masks of a dynamic background must be generated individually for each image—for example, if a static object appears in the same area of the frame of all images, when that object or area is masked in one image, the mask can be applied to other images as well.

In the following example instructions, water is masked in aerial images from the data collection for the workflow used for this report (fig. 7):

1. Go to File >> Import >> Import Masks.
2. If masks were created outside of Metashape, in the Import Masks dialog, open the "Method" dropdown and choose "From File".
3. Open the "Operation" dropdown. Select the option that specifies what to do when another mask is applied to the same image. This includes replacing, merging, taking the intersection, or only keeping the difference between the two masks.
4. Assess the default "Filename template" which is "{filename}_mask.png". Change it to reflect how the external masks have been named and stored. The "{filename}" text must match the image names being masked that are already in the Metashape project.
5. Click "OK" to navigate to the folder where the masks are stored.
6. If successful, the masks will appear on the images as a shaded part. Mask color and transparency settings can be accessed from Tools >> Preferences >> "Appearance" tab >> Photo View >> Masks.
7. To manually place a mask or edit a mask, double-click on the targeted image in the Photos panel to view it.
8. Choose a tool from the Photo menu to select the part of image to be masked (such as the breaking waves in fig. 7).
9. Because the white of the breaking waves is a high contrast in color to the tan of the beach and the brown of static water, using the Magic Wand tool with the tolerance set very high will select the white areas.
10. With the Magic Wand tool selected, click on the breaking waves in the image. This will create a polygon based on the tolerance settings.
11. Right-click on the image and select "Add Selection" to merge the polygon of newly selected breaking water with any already masked parts of the image. The other options available when right-clicking also allow the mask to be reset or to invert the masked and unmasked parts.
12. Repeat until all parts of the image are masked to satisfaction.

Adding Camera Positions

Camera positions measured with GNSS can aid in the alignment process. Without position information, a point cloud generated in the alignment process is referenced to a local coordinate system, that is, a system with an arbitrary origin, pixel scaling in three dimensions, and no correction for the curvature of the Earth. When Metashape searches for key points and tie points between the images, position information will speed up the process and reference the points to a global coordinate system. Many image collection systems (for example, most UAS systems) embed position information in the EXIF part of the image, but those data can also be provided

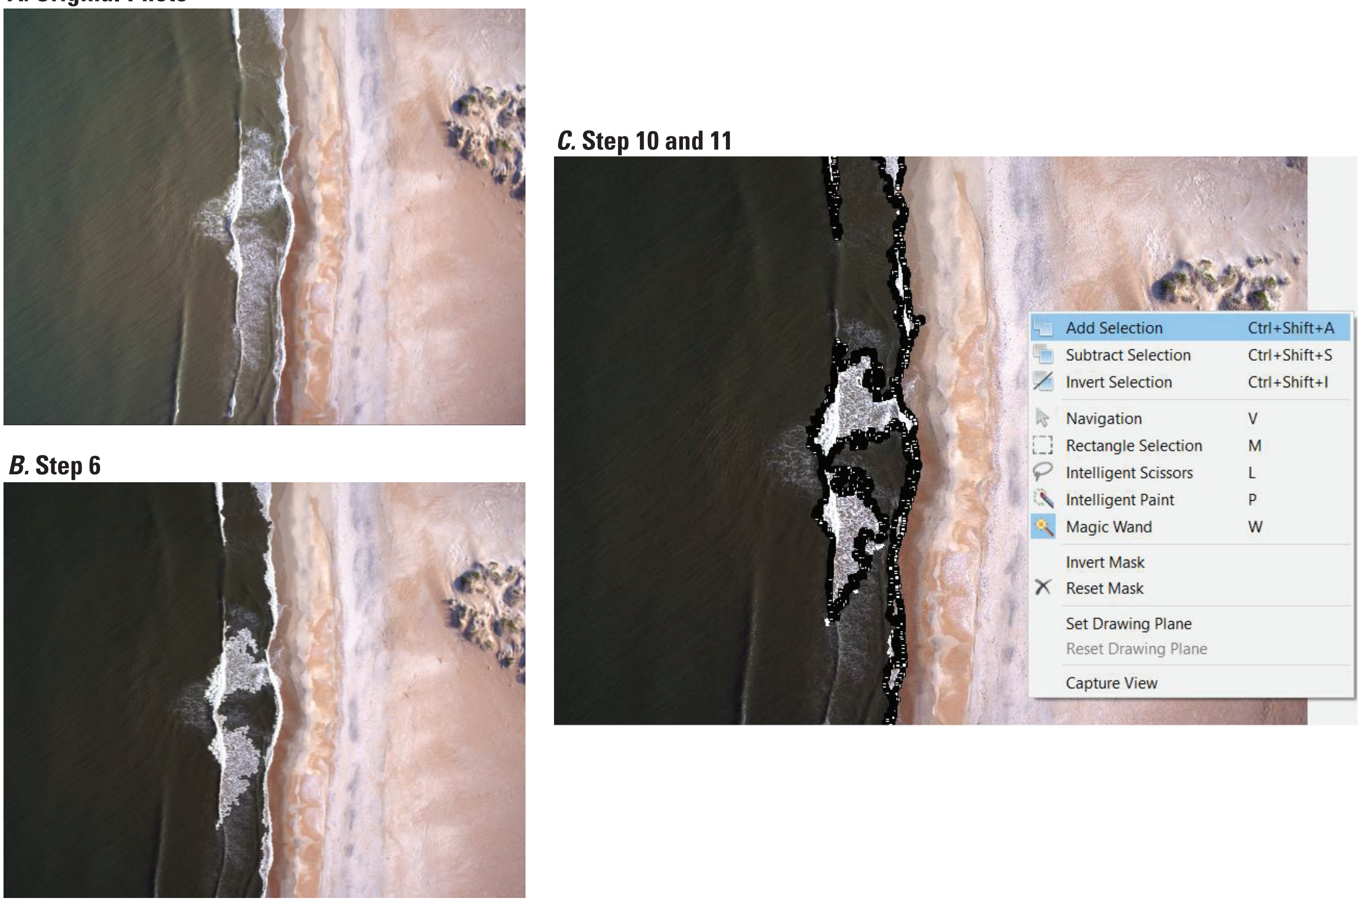

Figure 7. Screenshots of three stages in Agisoft Metashape Professional Edition of masking water from aerial imagery. *A*, The original image imported into Metashape; *B*, The water on the left side of the image has been masked—it appears darker in the Photo panel; and *C*, The Magic Wand tool has selected the white water, creating a polygon, with the nodes visible in the screenshot, and the "Add Selection" option that can be chosen to add the white water to the already masked water.

in an external file that matches image file names with location data. At any point in the workflow, position information of any image can be unchecked in the Reference panel; the positional data will not be used but the image is still incorporated into the products. Positional data might be suspect due to a lapse in collection conditions such as from weather or equipment failure. Images can be monitored in the Reference panel using the "Error (m)" column.

The camera positions in this workflow are in spherical coordinates: latitude, longitude, and ellipsoid height. For small survey areas, projected coordinates (for example, Universal Transverse Mercator (UTM) coordinates as northing, easting, and elevation) are sufficient, but do not account for the curvature of the Earth, which amounts to vertical deviations from the projected plane up to about 8 cm per kilometer. For larger survey areas, especially when precise camera altitudes are known, it is critical to work in spherical coordinates. This also prevents issues if the study area crosses into another UTM zone. Use as many decimal degrees as necessary to preserve desired measurement precision (six decimal places for decimeter precision, seven for centimeter precision, and eight for millimeter precision). Final products can be built or exported into a projected coordinate system (see section "Exporting Products and Error Reporting").

The following steps describe importing location data via a .csv or .txt file (fig. 8). Data input into steps for the workflow in this report are listed in brackets at the end of the appropriate step. Steps 1–8 will be repeated if positions are stored in multiple files. All camera location data are matched to the image using unique image labels, which includes the file extension.

1. From the Reference panel, select "Import Reference", which is an icon that looks like a spreadsheet with a small yellow folder in the bottom corner.

2. Navigate to the directory where the location information is stored and click "Open" to launch the "Import CSV" dialog.

3. Select an appropriate coordinate system.

4. Select the "Delimiter" box based on the data format.

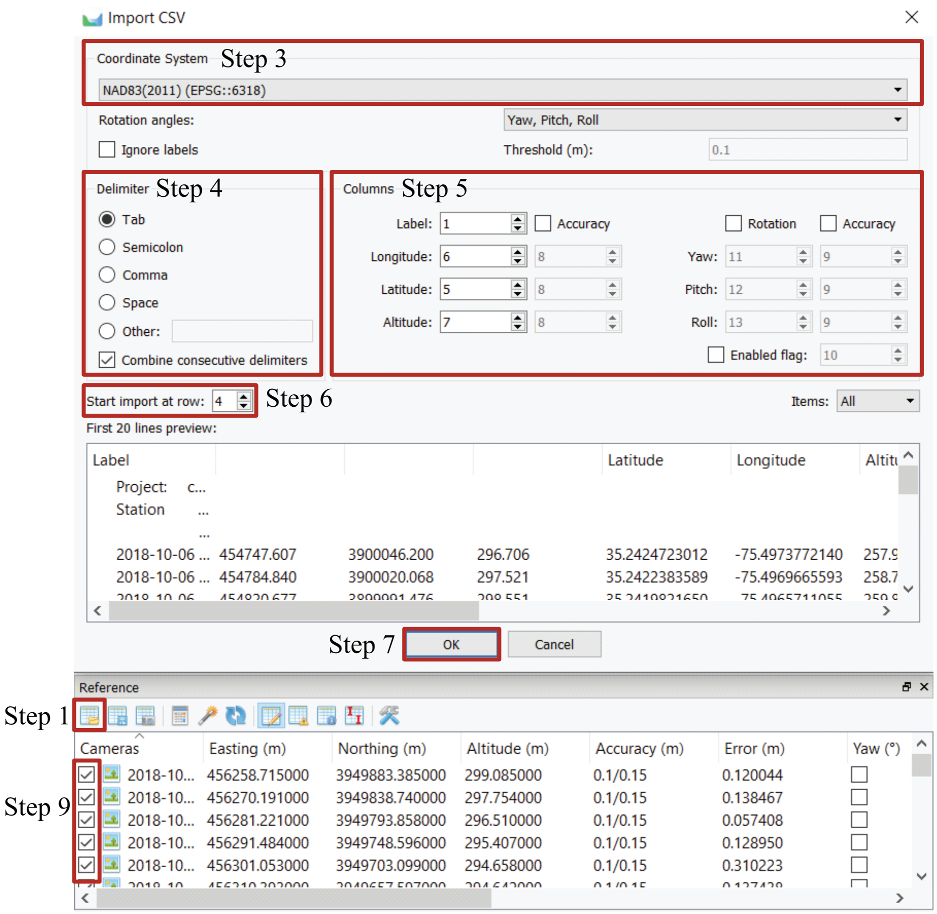

Figure 8. Screenshot of the Import CSV dialog in Agisoft Metashape Professional Edition that shows steps 1, 3–7, and 9 in the process to properly import camera positions into a project. The steps are outlined in red.

5. Under “Columns”, enter the column order listed in the camera location file of the x,y,z information as well as any rotation and accuracy values. If these values are unknown or unavailable, uncheck the respective “Rotation” and “Accuracy” boxes in the “Columns” section. Note that measurement accuracy values for camera positions and rotations should be changed to their best estimate that reflects any uncertainties (for example, from equipment, GNSS processing, and overall methodology) in the Reference Settings dialog (step 10). Even if not entered at this step, these values should be based off a quality assessment of the collection methodology (project based).

6. “Start import at row” tells Metashape at which row of data to start. Include header rows in the count of rows if there is header information in the camera location file.

7. Click “OK”.

8. A “Can’t find match for…” dialog appears if Metashape cannot match the loaded images’ label to the location file labels. This may be because the loaded imagery is a subsection of all the location information stored in the location file. Click “No to all” to prevent the nonmatches from being added to the project as “ghost” points. This dialog will not appear if each record in the location file has a corresponding image.

9. If the locations are added successfully, position information, including estimated errors, will be visible in the “Cameras” subsection on the right of the Reference panel and checkmarks will be shown next to the image names, indicating that their positions will be used in alignment and optimization.

10. The Reference Settings dialog can be accessed by clicking on the icon with the crossed hammer and wrench symbol. In this dialog, enter the applicable camera accuracies under the Measurement Accuracy >> “Camera accuracy (m)”. Note that x,y,z coordinate error can be independently specified, separated by a forward slash (/). For example, 0.1/0.15 indicates a 10-cm horizontal uncertainty and a 15-cm vertical uncertainty.

Aligning Images

The organized CCGs are now ready to proceed with alignment. At this stage, Metashape creates key points and tie points based on user-input parameters (fig. 9) and then produces the geometric reconstruction (position and orientation) of each image. If masks were applied to the imagery, they can be used at this stage to restrict where key points or tie points are created from the "apply masks to": dropdown of the Align Photos dialog. The selection of "None" will only mask areas when products are created. If this dropdown is grayed out, no masks were uploaded (fig. 9). Alignment outputs include a sparse point cloud and lens model. The alignment process is as follows (data input for steps are listed in brackets at the end of the appropriate step):

1. Access the Align Photos dialog from Workflow >> "Align Photos…"

2. Set "Accuracy" [High]. The accuracy should be set as high as the project allows or for the product quality desired. Lower accuracy settings will align faster and may be used if the project will be improved with the addition of GCPs and realigned. A "High" accuracy setting samples the images of the original size, "Medium" setting causes image downsampling by factor of 4 (a factor of 2 for each of the two dimensions), at "Low" accuracy source files are downsampled by a factor of 16, and the "Lowest" accuracy is downsampled by a factor of 84 (Agisoft LLC, 2020b). The "Highest" accuracy setting upsamples the image by a factor of 4 and substantially increases processing time.

3. "Generic preselection" [checked]. When this option is selected, images are preliminarily matched using the lower accuracy (downsampled) setting to find overlapping pairs of images before a second pass at the accuracy set in the dialog. This substantially speeds up the process but returns fewer overall tie points. The process does not do anything if the alignment level is already set at a low or lowest accuracy setting.

4. Select the "Reference preselection" checkbox and select "Source" from the dropdown. The option will be grayed out if no camera positions were added. When the "Reference preselection" option is available, three further options become available. The "Source" option allows Metashape to speed up alignment by matching overlapping pairs based on camera positions, the "Sequential" option will speed up alignment based on the camera label order, and the "Estimated" option will speed up alignment if this is a realignment and exterior orientations have already been calculated once. If working with oblique imagery specifically, see information on "Capture distance" when the "Source" option is in use in the Metashape user manual.

5. "Key point limit" [60,000]. A high limit or unlimited value will likely create more points, but the additional points may be lower quality and greatly increase processing time. The Metashape user manual suggests a limit of 40,000, but with the high-quality images used in the workflow documented in this report, a higher limit returned high-quality points.

6. "Tie point limit" [0] (unlimited). Tie points will be thinned after alignment. Too few tie points may cause holes in the point cloud, and too many tie points potentially cause memory problems (the result of a very large dataset); however, because tie points are thinned proportionately by quality in the optimization process, this is rarely an issue. Processing time increases with more tie points; when aligning very large datasets, adding a tie point limit may be practical.

7. "Adaptive camera model fitting" [unchecked]. When this option is unchecked, Metashape will only refine the default coefficients (f, cx, cy, k1, k2, k3, p1, and p2). This setting depends on the lens geometry of the image; see the "Alignment parameters" section of the Metashape user manual for details. The workflow used for this report does not enable this option to hold off refining coefficients that may not apply to the camera model (see the "Error Reduction—Optimization and Camera Calibration" section of this report).

8. Click "OK".

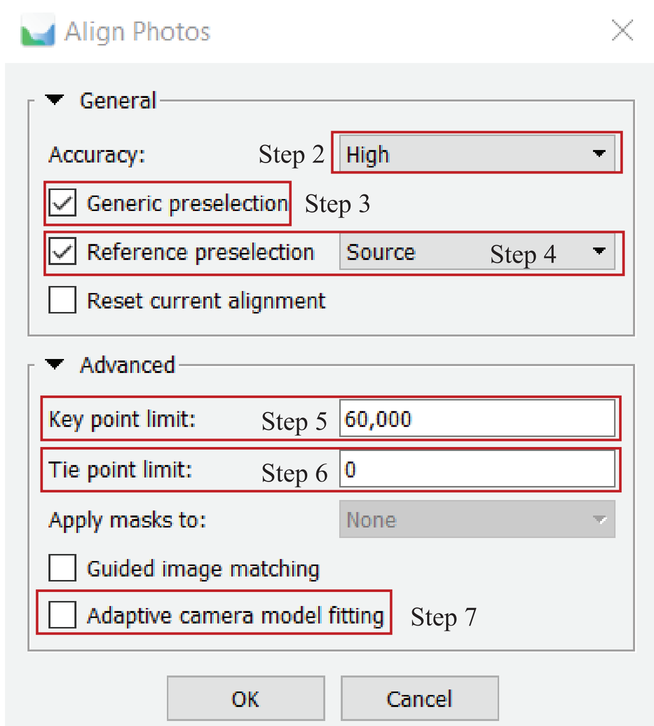

Figure 9. Screenshot of the Align Photos dialog in Agisoft Metashape Professional Edition, with the "General" and "Advanced" options expanded. Steps 2, 3, and 5–7 are outlined in red.

The alignment time depends on the size of the project. The Console panel will record the elapsed time of every stage in the alignment process, and if the "Write log to file" option is selected (see the "Metashape Preferences Dialog" section in this report), the times will also be recorded and saved there.

Adding Control Points and Check Points

Metashape recognizes "markers" as a way to georeference imagery. The data associated with the markers can be accessed from the Reference panel. Metashape refers to points that are used to reference the model as control points, and these can be an existing or purposely placed photograph-identifiable object or target with known location information. Scale bars can also be placed in a photographed scene and then digitized in Metashape. A check point is an unselected marker or control point that is not used by Metashape to optimize the camera alignment but can be used to check or validate the optimization results (Agisoft LLC, 2020b). The estimated values of check points can be compared with the known values in the "View errors" option in the Reference panel. After a control point is changed to a check point or the check point to a control point, click the "Update transform" button on the Reference panel toolbar to recalculate the estimated values. See the "What do the errors in the Reference pane mean?" section in the Metashape user manual for a full description on how the reference section points and markers are calculated (Agisoft LLC, 2020b).

Metashape has multiple ways to optimize the data it is given and can put different weights on marker and image positions based on the level of confidence the user places on their data. If image positions and markers are both available, the image (referring to the list of images under the "Cameras" subsection in the Reference panel) positions can be turned off, similar to a control point being made a check point. The decision is made on a case by case basis; for example, a small area with well-distributed markers and less accurate GNSS PPK locations may benefit from a camera model that leans more heavily on the markers.

Markers are represented in Metashape by three flag colors. On the Model and Ortho views, markers will always be blue flags and a selected marker will have a red dot. On individual images, flags can be blue, gray, or green (fig. 10).

- Blue flags in the Photos panel indicate that the markers were created by the user right-clicking on a location in the Model view and selecting "Add Marker". Metashape will put a blue flag on all the images at the marked location.
- Gray flags are not used in any calculation by Metashape; they represent where Metashape estimates a marker based on the coordinates provided or if done through automated marker detection.
- Green flags appear once the user manually moves or clicks on a gray or blue flag. This signals to Metashape that the user is confident in its location, at least to the user-input uncertainty values editable in the Reference Settings dialog.

Metashape also has an autodetect capability to locate user-placed markers placed in the scene. To include autodetection of markers, go to Tools >> Markers >> Detect Markers, and under the Parameters section the "Marker type" dropdown contains a variety of coded and noncoded markers that can be selected. Detecting markers with this method can improve accuracy compared with human-placed markers, although the results may underidentify or vastly overdetect the actual number of targets. Overdetection can be solved by sorting by the number of "Projections" in the Markers subpanel of the Reference panel and deleting any extraneous marker with a low number of projections. Alternatively, use the steps below to add and edit markers from an external file (fig. 10). Steps 1 to 6 are similar to the steps for importing camera locations (fig. 8).

1. Select the "Import Reference" icon; the symbol is a spreadsheet with a small yellow folder in the bottom corner.
2. Navigate to the directory where the marker data are stored. Click "Open" to launch the Import CSV dialog.
3. Select the appropriate coordinate system, making sure it corresponds to that of the camera positions if previously imported.
4. In the "Columns" section, enter the column order in the marker position file of the coordinates as well as any accuracy values. If accuracies of the marker positions are not in the file, uncheck the "Accuracy" boxes in the "Columns" section. Note that measurement accuracy values for marker positions should be changed to their best estimate in the Reference Settings dialog (step 9) even if not entered at this step; these values should be based off a quality assessment of the collection methodology (project based).
5. Check the "Delimiter" option if necessary.
6. "Start import at row" tells Metashape when to start reading data. Include header rows in the count of rows if there is header information in the external file.
7. Click "OK". If a "Can't find match…" dialog appears, select "Yes to all". Metashape is attempting to match every row of the external marker file with an existing marker, but all the markers are being created with this import. Metashape will automatically flag all the images that have a coincident estimated footprint with the marker.

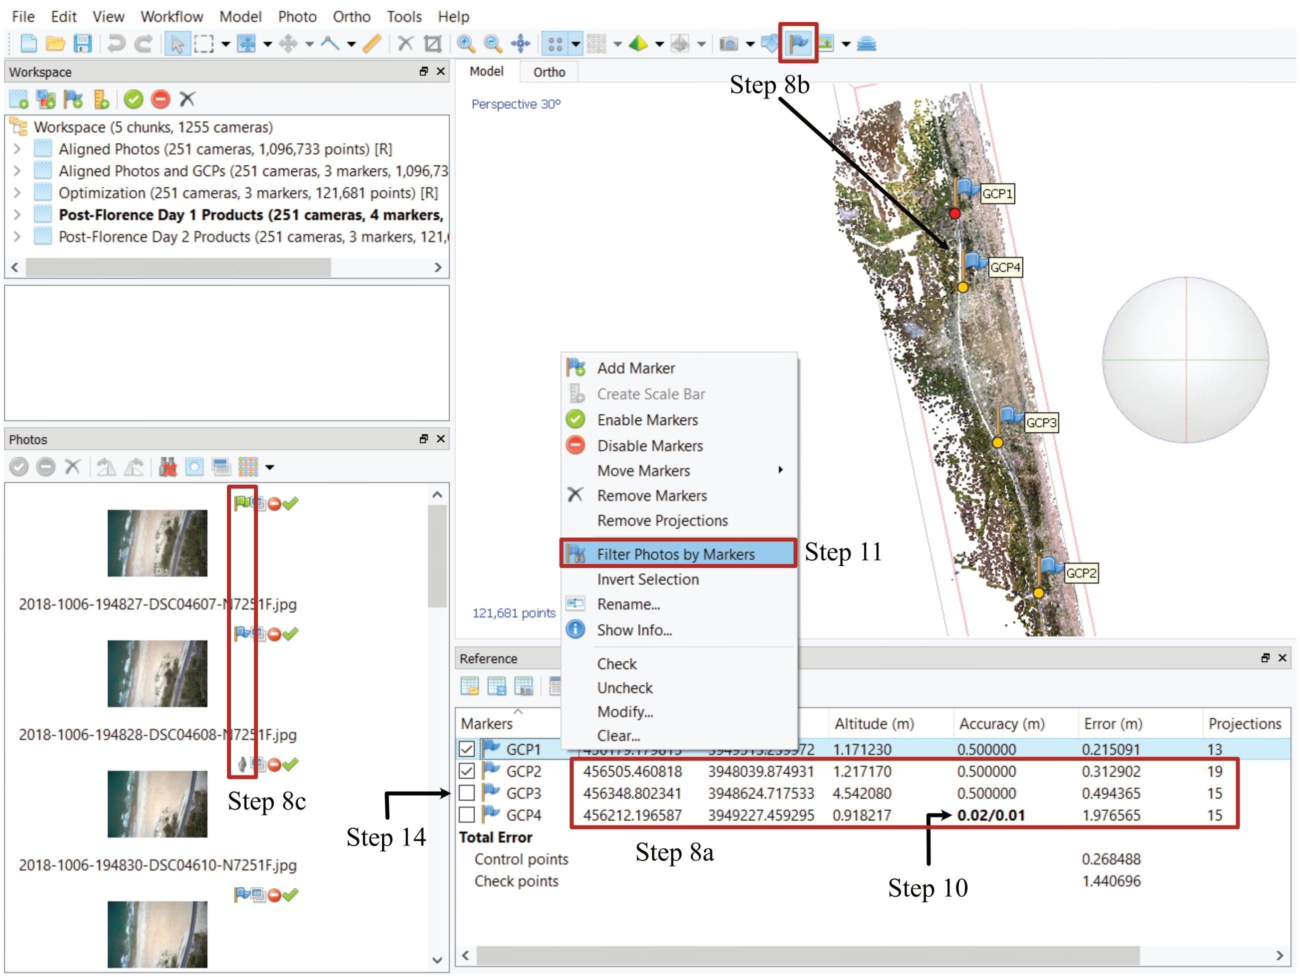

Figure 10. Screenshot of the Agisoft Metashape Professional Edition workspace showing steps for adding markers. Steps 8a–c, 10, 11, and 14 are indicated with red boxes and black arrows. The tall red box that represents step 8c also highlights the types of flags encountered in marker placement: a green flag, a blue flag, and a gray flag, respectively, from top to bottom.

8. If the markers are added successfully, check the following.

   a New information is visible in the "Marker" section in the Reference panel.

   b The Model view has blue flags; these will appear on the map if the locations are correct and if the model is already referenced to the same coordinate system. Make sure that the "Show Markers" icon, a blue flag, is turned on.

   c Flags will appear in the Photos panel.

9. After all markers have been added, navigate to the Reference Settings >> "Marker accuracy (m)" and "Marker accuracy (pix)". If the x,y,z coordinate accuracies are independent, separate the numbers with a forward slash (/). For example, 0.05/0.02 indicates a horizontal uncertainty of 5 cm, and a vertical certainty of 2 cm. Entering 0.05/0.07/0.02 indicates the uncertainty in the y direction is 7 cm. This will apply to all non-manually entered markers by default, but the values can be changed for each marker individually (see next step). In high-quality images, if the marker was an automatically detected target or a very carefully placed GCP, the marker accuracy is generally less than 0.5 pixel (the default value). The imagery for the workflow in this report supports a change to 0.1 pixel.

10. If markers have varying accuracies, double click on the accuracy cell of the specific marker in the Reference panel to manually edit the accuracy. Accuracies can also be separated by a forward slash (/) as in the previous step. Accuracies manually entered in this method will take on a bold typeface appearance and will not be affected if the Reference Settings dialog is changed.

11. Right-click on a marker and select "Filter Photos by Marker" to put every image with this specific marker into the Photos panel.

12. Double click on an image in the filtered Photos panel. Zoom in to the flag, which is likely gray because it was placed by Metashape and not confirmed by the user. If the flag has been placed correctly in the image, right-click on it and select "Place Marker" to turn the flag green. If it is not in the correct place, manually click and drag the marker. Note, there is no undo button and once a flag has been moved it cannot be put back exactly where it was.

13. If the placement of the marker is not acceptable or correct for an image, right-click on the marker in the image and select "Remove Marker" to remove it in just that image; the flag will revert to a gray flag. To remove the marker completely, right-click on it in the Reference panel and select "Remove Markers" to remove it from all images.

14. To turn a marker from a control point into a check point, uncheck its corresponding box in the "Markers" column of the Reference panel. It can be rechecked to be turned back into a control point.

15. To add a marker in an image, right-click on the location and choose "Add Marker".

Error Reduction—Optimization and Camera Calibration

Error reduction involves the selection and removal of low-quality tie points. The quality metrics used are based on the camera geometry of the images. The goal is to derive a sparse point cloud made up of only high-quality tie points and repeatedly optimize the camera model. This is the most subjective section of the workflow and testing how many points can be removed at each stage may be necessary to create a successful end product. Poor geometric relations between cameras result in points that are selected using the reconstruction uncertainty metric. Tie points to which Metashape has internally given poor match accuracies are selected using the projection accuracy criterion. Lastly, tie points that are the result of false matches are selected using the reprojection error criterion. The selection and elimination processes are iterative. The removal of poor tie points will improve the estimated internal and external orientation parameters but each time tie points are removed, the accuracies of the remaining tie points change, and the project requires reoptimization before continuing. Estimates of lens coefficient values after an optimization are viewed in the Camera Calibration dialog under the "Adjusted" tab (fig. 11). This process can be made more time efficient with the use of a script (for example, Logan and Ritchie, 2020).

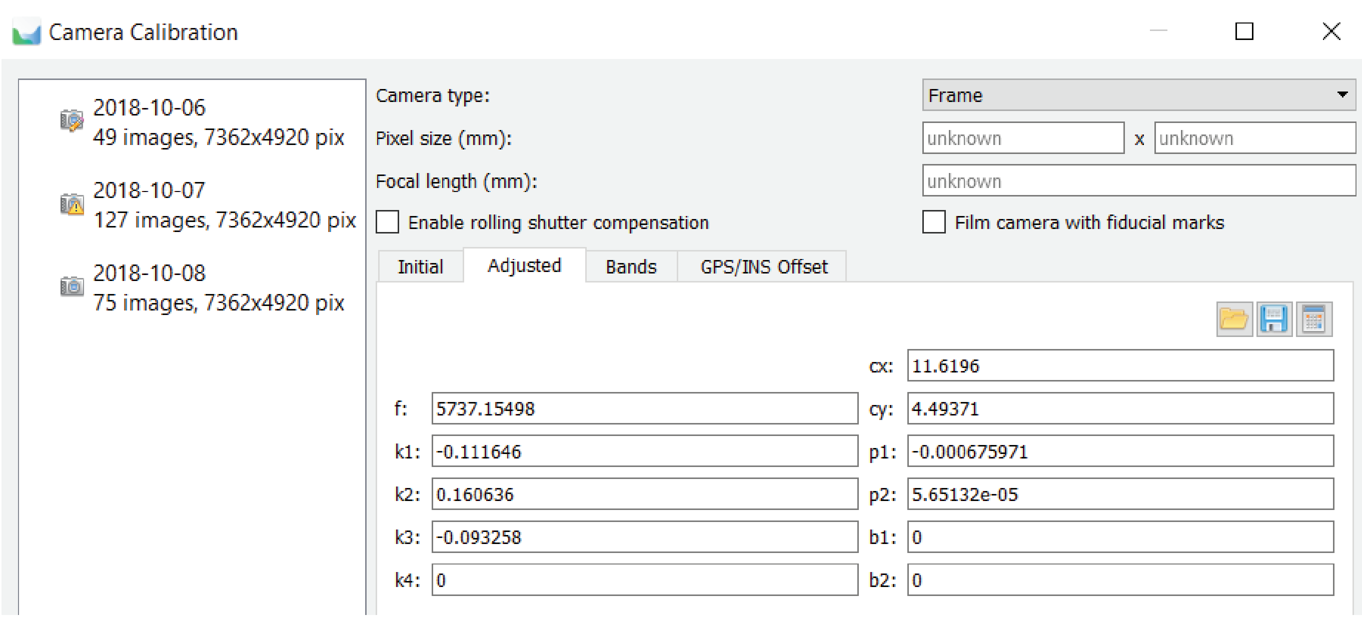

Figure 11. Screenshot of the Camera Calibration dialog in Agisoft Metashape Professional Edition, after optimization with the default parameters selected. The data for the "Adjusted" tab for the 2018-10-07 camera calibration group have been filled in for the lens coefficients selected in the optimization [f, k1, k2, k3, cx, cy, p1, p2]. The yellow alert exclamation icon on the 2018-10-07 camera symbol indicates that no camera information has been filled in; note the gray "unknown" text in the "Pixel size (mm)" and "Focal length (mm)" fields. The pencil icon on the camera symbol to the left of the 2018-10-06 camera calibration group indicates that a precalibration file was imported.

The error reduction phase relies on robust tie point, control point, and camera position error estimates. Accurate “Reference Settings” inputs (steps 9 and 10 of the “Adding Control Points and Check Points” section and step 5 of the “Adding Camera Positions” section) prevent misleading error statistics; incorrect error estimates lead to incorrect error models, and lens coefficients are very sensitive to the accuracies of the camera, marker, and tie points. The tie point accuracy parameter, measured in pixels, in the Reference Settings dialog corresponds to the normalized accuracy of tie point projection detected at a scale equal to 1 pixel; tie points detected at a different scale will have a proportional accuracy (Agisoft LLC, 2020b). A high-accuracy alignment supports leaving the default tie point projection at 1 pixel for initial optimization. A low-accuracy alignment with a value of 1 pixel would put more weight on the coordinate information of the cameras or markers rather than the tie points because the key point and tie points would be matched at a subsampled scale, and consequently their accuracy would be lower. Because tie points are not all created equal and scale individually (see the “Processing Parameters” section on page 63 of the Metashape user manual), the tie point accuracy is largely unconstrained, reducing, as a rule of thumb, with image-object resolving power. Conversely, marker accuracies should not be given unreasonably optimistic or unconstrained values in the Reference panel or Reference Settings dialog; they should reflect their known and reliable uncertainty values. In practice, if optimization after changing the tie point accuracy from 1 pixel to a lower value increases the marker “Error (m)” column to a level greater than what the “Accuracy (m)” column is set to, then the camera calibration is suspect. To avoid this issue, only change the tie point accuracy estimate in the reprojection error filtering steps; at this stage in the process, error should not be redistributed or extended to the control points.

A method to monitor the camera model refinement is to look at the Console panel. During the optimization process, the console displays a date and time stamp, followed by “adjusting”: and a string of “xxx” or “!” and two numbers (fig. 12). Each “x” in the string represents an iteration in the model optimization procedure required to solve for the lens coefficients, and each “!” represents a skipped iteration. A long string of “x”s indicates that individual solution sets are not converging quickly, or at all, in the model optimization. If the length of the string of “x”s does not begin to shorten in the process, then the solution may be diverging or insufficient lens coefficients have been modeled, failing to reach an internal trigger that would otherwise end the optimization process. If repeated optimization with the same coefficients does not reduce the number of “x”s, then the camera calibration model might have been overfitted and the chunk may need to be reprocessed from the beginning of postalignment, rather than trying to change parameters and continue reoptimizing. The final number after the “x”s is referred to as the “Sigma0” value by Agisoft and is equivalent to the photogrammetric adjustment quality indicator sigma naught (σ_0), which is the SEUW. The value is not found anywhere else in the GUI but can be accessed using the “chunk.meta” API (Agisoft LLC, 2020c). This value is the degree of deviation from an assumed accuracy or how closely the RMS reprojection errors match the estimated error values. In Metashape, the farther the SEUW is from 1, the poorer the tie point accuracy is estimated to be. The SEUW value is not part of the error reporting in the workflow because the weighting it relies on is not well documented and its default value is left at an unrealistic value until the final error reduction stages. The relative value is used to fine tune the tie point accuracy (see the “Error Reduction Step 3: Filtering by Reprojection Error” section of this report).

```
2020-03-10 12:56:40 adjusting: xxxxxxxxxxxxxxxxxxxx 0.213735 -> 0.211832
2020-03-10 12:56:43 point variance: 0.216155 threshold: 0.648464
2020-03-10 12:56:43 adding 0 points, 1101 far (0.648464 threshold), 0 inaccurate, 0 invisible, 0 weak
2020-03-10 12:56:43 adjusting: xxxxxxxxxxxxxxxxxxxx 0.128458 -> 0.126215
2020-03-10 12:56:47 point variance: 0.128826 threshold: 0.386478
2020-03-10 12:56:47 adding 0 points, 2180 far (0.386478 threshold), 0 inaccurate, 0 invisible, 0 weak
2020-03-10 12:56:47 adjusting: xxxxxxxxxxxxxxxxxxxx 0.114372 -> 0.109792
2020-03-10 12:56:49 point variance: 0.111348 threshold: 0.334045
2020-03-10 12:56:49 adding 75 points, 1348 far (0.334045 threshold), 0 inaccurate, 0 invisible, 0 weak
2020-03-10 12:56:49 adjusting: xxxxxxxxxxxx 0.105314 -> 0.105255
2020-03-10 12:56:51 point variance: 0.106349 threshold: 0.319048
2020-03-10 12:56:51 adding 40 points, 608 far (0.319048 threshold), 0 inaccurate, 0 invisible, 0 weak
2020-03-10 12:56:51 adjusting: xxxxxxxxxx 0.103644 -> 0.103631
2020-03-10 12:56:52 point variance: 0.104529 threshold: 0.313588
2020-03-10 12:56:52 adding 80 points, 300 far (0.313588 threshold), 0 inaccurate, 0 invisible, 0 weak
2020-03-10 12:56:52 optimized in 12.059 seconds
2020-03-10 12:56:52 f 5749.36, cx -3.92288, cy 6.71247, k1 -0.101684, k2 0.112598, k3 -0.00829077
2020-03-10 12:56:52 f 5767.3, cx 31.9815, cy 19.0754, k1 -0.103073, k2 0.116441, k3 -0.0136911
2020-03-10 12:56:52 finished sfm in 22.016 seconds
```

Figure 12. Screenshot of an example of the Console panel in Agisoft Metashape Professional Edition that highlights a series of adjustments after an optimization. Note the length of the string of ”x”s, symbolizing remaining iterations, getting shorter and the standard error of unit weight values after the string moving further from an accuracy value of 1.

A secondary method to monitor the camera model refinement is to observe the number of projections for each image (see “Cameras” subsection) in the Reference panel. Projections are the total number of valid tie points found on a given image (Agisoft LLC, 2020b). If the number of projections for any given image drops below 100 projections, then the image will not be used in making final products. When an image’s projections drop below the Metashape projection standards or reaches zero, a dialog will pop up and ask “Some cameras have insufficient number of projections and will be reset. Continue?” Select “Yes”; which will allow the image to be viewed in the Reference panel to determine why it was removed. For example, the image may be blurry, or it may have insufficient texture or overlap with other images to find key points. To view the number of points that have been removed throughout the error reduction process, right-click on the active chunk, then Show Info >> “Point Cloud” property >> “Points” value to list the current and the original numbers of tie points.

The last method covered to monitor the state of the camera model refinement is the chunk “Point Cloud” property that displays the Root Mean Square (RMS) reprojection errors, which are calculated using tie point projections. This value will be used in the following subsections of this report to assess the quality of the optimization. The tie points are located at different map scales to improve the robustness of the project, especially if there are blurred images or images with few distinct features or textures. Metashape uses information about the scale the tie points were matched at to weight the tie point reprojection errors (Agisoft LLC, 2020b); this helps the bundle adjustment but convolutes the meaning of the first reported RMS reprojection error because the scaling factors and weighting methods are not reported in Metashape. Of the two reported RMS reprojection errors in the chunk’s Show Info dialog, the weighted value is the first value in the units of key point scale, and the unweighted value, reported in pixels, is the second value in parentheses (fig. 13). A lower unweighted value, in general less than 0.3 pixel, is considered ideal. The process of filtering out tie points in the subsequent processes may result in the unweighted RMS reprojection error increasing if the camera model moves away from a certain solution because too many tie points have been removed.

There are two ways to visualize the uncertainty in the camera model after each optimization in the subsequent processes. View the numerical error in the Reference panel “Cameras” error subsection, or toggle the “View Variance” icon to display the standard deviation values (sigma or σ) for the image position coordinates and rotations angles. The square of the standard deviations are the variances, which Metashape calculates in a covariance matrix when the “Estimate tie point covariance” is selected on the Optimize Camera Alignment dialog. The matrix measures the uncertainty of the bundle adjustment transformations; the diagonal elements are the variances, and the off-diagonal elements are the covariances between the coefficients used to solve the camera model (Agisoft LLC, 2020b). The matrix is calculated for each tie point, and the direction and value of the largest error for the tie point estimated positions, determined by the covariance values, can be displayed in vector form using the Point Cloud Variance view (detailed in step 6 in the next paragraph; fig. 13).

To initiate the actual optimization process, select the Reference panel or Tools >> Optimize Cameras. The resulting Optimize Camera Alignment dialog allows for a selection of lens coefficients used in the optimization (fig. 13). Do not solve for all camera coefficients until subsequent optimizations reduce the error or the calibration check of the camera model results in a suggestion for additional corrections (see the “Camera Calibration Check” section of this report). Unchecking “Adaptive camera model fitting” is preferred; if checked, the camera model solutions are unpredictable because the software attempts to find the “best” combination of camera model coefficients that reduce the camera model error. Letting the software choose for you may result in lower error but is also an easy way to overfit the data or create an overly complex camera model that does not accurately reflect the equipment used, leading to underestimated or increased actual error. Note that the optimization process discards active depth maps, dense cloud, and mesh models in the chunk. These must be rebuilt to reflect any changes. Repeated optimizations may have diminishing returns or negative consequences to the camera model, and it is encouraged to start the error reduction steps over in a duplicated chunk from the last step rather than attempting to continue optimizing if the camera model appears to be diverging (using any of the methods outlined in this section as indicators of divergence). The following are the steps for a camera optimization after alignment (data input into steps for the workflow in this report are listed in brackets at the end of the appropriate step; fig 13):

1. Right-click on the active chunk and select “Duplicate…”
2. If no more images are expected to be added to the project, deselect “Key Points” from the Duplicate Chunk dialog. This operation will save file space. Click “OK”.
3. Verify that the combination of cameras and markers to be used is checked or unchecked in the Reference panel; all cameras will be used (checked) by default if their positions are added. Unchecked markers are used as check points. The decision to use camera positions or markers for georeferencing is based on the quality of the positions of both and the distribution of the markers.
4. Access the Optimization dialog by either the “Optimize Cameras” wand icon in the Reference panel or from Tools >> Optimize Cameras.
5. From the Optimize Camera Alignment dialog, keep the “General” coefficients checked [f, k1, k2, k3, cx, cy, p1, and p2].
6. In “Advanced” options, check “Estimate tie point covariance”. This will make it possible to visualize tie point error in the Model view.

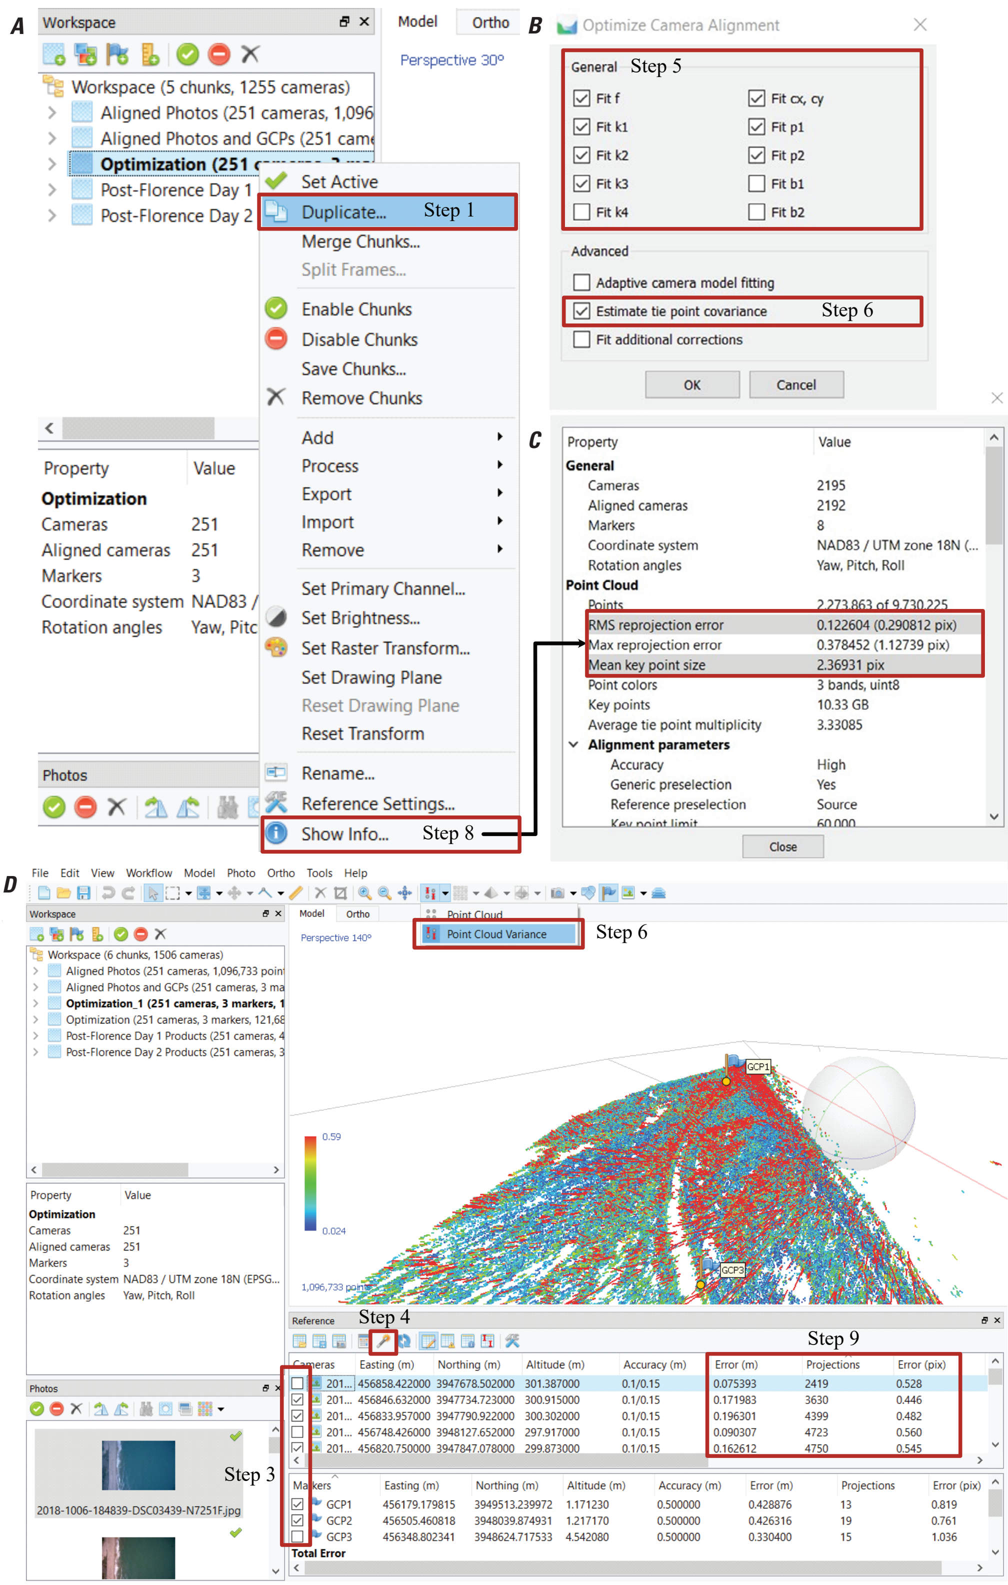

Figure 13. Screenshots of panels and dialogs in a workflow in Agisoft Metashape Professional Edition that depict the process and steps to perform the initial optimization or bundle adjustment and check the results; steps are outlined in red. *A*, Right-click a chunk in the Workspace panel and select "Duplicate…" (fig. 3) followed by "Show Info"; *B*, The Optimization Camera Alignment dialog with "General" camera coefficients selected; *C*, The Show Info dialog, with the two chunk properties outlined in red; and *D*, The Optimization wand tool in the Reference panel, the "Projections" and "Error" columns that appear for each camera after optimization, and the "Point Cloud Variance" option in the Model view.

7. Click “OK”. The optimization will be performed.

8. Right-click on a chunk, then Show Info… >> Property >> Point Cloud; note the unweighted “RMS reprojection error (pix)”.

9. Observe the error and number of projections of the cameras in the Reference panel.

Note that the affinity and nonorthogonality or skew coefficients (b1 and b2) are not selected by default; the user should initially suppress these coefficients and only include them if an inflated RMS reprojection error value indicates the possible presence of in-plane distortion and their inclusion in the camera model substantially affects the RMS reprojection error (Fraser, 2013).

Error Reduction Step 1: Filtering by Reconstruction Uncertainty

The first phase of error reduction is to remove points that are a result of poor camera geometry. Reconstruction uncertainty is a numerical representation of the uncertainty in the position of a tie point based on the geometric relationship of the cameras from which that point was projected or triangulated, considering geometry and redundancy. Reconstruction uncertainty can also be thought of as the ratio between the largest and smallest semiaxis of the error ellipse created when triangulating 3D point coordinates between two images (Agisoft LLC, 2020b). Points constructed from image locations that are too close to each other have a low base-to-height ratio and high uncertainty. Removal of these points does not affect the accuracy of optimization but reduces the noise in the point cloud and prevents points with a large uncertainty in the z axis from influencing other points with good geometry or being incorrectly removed in the reprojection error step. A reconstruction uncertainty of 10, which can be selected with the gradual selection filter and set to this value or level, is roughly equivalent to a good base-to-height ratio of 1:2.3 (parallax angle of about 23 degrees [°]), whereas 15 is roughly equivalent to a marginally acceptable base-to-height ratio of 1:5.5 (parallax angle of about 10°). Previous guidance (U.S Geological Survey, 2017) directs the reconstruction uncertainty selection procedure to be repeated two times to reduce the reconstruction uncertainty toward 10 without having to delete more than 50 percent of the tie points each time. If a reconstruction uncertainty of 10 is reached in the first selection attempt and less than 50 percent of the tie points are selected, a single optimization after deleting points is sufficient. Repeated selections have diminishing returns and the added optimizations may overfit the camera model before more poor-quality tie points are removed. Be cautious repeating the following steps (fig. 14); the use of multiple chunks to test the best selection and optimization combinations is encouraged (data input into steps for the workflow in this report are listed in brackets at the end of the appropriate step):

1. Model >> Gradual Selection.

2. In the Criterion dialog, select “Reconstruction uncertainty”.

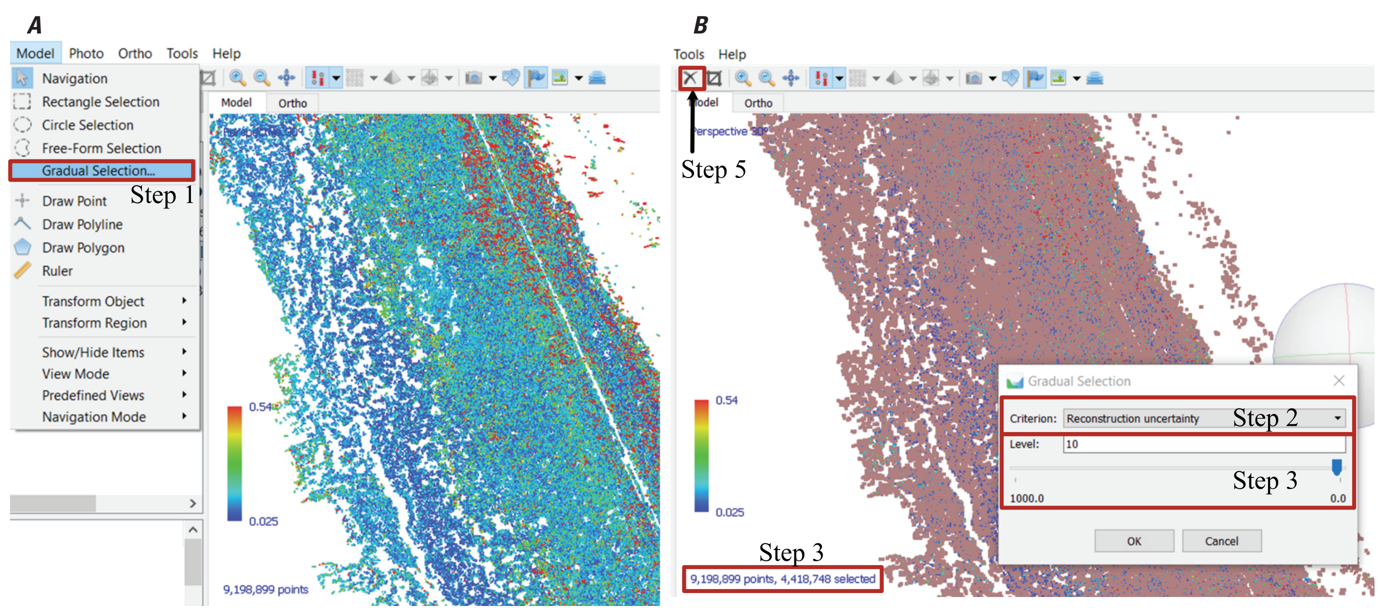

Figure 14. Screenshot of steps in Agisoft Metashape Professional Edition, indicated with red boxes and a black arrow, showing how to filter points based on reconstruction uncertainty. *A*, Select the Gradual Selection tool to remove tie points; and *B*, Adjust the reconstruction uncertainty criterion in the Gradual Selection dialog. The point cloud in the background shows the points selected in dark pink, on the basis of the criteria and the number of points shown on the screen (step 3).

3. Slide the level bar or directly enter the targeted number [10]. If more than about 50 percent of the tie points are selected, increase the value in 0.1 increments until less than 50 percent of the points are selected.
4. Click “OK”.
5. Delete selection.
6. “Optimize Cameras” with default coefficients [f, k1, k2, k3, cx, cy, p1, p2] and “Estimate tie point co-variance” selected (see steps 4 to 6 in the “Error Reduction—Optimization and Camera Calibration” section of this report).
7. Monitor lowest camera projections and unweighted RMS reprojection error (in pixels; see steps 8 and 9 in the “Error Reduction—Optimization and Camera Calibration” section of this report).

Error Reduction Step 2: Filtering by Projection accuracy

The second error reduction phase removes points based on projection accuracy, which is a measure of the “Mean key point size”; the value is accessed from the chunk information dialog (fig. 13*C*). Key point size (in pixels) is the standard deviation value (in σ) of the Gaussian blur at the scale at which the key point was found; lower standard deviation values are more precisely located in space (Agisoft LLC, 2020b). Thus, the smaller the mean key point value, the smaller the standard deviation and the more precisely located the key point is in the image. Projection accuracy is essentially a representation of the precision with which the tie point can be known given the size of the key points that intersect to create it. Metashape saves an internal accuracy and scale value for each tie point as part of the correlation process. The highest accuracy points are assigned to level 1 and are weighted based on the relative size of the pixels. A tie point assigned to level 2 has twice as much projection inaccuracy as level 1. Not all projects can tolerate gradual selection and removing points at a level of 2 to 3, particularly if the images have undergone compression or are of a lower quality from noise or blur. A gradual selection level of 5 or 6 may be the best that can be obtained. Previous guidance (U.S Geological Survey, 2017) directs the projection accuracy selection procedure is to be repeated without having to delete more than 50 percent of the points each time until level 2 is reached, and at this level there are few to no points selected. Here we clarify that after level 3 is reached, repeated selections have diminishing returns and the added optimizations may overfit the camera. If the project can support an initial projection accuracy selection of level 3 without selecting more than 50 percent of the tie points, a single optimization after deleting the points is sufficient. The camera model may benefit from repeating the following steps to get closer to level 2 (fig. 15), but there is a balance to be

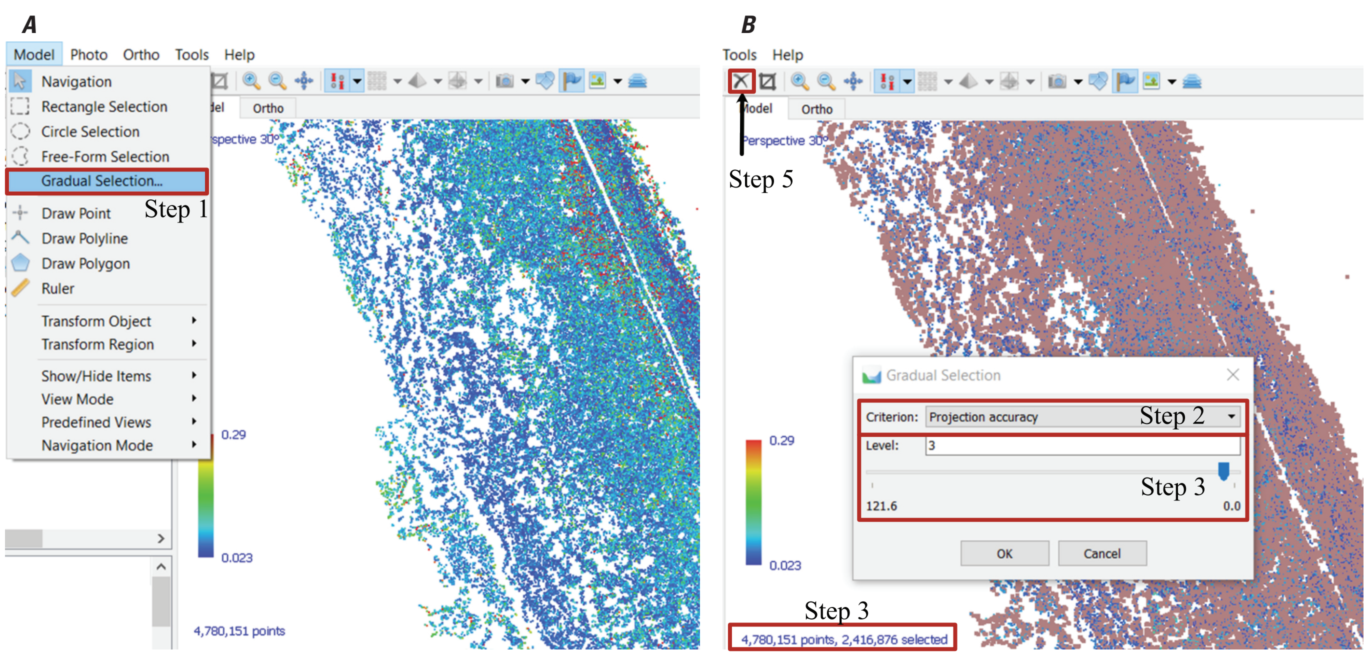

Figure 15. Screenshot of steps in Agisoft Metashape Professional Edition, indicated with red boxes and a black arrow, showing *A*, The Gradual Selection tool to remove tie points; and *B*, Adjust the projection accuracy criterion in the Gradual Selection dialog. The point cloud in the background shows the points selected in dark pink based on the criteria and the number of points shown on the screen (step 3).

maintained to prevent deleting too many tie points (data input into steps for the workflow in this report are listed in brackets at the end of the appropriate step):

1. Model >> Gradual Selection.
2. In the Criterion dialog, select “Projection accuracy”.
3. Slide the level bar or directly enter the targeted number [3]. If more than about 50 percent of the tie points are selected, increase the value in 0.1 increments until less than 50 percent of the points are selected.
4. Click “OK”.
5. Delete selection.
6. “Optimize Cameras” with default coefficients [f, k1, k2, k3, cx, cy, p1, p2] and “Estimate tie-point co-variance” selected (see steps 4–6 in the “Error Reduction—Optimization and Camera Calibration” section of this report).
7. Monitor lowest camera projections and unweighted RMS reprojection error (in pixels; see steps 8 and 9 in the “Error Reduction—Optimization and Camera Calibration” section of this report).

Error Reduction Step 3: Filtering by Reprojection Error

The final error reduction phase is removing points based on reprojection error. This is a measure of the error between a 3D point’s original location on the image and the location of the point when it is projected back to each image used to estimate its position (Agisoft LLC, 2020b). Error values are normalized based on key point size (in pixels). High reprojection error usually indicates poor localization accuracy of the corresponding point projections at the point matching step (Agisoft LLC, 2020b). Reprojection error can be reduced by iteratively selecting and deleting points, then optimizing until the unweighted RMS reprojection error (in pixels) is between 0.13 and 0.18 (see the “Error Reduction—Optimization and Camera Calibration” section of this report; fig 13) with additional optimization coefficients checked (that is, flagged as required) and with the use of the “Fit additional corrections”. In Metashape version 1.6.x, the scaled (weighted) RMS reprojection error is calculated such that not all points selected by targeting a specific reprojection error level will remove the same level of estimated error. When a value of 0.13 to 0.18 unweighted RMS reprojection error (in pixels) is reached, there is 95 percent confidence that remaining points meet the estimated error. Most projects should be sufficiently optimized at a reprojection error of 0.3 pixels, and there should be no noticeable differences in products with additional optimization. However, high-quality images may benefit from more rigorous error reduction to the lower levels and produce a better camera calibration. At this step, the tie point accuracy (in pixels) could also be revisited (step 8), while constantly monitoring the camera and marker total errors to assist in the decision to stop the error reduction and optimization process. The following steps are the process to reduce error by filtering points based on reprojection error (data input into steps for the workflow in this report are listed in brackets at the end of the appropriate step; fig. 16):

1. Reference panel >> Reference Settings >> Image Coordinates Accuracy >> Tie point accuracy [1] (default).
2. Go to Model >> Gradual Selection.
3. In Criterion dialog, select “Reprojection error”.
4. Slide the bar until about 10 percent of the tie points are selected or the level is 0.3.
5. Click “OK”.
6. Delete selection.
7. “Optimize Cameras” with one of the following options.
   - default coefficients selected [f, k1, k2, k3, cx, cy, p1, p2];
   - all coefficients selected [f, k1, k2, k3, k4, cx, cy, p1, p2, b1, b2] and “Fit additional corrections”; or
   - some other combination of coefficients that fits the data and camera.
8. Repeat steps 2 to 7 until the RMS reprojection error level is at 0.3, and zero or close to zero tie points are selected. Be cautious of overoptimizing in this step, as there are diminishing returns. Monitor the total camera and marker errors; if they ever start to exceed the “Accuracy (m)” column or value set in the Reference Settings, consider stopping the process.
9. Check the unweighted RMS reprojection error (fig. 13). If it is below 0.18 pixel, stop (the process ends); otherwise, continue these steps. If the SEUW has deviated from 1, continue to step 10 to reduce the deviation. If the SEUW has not deviated from 1 during the optimizations, or you do not wish to improve the deviation, go to step 13.
10. “Tie point accuracy” can be changed from the default in step 1. Change the estimate to a smaller value [0.3–0.1].
11. “Optimize Cameras” with one of the following options:
    - default coefficients selected [f, k1, k2, k3, cx, cy, p1, p2];
    - all coefficients selected [f, k1, k2, k3, k4, cx, cy, p1, p2, b1, b2] and “Fit additional corrections”; or
    - some other combination of coefficients that fits the data and camera.

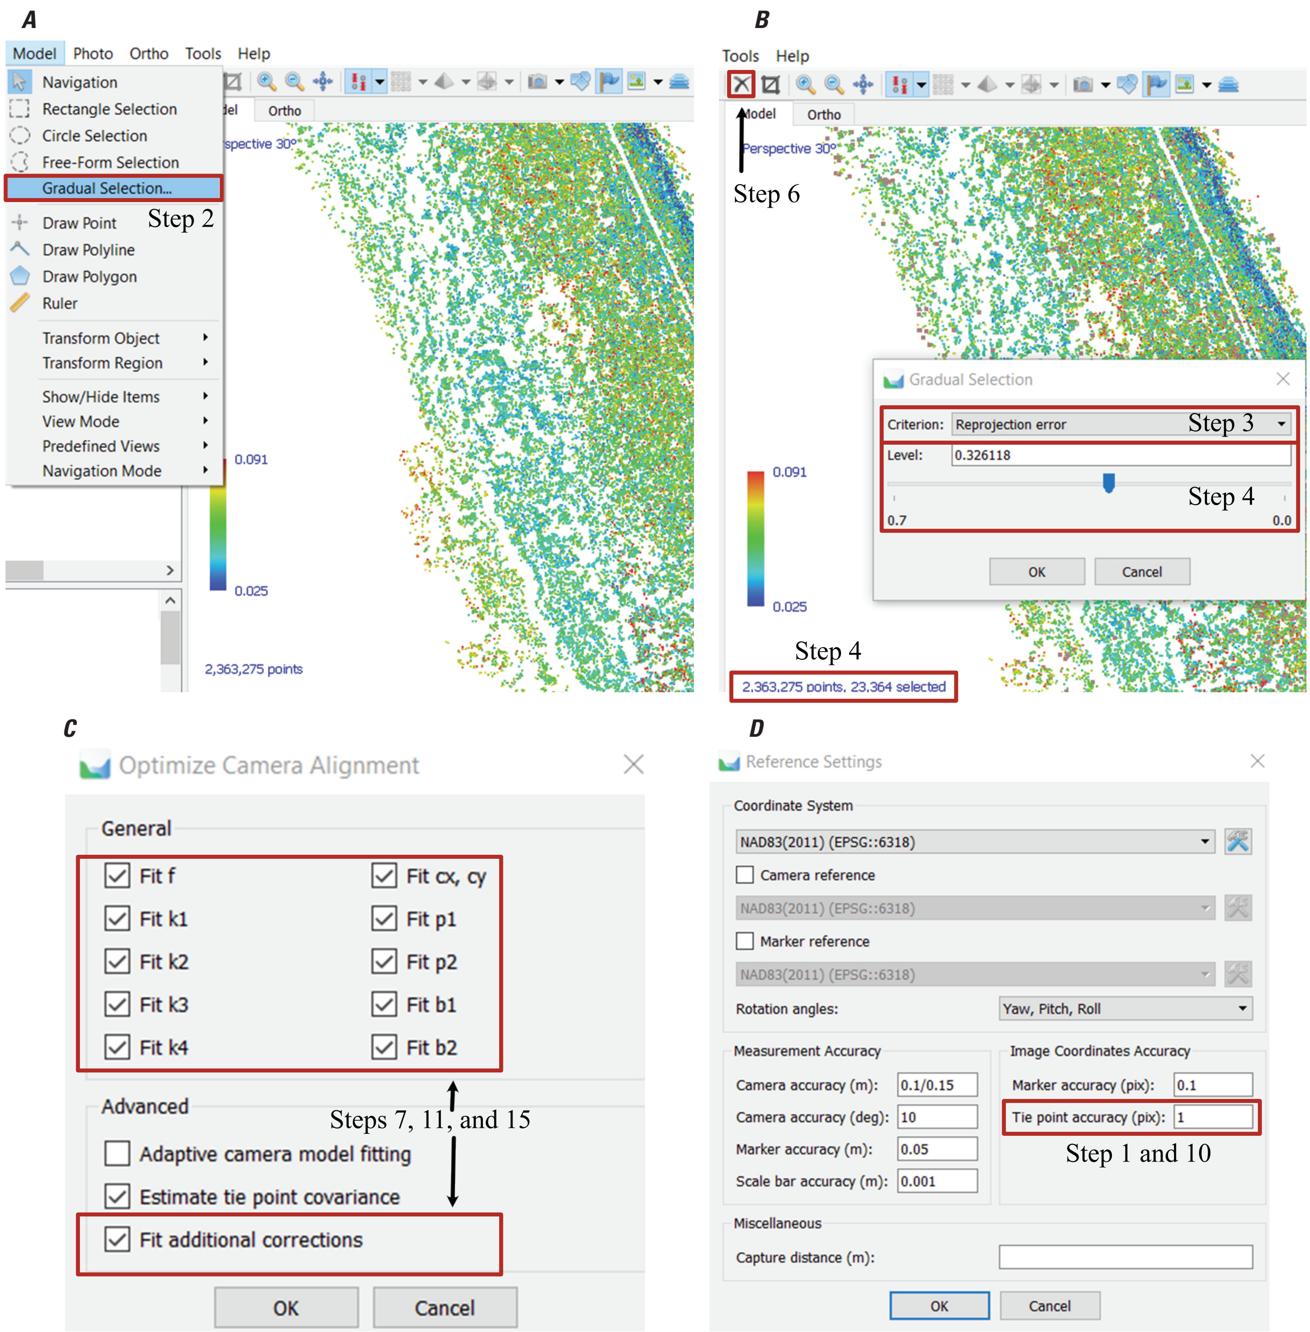

Figure 16. Screenshot of steps in Agisoft Metashape Professional Edition, indicated with red boxes and a black arrow, showing how to reduce error by filtering points based on reprojection error. *A*, Use the Gradual Selection tool to remove tie points; *B*, Adjust the reconstruction error criterion (rather than typing in a level, observe the number of tie points selected in the bottom left corner); *C*, Optimize Camera Alignment dialog where additional camera coefficients may be selected; and *D*, Modifying the "Tie point accuracy (pix)" parameter in the Reference Settings dialog.

12. Check the SEUW, total marker error (in meters), and total camera error (in meters). The total error values are the bolded values at the bottom of the "Cameras" and "Markers" subsection in the Reference panel. Repeat steps 10 and 11 until the total camera error and marker errors do not exceed the "Accuracy (m)" column or value set in the Reference Settings dialog and the SEUW is closer to 1.

13. Make a decision on the project on the basis of the answers to the following questions. If the answers are mostly yes, continue to step 14; otherwise, stop (processing is completed).

    - Are the lowest numbers in the "Projections" column in the Refence panel "Cameras" subsection still in the hundreds, disregarding images that may have been removed in earlier optimizations due to blurriness or subject matter?
    - Was the project able to set the initial level of reconstruction uncertainty to 10 and the projection accuracy to 3 without selecting more than the suggested number of tie points?
    - Did the number of "x"s decrease in the Console panel during the optimization process (fig. 12)?
    - Was there no evidence of overfitting (see the "Camera Calibration Check" section of this report)?
    - Are there still more than 15–20 percent of the original number of tie points? View this in the chunk information dialog.

14. Go to Model >> Gradual Selection >> Reprojection error.

    a Slide the level bar until 10 percent of tie points are selected.

    b Click "OK".

    c Delete points.

15. "Optimize Cameras" with one of the following options:

    - default coefficients selected [f, k1, k2, k3, cx, cy, p1, p2];
    - all coefficients selected [f, k1, k2, k3, k4, cx, cy, p1, p2, b1, b2] and "Fit additional corrections"; or
    - some other combination of coefficients that fits the data and camera.

16. Check the unweighted RMS reprojection error (in pixels).

17. Repeat steps 14 to 16 until the unweighted RMS reprojection error (in pixels) is between 0.13 and 0.18. If the value increases consistently at any point, the project may not support this level of error reduction the chunk may need to be reoptimized with a different set of parameters. To reoptimize, return to step 1 in the "Error Reduction—Optimization and Camera Calibration" section of this report.

For guidance on coefficient selection during optimization, see the discussion in the "Error Reduction—Optimization and Camera Calibration" section of this report or refer to the Metashape user manual for reasons to select or deselect coefficients and reduce the chance of overfitting the camera model. Note that the unweighted RMS reprojection error (in pixels) may not fall below 0.18 if only the default coefficients are selected or the images in the project do not support this level.

Camera Calibration Check

Once the reconstruction uncertainty, projection accuracy, and reprojection error processes are finalized, it is good practice to check the Camera Calibration dialog and look at the adjustments and distortion plots (fig. 17). Note that the plots are averaged across camera calibration groups so this may have to be repeated for each group.

1. Tools >> Camera Calibration.

2. Right-click on a CCG, then select Distortion Plot.

3. Under the "Distortion" tab, note the following:

    - The "Total" radio button displays the modeled lens distortion based on the adjusted calibration coefficient values. "Total" includes the radial and tangential ("Decentering") distortion, note that the effect of the radial component of distortion is normally much greater on the overall model.
    - The "Residuals" radio button displays a plot of the quality of the optimization and how adequately the camera geometry is described by the applied model (Agisoft LLC, 2020b). A pattern to the residual error indicates some sort of systematic issue with the project. Ideally, the majority of the residual error vectors should be less than 0.3 pixels long.

4. The "Correlation" tab provides useful information about the standard deviation of the camera calibration coefficients used in the interior orientation camera model, as well as the covariance (the degree of correlation between different pairs of coefficients in the corresponding row or column for the elements off of the diagonal). If the standard deviation is approaching the same value (or same magnitude) as the coefficient, that is a good indication of overfitting and may indicate that the coefficient may not be contributing substantially to the accuracy of the camera model.

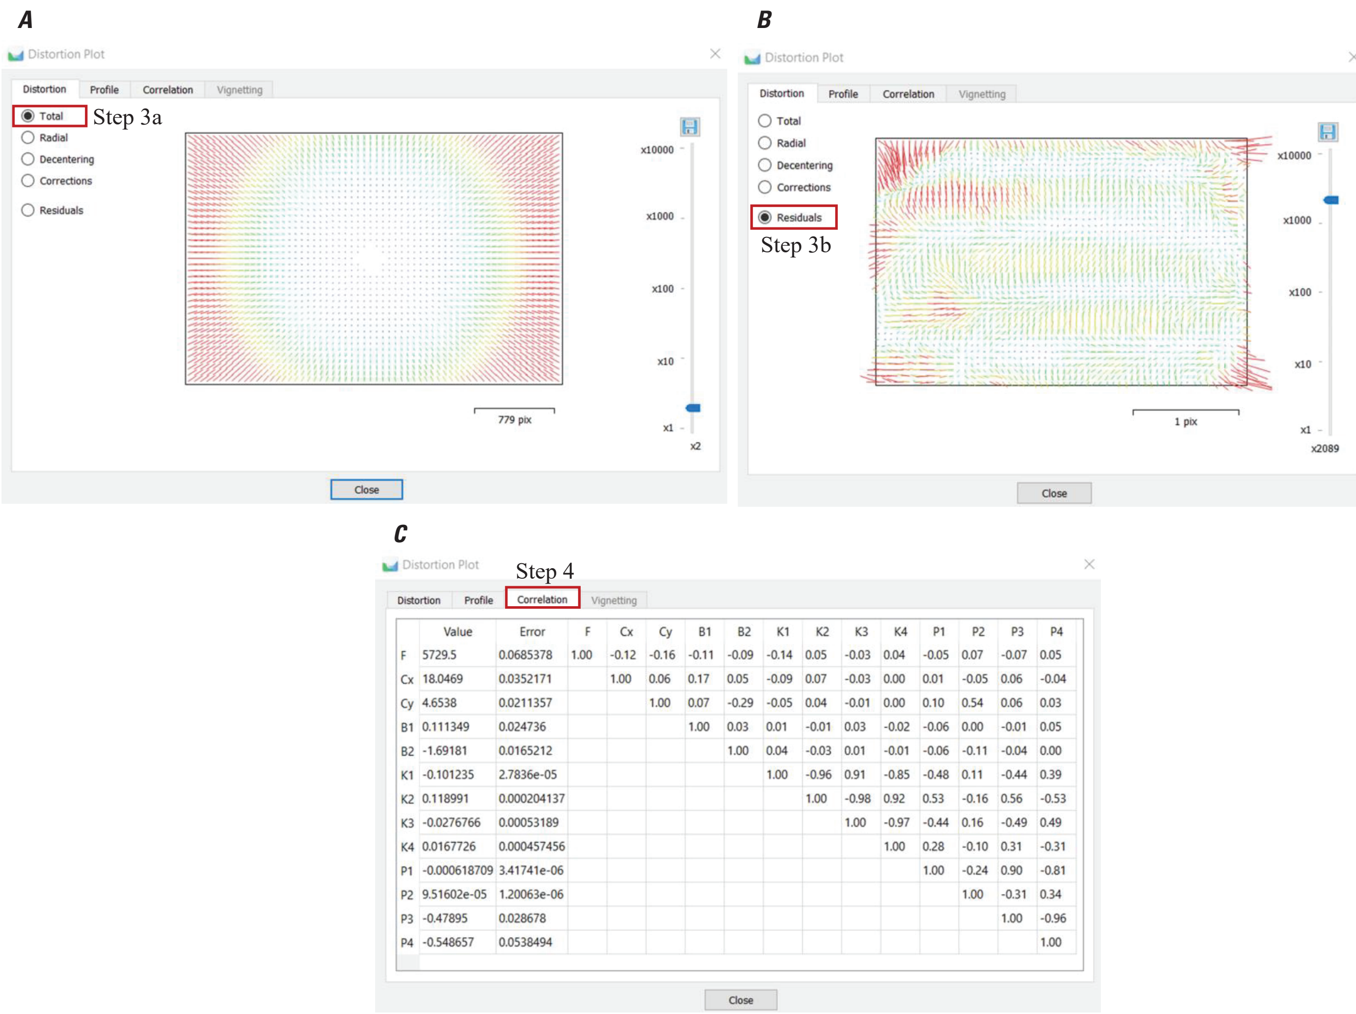

	Value	Error	F	Cx	Cy	B1	B2	K1	K2	K3	K4	P1	P2	P3	P4
F	5729.5	0.0685378	1.00	-0.12	-0.16	-0.11	-0.09	-0.14	0.05	-0.03	0.04	-0.05	0.07	-0.07	0.05
Cx	18.0469	0.0352171		1.00	0.06	0.17	0.05	-0.09	0.07	-0.03	0.00	0.01	-0.05	0.06	-0.04
Cy	4.6538	0.0211357			1.00	0.07	-0.29	-0.05	0.04	-0.01	0.00	0.10	0.54	0.06	0.03
B1	0.111349	0.024736				1.00	0.03	0.01	-0.01	0.03	-0.02	-0.06	0.00	-0.01	0.05
B2	-1.69181	0.0165212					1.00	0.04	-0.03	0.01	-0.01	-0.06	-0.11	-0.04	0.00
K1	-0.101235	2.7836e-05						1.00	-0.96	0.91	-0.85	-0.48	0.11	-0.44	0.39
K2	0.118991	0.000204137							1.00	-0.98	0.92	0.53	-0.16	0.56	-0.53
K3	-0.0276766	0.00053189								1.00	-0.97	-0.44	0.16	-0.49	0.49
K4	0.0167726	0.000457456									1.00	0.28	-0.10	0.31	-0.31
P1	-0.000618709	3.41741e-06										1.00	-0.24	0.90	-0.81
P2	9.51602e-05	1.20063e-06											1.00	-0.31	0.34
P3	-0.47895	0.028678												1.00	-0.96
P4	-0.548657	0.0538494													1.00

Figure 17. Screenshots of the Distortion Plot dialog in Agisoft Metashape Professional Edition that visualizes and reports error of the camera model. Options discussed are outlined with red boxes. *A*, The "Distortion" tab "Total" plot; *B*, The "Distortion" tab "Residuals" plot; and *C*, The "Correlation" tab matrix.

Building Products

Multiple products can be created after the point cloud has been edited and the camera model has been optimized and calibrated to its practical limit. Dense point clouds, mesh models, orthomosaics, and DEMs are common products for coastal change projects such as the USGS hurricane relief and recovery effort detailed in this report. Mesh building is not included in this report because the Metashape user manual covers the topic extensively and meshes are not currently in our workflow. Note that depth maps, dense point clouds, and mesh models are discarded from the active chunk if the chunk is optimized again, and products will need to be rebuilt.

Before creating products, it can be useful to edit the region, the 3D frame seen in the Model view, to refine the boundaries of the products. Only the area within the region will be processed or exported. The default region that Metashape creates is arbitrary in size and orientation. When editing, the vertices of the framework can be moved with the mouse or rotated or moved as a whole, using the options from the Move Region action (fig. 18).

If creating products from a 4D project, it is important to remember that the 4D part of the project is for alignment, not for generating products. After image alignment, 4D projects are returned to the 3D workflow in order to create products. To create the 4D products, it is necessary to disable or delete "Camera Groups" from different dates or collections and only produce products with one collection enabled at a time; this is the equivalent of what would be processed in a 3D project (fig. 19). This can be done either by serially disabling different collections and producing products in a single chunk, or by duplicating chunks and disabling or deleting Camera Groups for all but one collection. The latter is more useful if products are made in a batch process (see the section "Batch Processing"). The following are the steps to prepare a 4D project for product creation (data input into steps for the workflow in this report are listed in brackets at the end of the appropriate step):

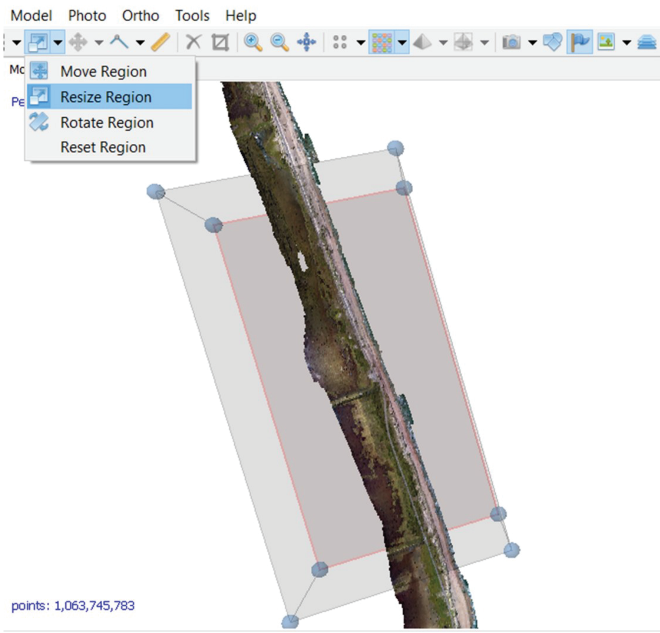

Figure 18. Screenshot of the Agisoft Metashape Professional Edition workspace during region editing; the three-dimensional framework as depicted by the eight blue spheres can be rotated, moved, and resized to determine the export dimensions.

1. Duplicate the latest optimized chunk: right-click, then select Duplicate.
2. Rename the duplicated chunk(s) to reflect the reason they were split (for example, 4D_Products if following step 3a, or 4D_2018-10-06_Products and 4D_2018-10-07_Products if following step 3b).
3. Disable or delete the camera groups not relevant to the chunk for processing. This can be done in the Workspace panel by right-clicking on the targeted camera group, then selecting:
   a the "Disable Cameras" option. Disabled camera groups will not be used in product creation and are identified in the active chunk by the red disabled symbol on the folder and in the individual images in the Reference and Photos panels or
   b "Remove Cameras" option.
4. Repeat steps 1–3 for each chunk that is going to be producing products separately. Or, repeat step 3a and change which camera group(s) is disabled or enabled if producing products in the same chunk; right-click on the first and subsequent products to deselect the "Set as default" option to prevent the product from being overwritten. The product name should become italicized if successful.

The initial step in producing high-resolution products creates depth maps, also known as disparity maps, that can be viewed for each individual image when selected from the Photos panel (fig. 20). The depth maps are created based on intersecting points from discrete camera positions. The relation between depth and disparity is based on camera baseline separation.

Depth maps are generated only for cameras with a certain number of projections and is controlled by a quality setting and a filtering mode. The quality setting controls the resolution relative to the source image, with ultrahigh (the highest setting) being full resolution, and each lower setting providing a resolution downsampled by a factor of 4 (2 times by each side) of the next higher resolution. Filtering modes control the amount of noise filtering performed, and include "none", "mild", "moderate", and "aggressive". The filter level controls the maximum size of connected components discarded by the filter (Agisoft LLC, 2020b). The "mild" setting preserves the finest detail in the landscape while also removing excessive noise. When creating products that depend on depth maps of the same filter level, selecting "Reuse depth maps" will shorten the processing time.

Build Dense Point Cloud

The dense point cloud uses the estimated camera positions generated during sparse point cloud based (tie point) matching and the depth map for each camera. Metashape produces very dense point clouds, which are of equal to or of a greater density than airborne lidar point clouds. The dense point cloud can be classified into vegetation, manmade objects, roads, and so on, by manual selection, color selection, or other means, as detailed in the Metashape user manual. Products can be created that use the entire dense point cloud or only certain point classes, allowing the project to be directly comparable to datasets such as ground-classified lidar points. If comparing directly with lidar point clouds, it can be useful to limit lidar points to first-return data. The settings for "Dense Point Cloud" creation are as follows (data input into steps for the workflow in this report are listed in brackets at the end of the appropriate step; fig. 21):

- Select from "Quality" dropdown [High]. The quality is based on the products the project is expected to produce.
- Select from "Depth filtering" dropdown [Mild]. This parameter defines the level of detail or texture kept in the scene. Images with poor texture or fine detail should avoid both aggressive filtering (which will remove the detail) and "Disabled" [none] (which will produce a very noisy product). Experimentation is likely needed to find the most suitable setting for a project.
- "Reuse depth maps" [checked]. If depth maps are available in the chunk, enabling this feature will save processing time; if depth maps are not available, deselect the option to create new depth maps. If the option is grayed out, depth maps of the "Quality" selected will be created.

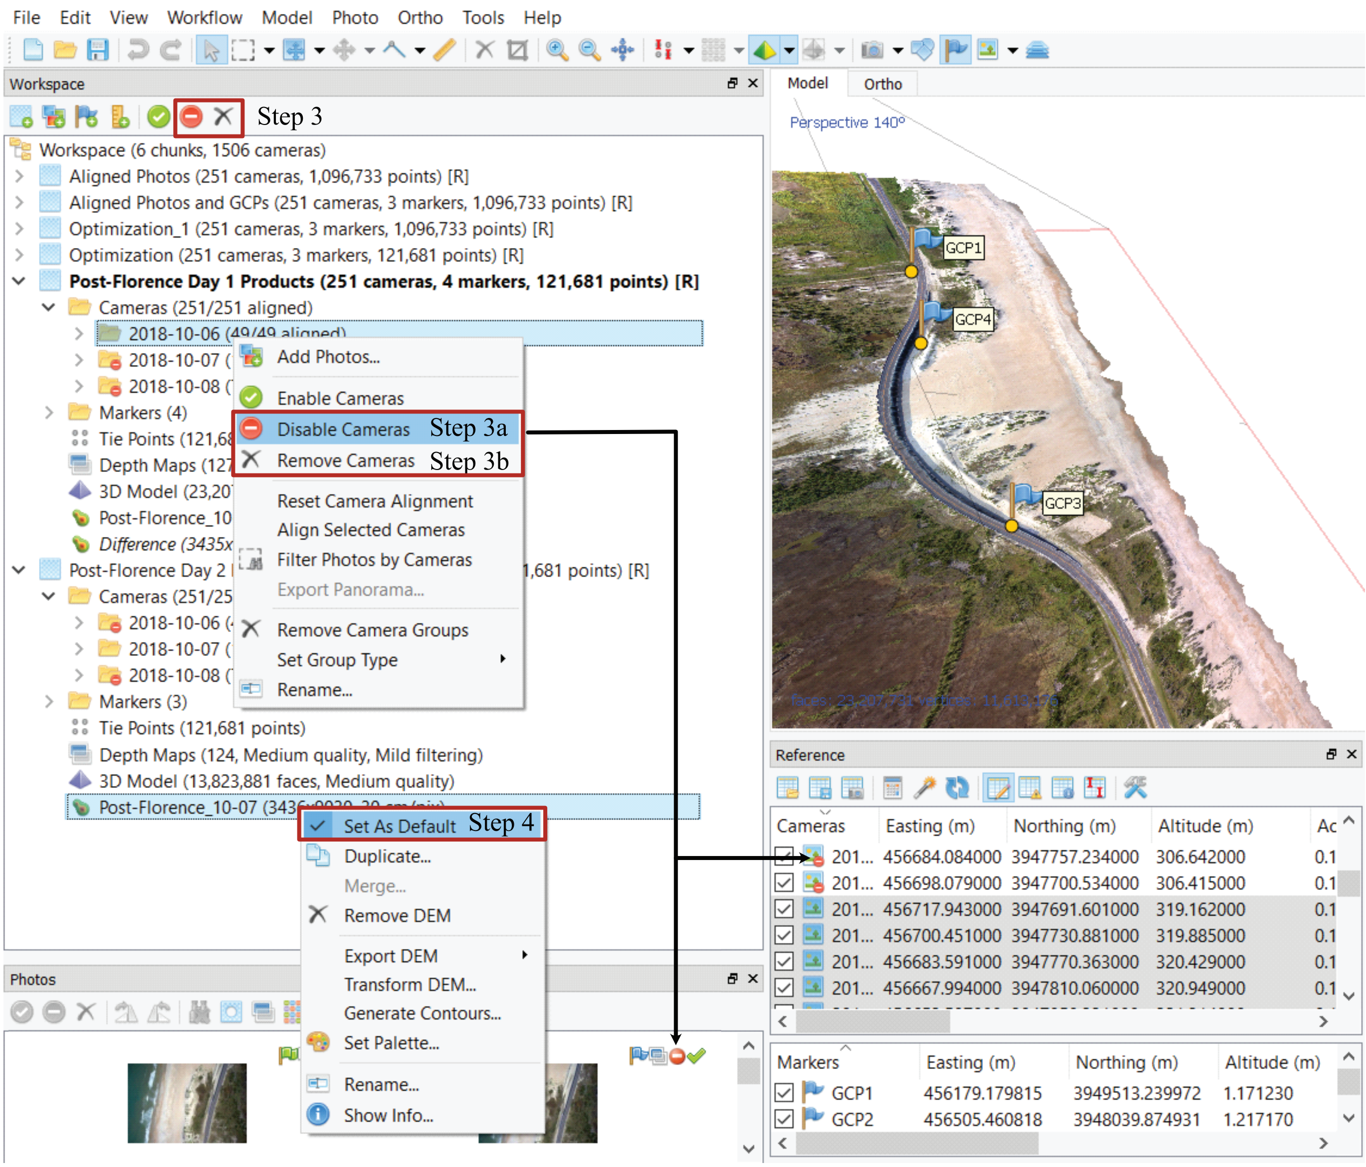

Figure 19. Screenshot of the Workspace panel in Agisoft Metashape Professional Edition that shows the location of the icons to disable and delete cameras and the Disable Cameras and Remove Cameras options when a camera group is right-clicked. Steps are outlined in red. If images are disabled, the red Disable Cameras icon appears in the Photos and Reference panels (black arrows).

- "Calculate point colors" [checked]. Checking this option will result in RGB values for each point, which can be used to classify the dense cloud.
- "Calculate point confidence" [checked]. This option counts how many depth maps have been used to generate each dense cloud point. The parameter can be used for dense cloud filtering.

Once the dense cloud has been created, it generally needs to be edited to remove noise. There are multiple ways to observe, delete, or classify poor-quality points before continuing in the workflow. These include:

- use of the "Dense Cloud Confidence" filter in the Tools >> Dense Cloud >> Filter by Confidence >> Select Confidence Range dialog (fig. 22). On a scale of 0 to 255, values closer to 0 have fewer depth maps involved in the point generation. This option is highly dependent on flight configuration, but deleting points based on confidence levels can remove noise. Be careful to note when the range starts to filter points that are a part of the landscape detail (for example, steeper objects like buildings and cliffs). The filtered points could also be classified as noise and be excluded from end products rather than be deleted;

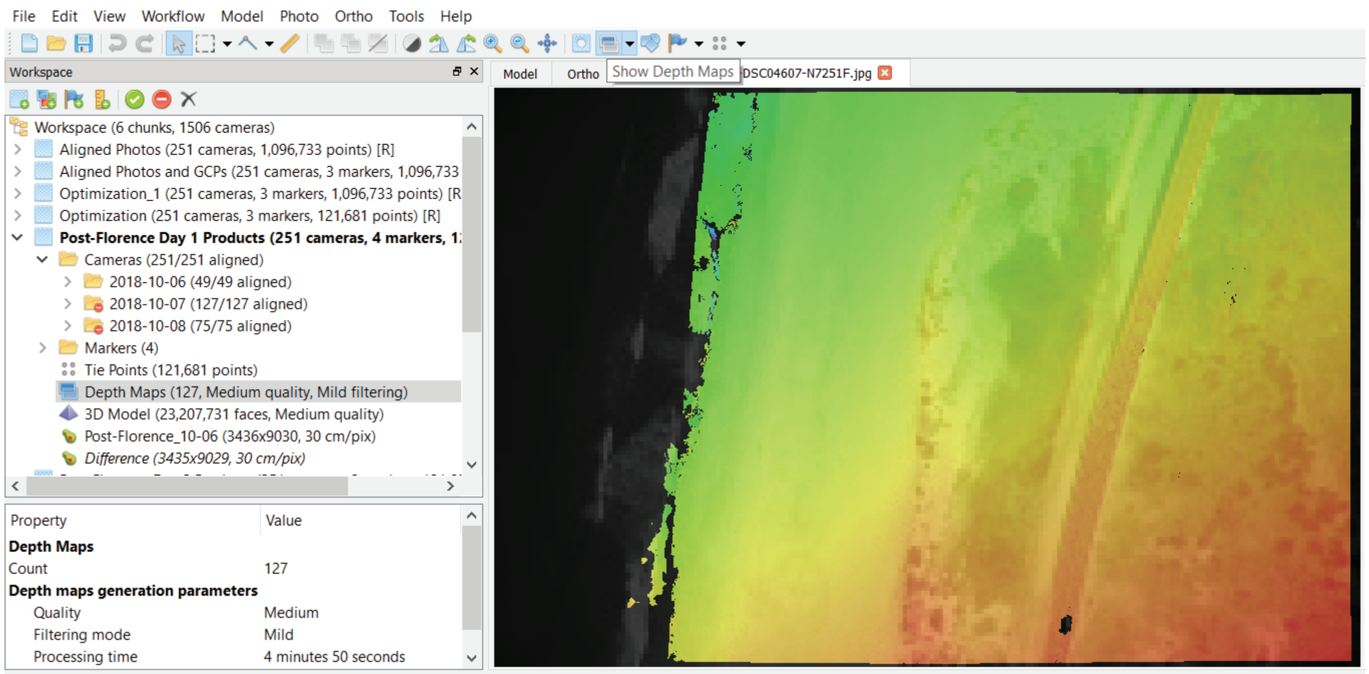

Figure 20. Screenshot of the Agisoft Metashape Professional Edition workspace showing a depth map, as well as the location of the Show Depth Maps icon on the main toolbar. Depth map properties are shown in a separate pane beneath the folder navigation pane (bottom left of Workspace panel). The Show Depth Maps button is only activated when an image has been opened in the working space.

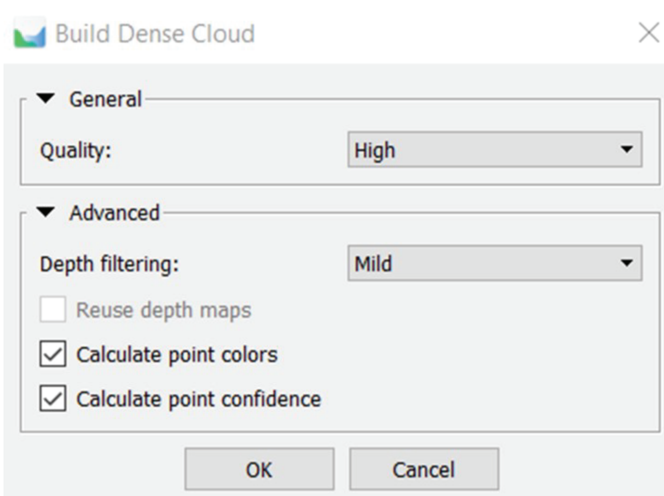

Figure 21. Screenshot of the Build Dense Cloud dialog in Agisoft Metashape Professional Edition, with the “General” and “Advanced” setting panes expanded.

- hand editing “flyers” and “sinkers”, which are unrealistically high and low points when viewing the point cloud horizontally. These can be removed by using the Selection tool and deleting (fig. 22*A*, pink shaded box);
- exporting the dense cloud to a .laz file to be edited in a separate software, such as LAStools (rapidlasso GmbH, 2020). The dense cloud confidence data can be exported using the Chapter 2: Application Modules “exportPoints” function “save_confidence” parameter in the API for use in separate software (Agisoft LLC, 2020c);
- use of Metashape classification tools (see the Metashape user manual); or
- application of masks that could decrease areas of consistent noise (see the “Masks” section of this report).

Build DEM

A DEM is a common final product. Any point classifications can be carried forward into creation of the DEM. The most accurate results are calculated from the dense point cloud data. When a DEM includes both terrain and nonground artifacts, such as buildings and trees, the model is a digital surface model (DSM); when the nonground artifacts are removed, it would be a digital terrain model (DTM). To save more than one DEM in a chunk, right-click on all previously built DEMs and uncheck the “Set as default” option. Use Workflow >> Build DEM and the following settings to create a DEM (data input into steps for the workflow in this report are listed in brackets at the end of the appropriate step; fig. 23):

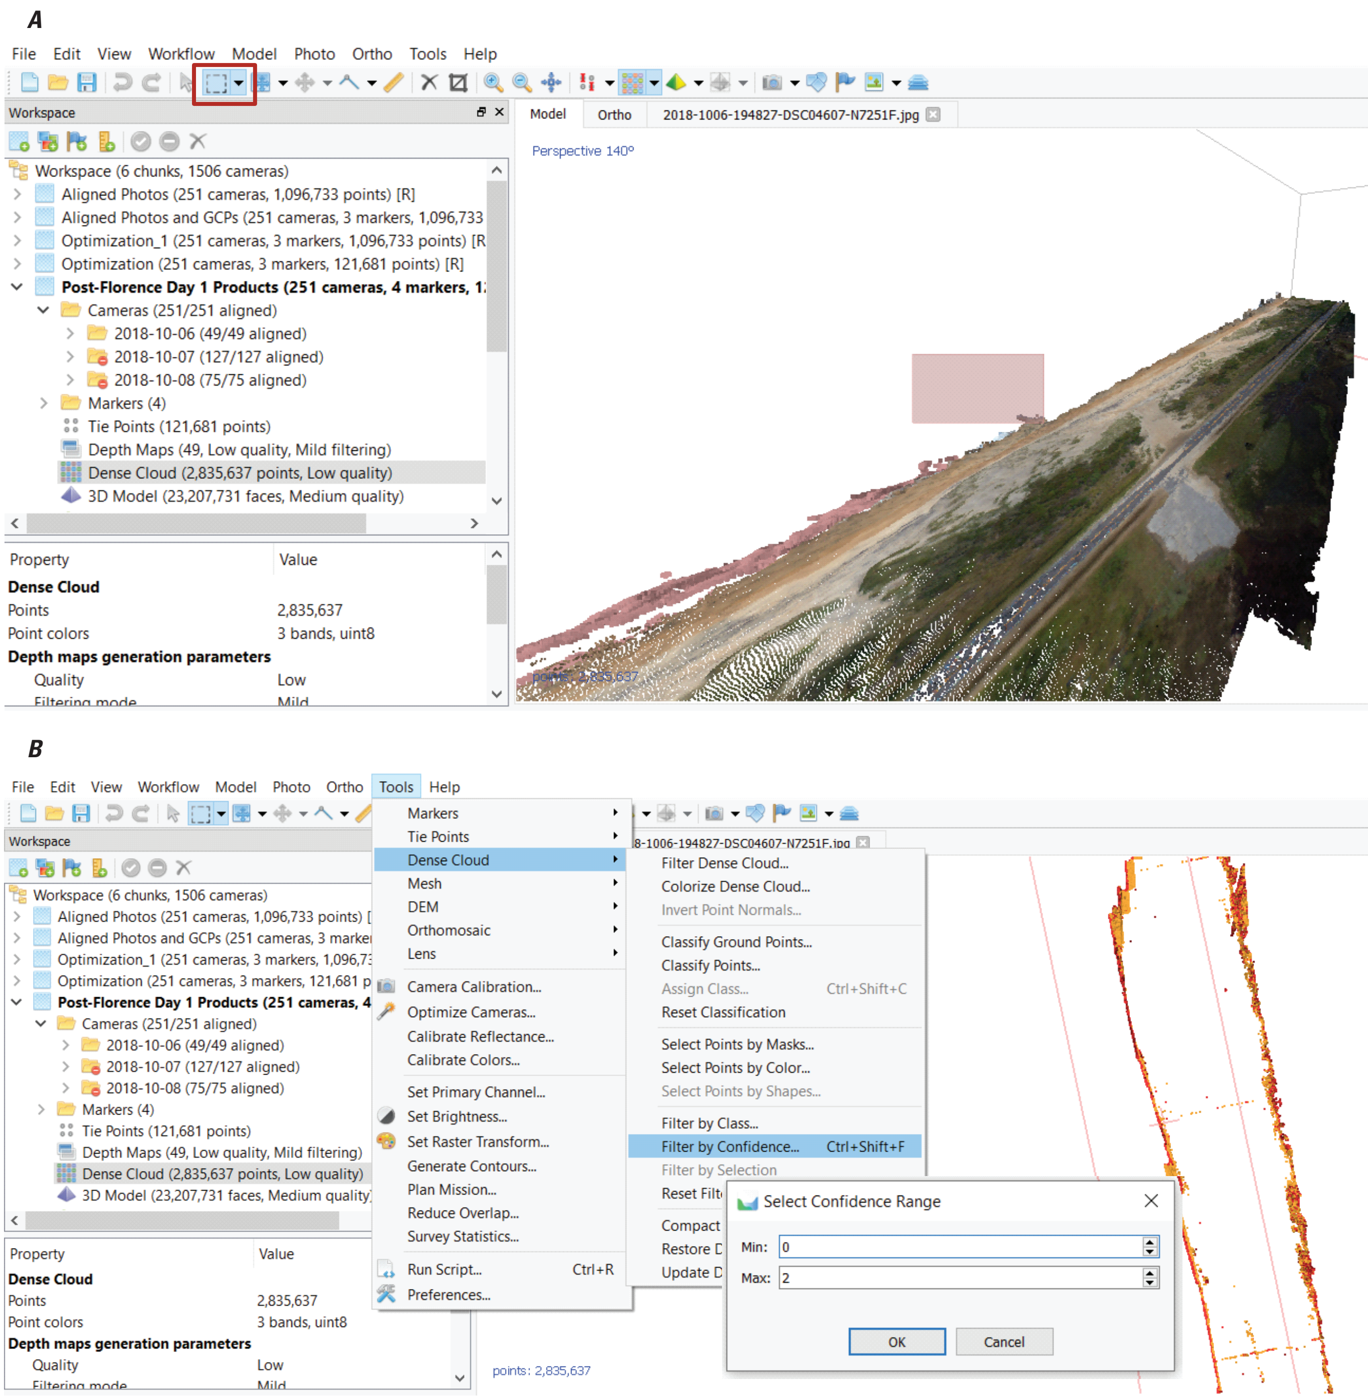

Figure 22. Screenshots of the Agisoft Metashape Professional Edition workspace showing methods to edit and analyze the dense cloud. *A*, The Selection tool, outlined in red, can be used to select noise; the points in the dark pink shaded box are inaccurate reconstructed points, "flyers" and "sinkers"; and *B*, Points can also be isolated using the dense cloud confidence filter and choosing a minimum ("Min") and maximum ("Max") for the range; this limits the points displayed in the workspace, which can then be selected using the Selection tool and then either deleted or reclassified.

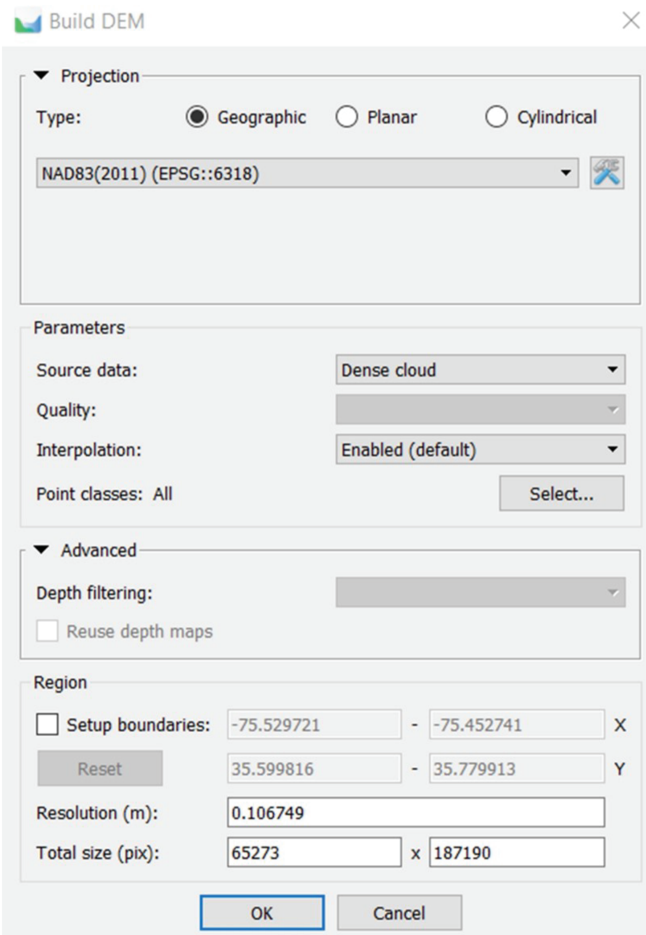

Figure 23. Screenshot of the Build DEM dialog in Agisoft Metashape Professional Edition with all the options expanded.

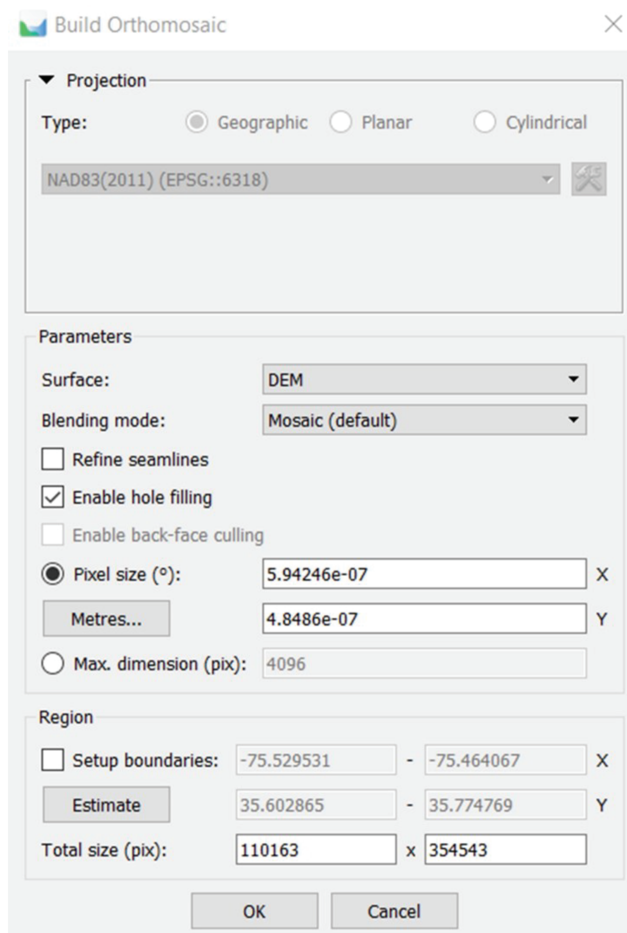

Figure 24. Screenshot of the Build Orthomosaic dialog in Agisoft Metashape Professional Edition.

- Set “Type” [Geographic]. At this point, the coordinate system of the product can be designated in projected or geographic coordinates with ellipsoidal or orthometric heights.
- Select from “Source Data” dropdown [Dense Cloud]. A DEM can also be constructed using the sparse cloud, mesh, or depth maps.
- Select from “Interpolation” dropdown [Disabled]. This option can be enabled to create a DEM and subsequent products, such as an orthomosaic, with interpolation across areas without dense point data. It can be beneficial to create two DEMs: one interpolated for building the orthomosaic without holes, and one without interpolation that is truer to the SFM elevation data reconstruction.
- Select “Point Classes” [All]. If points were classified, for example, when filtering by dense cloud confidence, the points to be kept out of the DEM can be deselected. Option will be grayed out if no points were classified.
- Go to “Region” >> “Setup Boundaries” and “Resolution (m)” and set to [Default]. The default values are the corners of the 3D framework and maximum resolution. These options can be changed when the product is exported.

Build Orthomosaic

The last product commonly generated in the workflow is an orthomosaic generated from the aerial imagery and projected onto a surface. Use Workflow >> Build Orthomosaic and the following settings to create an orthomosaic (data input into steps for the workflow in this report are listed in brackets at the end of the appropriate step; fig. 24):

- Set “Projection Type” [Geographic]. The default projection will be whatever the last surface (DEM, mesh, and so on) was created in. This option can be changed to a projected coordinate system.
- Select from “Surface” dropdown [DEM]. This can also be generated from a mesh.

- Select from “Blending Mode” dropdown [Mosaic] (default). The “Mosaic” option blends the images by dividing images from each camera into multiple frequency domains, using only the highest frequency domain for the seamline, making the seam between individual photos difficult to detect. The average mode uses a weighted average RGB value for all pixels from individual images. In “Disabled” mode, the color of the pixel is taken from the camera view closest to perpendicular to the reconstructed surface.
- “Refine seamlines” [unchecked]. If turned on, this option will create seamlines that bypass complex objects, such as buildings, to avoid visual artifacts. If the surface is a DTM (see the “Build DEM” section of this report), turn this option off.
- “Enable hole filling” [checked]. This is relevant mostly for small scale or nonlandscape projects to prevent data gaps in the output.
- “Pixel size” (in meters or degrees) [Default]. Refers to the maximum ground sampling resolution. Leave as is, the resolution can be changed when the product is exported. Note that projected coordinate reference systems are measured in meters and geographic systems are measured in degrees.
- “Setup Boundaries” [Default]. The numbers represent the boundaries of the 3D framework. Leave as is; the boundaries can be changed when the product is exported.

Exporting Products and Error Reporting

Products can be exported from the File menu or by right-clicking on the desired chunk >> “Export…”. Ideally, the quality of the results and products are always reported; refer to James and others (2019) for guidance on reporting the quality. Error metrics that describe bias, accuracy, and precision should distinguish between systematic and random error (Höhle and Höhle, 2009; James and others, 2017). Metashape provides a full report that includes markers, camera (image) locations, and point cloud errors (see the “Exporting Project Report” section of this report), though the accuracies assigned in the “Reference Settings” window are also relevant (see for example, Sherwood and others, 2018). Any correlation between lens coefficients and residual errors (see section “Camera Calibration Check” section of this report) is pertinent to the integrity of the camera model and may require further evaluation, though an indepth review of this task is beyond the scope of this report.

Exporting Dense Cloud Products

The workflow used for this report does not include dense point clouds as a final product. Dense point clouds may be exported before building the DEMs and other products, edited externally, and reimported if Metashape does not have the capabilities to process the points cloud to the extent the project needs. Use File or right-click on a chunk >> Export >> Export Points. See the Metashape user manual for full documentation on the supported file types.

Exporting DEMs and Orthomosaics

The elevation products from the workflow used for this report are exported as GeoTIFFs, from File or right-click on chunk >> Export >> Export DEM >> “Export TIFF/BIL/XYZ…”; and an orthomosaic is exported from File or right-click on chunk >> Export >> Export Orthomosaic >> “Export JPG/TIFF/PNG…”. At this point in the workflow, the products can be adjusted to export at a lower resolution (in meters or degrees) in the Export DEM dialog (fig. 25). In order to ensure that multiple surveys are comparing data for the same area, and to more easily estimate error, products are exported with coincident bounds. This is most easily done by limiting bounds to an integer multiple of the cell size in units of 10–100 meters.

Exporting Project Report

Metashape provides a comprehensive project report that details the survey data (for example, number of images, coverage area), results of the camera calibration, error estimates of GCPs and camera locations, the processing parameters in the “Show Info” tab, and information about any products created. The report can be generated from File >> Export >> “Generate Report…”. The report may be used for quality control of the overall results of the SFM project. The individual and total errors of the camera locations and GCPs should be reported when publishing products. In general, a project may be deemed “good” if the following goals have been achieved (this list is not comprehensive):

- Point cloud reprojection error or unweighted RMS reprojection error is less than 0.3 pixel.
- The residual error vectors (see the “Camera Calibration Check” section of this report) are less than 1 pixel long.
- The total error of the camera locations and markers is equal to or less than the prescribed accuracy in the Reference Settings dialog.
- The above goals were achieved and about 10 percent of the original number of tie points remain. A lot more points (more than 25 percent) may suggest that further reduction of error is possible and less than that may indicate that too many points were removed.

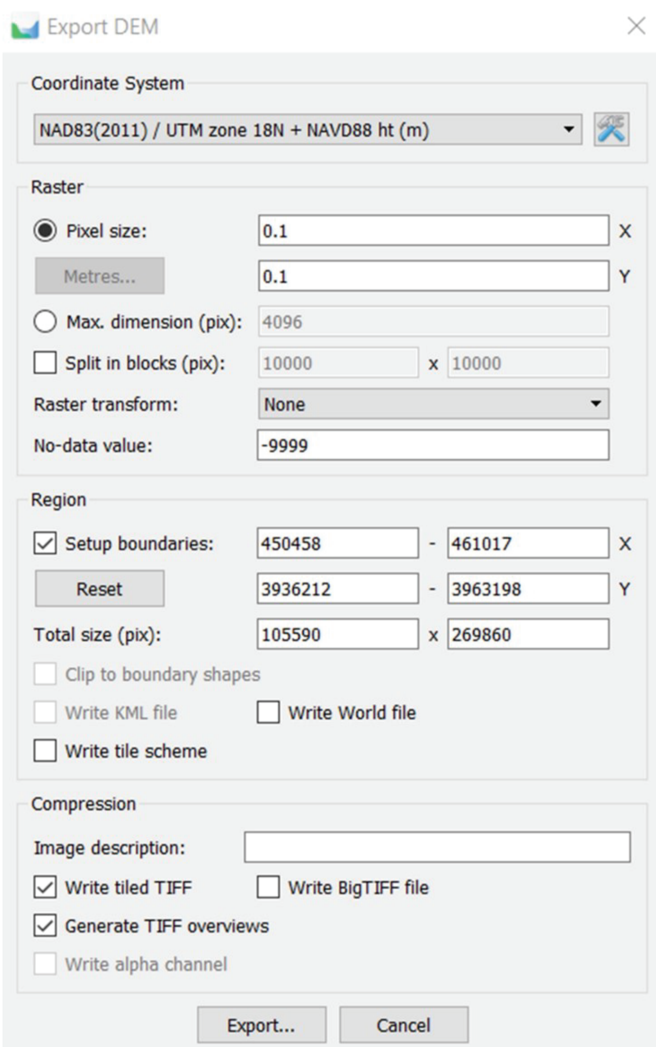

Figure 25. Screenshot of the Export DEM dialog in Agisoft Metashape Professional Edition; the "Coordinate System" dropdown has been changed to a compound system; the "Raster" resolution has been adjusted to a round number, and the "Region" section has been edited to reflect rounded values in the "Setup boundaries".

Creating Compound Coordinate Reference Systems

The products in this workflow are exported using a compound coordinate system. For example, the Hurricane Florence data were made up of the North American Datum of 1983 (2011 update; horizontal datum; NAD83[2011]) and then exported in NAVD88(2011) and Universal Transverse Mercator Zone 18N projection (UTM Zone 18N), and the elevations (height) in the vertical datum were changed from ellipsoid height to the North American Vertical Datum of 1988 (NAVD88) measured in meters, resulting in the compound coordinate system of "NAD83(2011) / UTM Zone 18N + NAVD88 height (m)". This is not natively in the Metashape library, but the datum can be created in Metashape (fig. 26). If the error "Selected vertical datum is unavailable" occurs, then the most likely explanation is that the proper geoid was not placed in the Metashape directory. Use the following steps to create a compound coordinate reference system:

1. Download the necessary geoid (Agisoft LLC, 2021b); the example in figure 26 is "GEOID 12B (ESPG::5103) for NAD83(2011; ESPG::6318)". Currently (2021), the most up-to-date geoid is GEOID 18.
2. Place the downloaded geoid into the correct Metashape directory, based on where the Metashape files were placed during installation; for example, C:\Program Files\Agisoft\Metashape Pro\geoids\.
3. Restart Metashape.
4. Reference panel >> Reference Settings >> "Coordinate System" dropdown >> "More…".
5. Use the "Filter" input bar or in Projected Coordinate Systems, navigate to the correct datum, in the example, NAD83(2011) / UTM zone 18N.
6. In the Select Coordinate System dialog, select the "Edit Coordinate System" pencil icon.
7. Change the "Name" to reflect the compound system being created.
8. In this example, the "Vertical CS" field was changed from the previously selected "Ellipsoidal Height" to "NAVD88 height (EPSG::5703)" in the dropdown list. This completes the switch to orthometric heights based on the downloaded geoid.
9. Click "OK".
10. The new system is now saved under "User-Defined Coordinate Systems" accessed from the Select Coordinate System dialog.

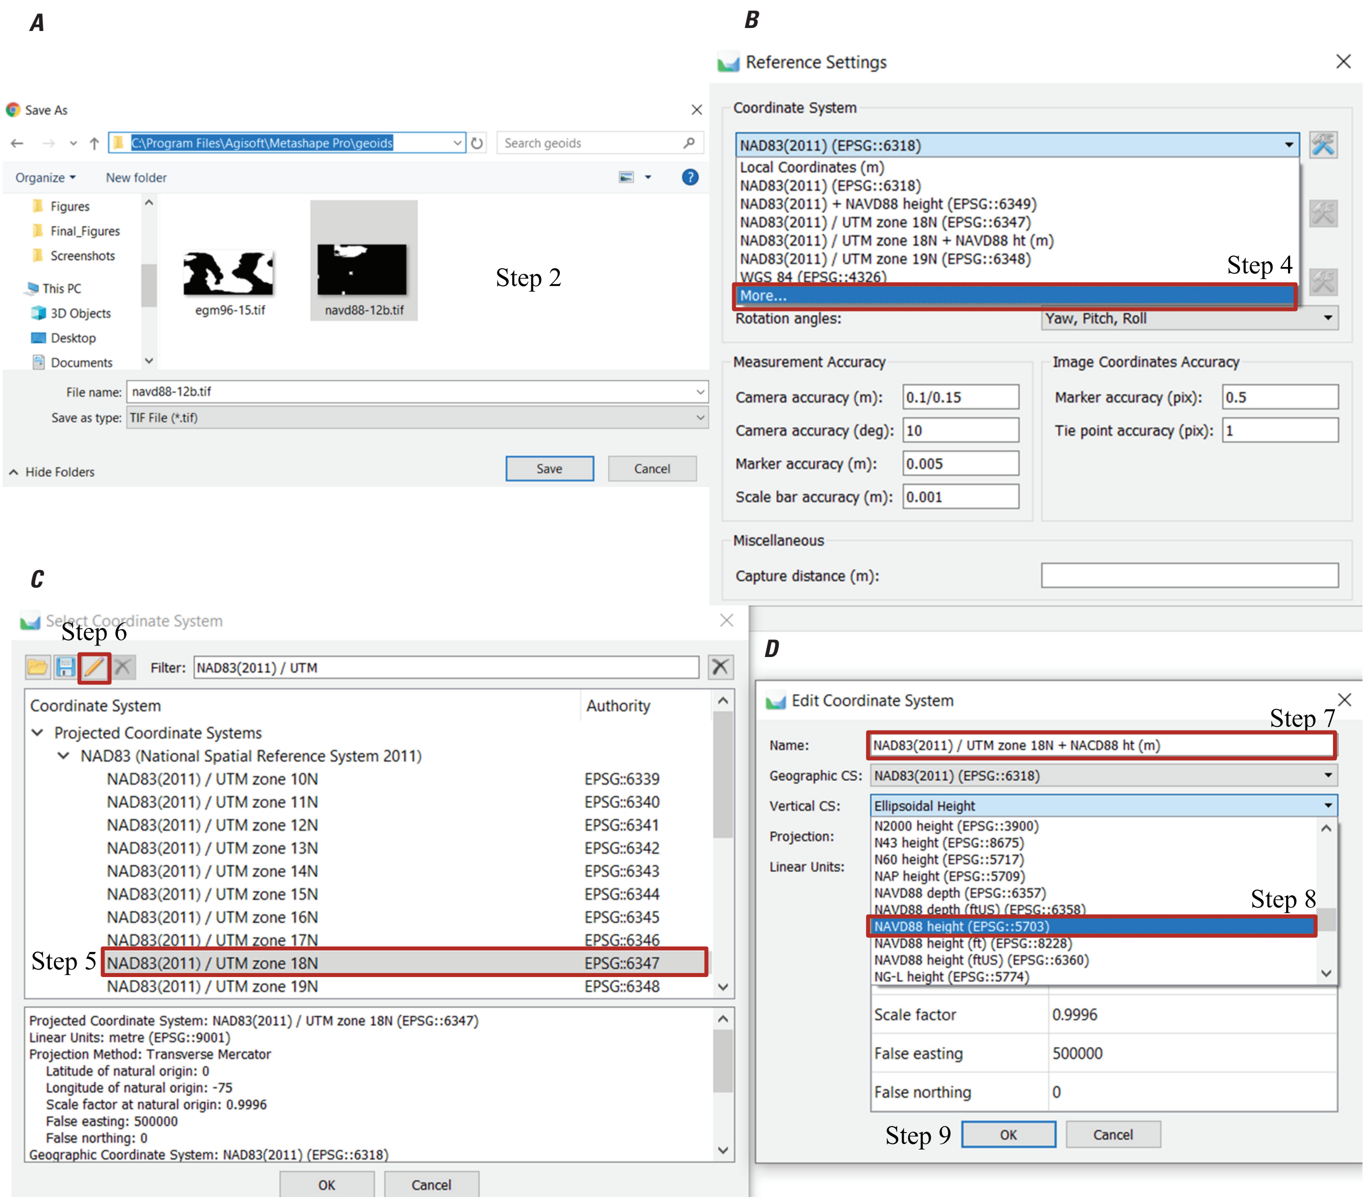

Figure 26. Screenshots of the process in Agisoft Metashape Professional Edition to create a compound coordinate reference system, with specific steps outlined in red. *A*, Download the correct geoid and place it in the correct directory; *B*, Navigate to Reference Settings and select “More…” from the “Coordinate System” dropdown; *C*, From the Select Coordinate System dialog select the coordinate system and then choose the edit (shaped like a pencil) button to bring up *D*, The Edit Coordinate System dialog, where the different parameters can be selected to build the compound coordinate reference system.

Batch Processing

Metashape can run multiple steps of the SFM workflow on multiple chunks without user intervention by a batch process, which can keep processing steps consistent between projects or users. A batch process can be saved and run for many parts of the workflow, including alignment, optimization, building, and exporting. The batch process step progression contains the same settings as the steps in the "Building Products" section of this report (fig. 27). See the "Automation" section in the Metashape user manual for a full description of batch processing.

1. Workflow >> Batch Process >> "Add…".
2. Select "Build DEM" from "Job type".
3. Target the desired chunk(s) by using the "Apply to" dropdown to choose "Selection".
4. Change or accept the job type settings:
   a "Point classes" set to "All". If points were classified and are meant to be excluded, double-click on "All" >> "…" >> and deselect the classes.

   b "Projection" is not added by default; double-click to set.

   c "Resolution (m/pix)" set to 0 means the native resolution will be used to create the product. This can be changed to a lower resolution immediately or when exporting the product; note that there is a bug in the early 1.6.x versions that would accept a lower resolution but still build at native resolution.
5. When done setting the property values, click "OK".
6. Add another job. In this example, add a "Build Dense Cloud" step. Target the same chunk(s) as step 3.
7. Change or accept the job type settings; the value in bracket is the setting selected in the example:
   a "Quality" [High]; and

   b "Depth filtering" [Mild].
8. Reorder the jobs in the batch dialog using the green arrows; the DEM is built second based on the Dense Cloud.
9. Add a final job "Build Orthomosaic". See the setting choices in the "Build Orthomosaic" section of this report. "Projection" is not added by default; double-click to set.
10. Save the batch process with all the jobs added to a file so it can be loaded for other projects. Use the "Agisoft Batch Jobs .xml" as the file type.

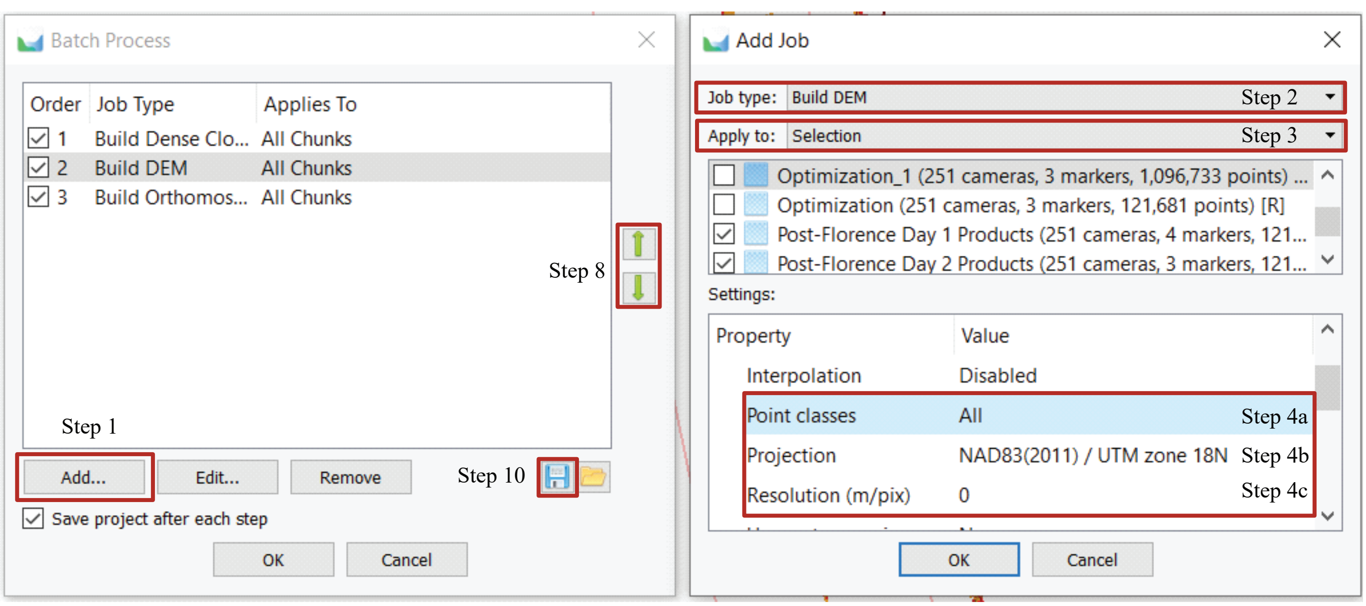

Figure 27. Screenshot of the process in Agisoft Metashape Professional Edition to create a Batch Process, with specific steps outlined in red. *A*, Use the Batch Process dialog to add, edit, remove, and reorder jobs; and *B*, If the "Add…" button is selected, the Add Job dialog appears, where the job settings can be changed.

Limitations in the SFM Workflow

This document is not an exhaustive study in all possible permutations for an SFM workflow. The SFM workflow established for the imagery collected as part of the USGS hurricane relief and recovery efforts project is meant to guide rather than be followed strictly. Many decisions depend on the nature of the imagery and the resolution at which geospatial outputs are required. Therefore, the diversity of imagery and projects that find this guide useful should not be inhibited by the narrow processing scope in this report, which is intended for coastal aerial datasets.

Agisoft updates the Metashape software frequently, and the Metashape user manual should be consulted when updates become available. Although the concepts of photogrammetry do not change, the way software refers to them sometimes does. Changes in the application of technology to photogrammetric processes lead to updates in software and workflows.

References Cited

Agisoft LLC, 2020a, Agisoft Metashape 1.6.5—Professional edition: Agisoft LLC software, accessed October 1, 2020, at https://www.agisoft.com/downloads/installer/.

Agisoft LLC, 2020b, Agisoft Metashape user manual—Professional edition, version 1.6: Agisoft LLC, 166 p., accessed January 22, 2021, at https://www.agisoft.com/pdf/metashape-pro_1_6_en.pdf.

Agisoft LLC, 2020c, Metashape Python reference—Release 1.6.5: Agisoft LLC, 205 p., accessed February 10, 2021, at https://www.agisoft.com/pdf/metashape_python_api_1_6_5.pdf.

Agisoft LLC, 2021a, Agisoft forum: Agisoft LLC web page, accessed January 25, 2021, at https://www.agisoft.com/forum/.

Agisoft LLC, 2021b, Downloads: Agisoft LLC web page, accessed January 27, 2021, at https://www.agisoft.com/downloads/geoids/.

Alfredsen, K., Haas, C., Tuhtan, J.A., and Zinke, P., 2018, Brief communication—Mapping river ice using drones and structure from motion: The Cryosphere, v. 12, no. 2, p. 627–633. [Also available at https://doi.org/10.5194/tc-12-627-2018.]

Brown, D.C., 1958, A solution to the general problem of multiple station analytical stereotriangulation: U.S. Department of the Air Force, RCA Data Reduction Technical Report 43, AFMTC–TR–58–8, 57 p. [Also available at https://digital.hagley.org/08206139_solution.]

Brown, D.C., 1976, The bundle adjustment—Progress and prospect, *in* International Society for Photogrammetry and Remote Sensing, congress, 13, Helsinki, Finland, July 12–23, 1976, invited paper: Committee of the XIII International Congress of Photogrammetry Helsinki 1976, pt. 3, paper 3–03–180, 33 p. [Also available at https://goobi.tib.eu/viewer/image/1066535345/502/.]

Brown, J.A., Sherwood, C.R., Martini, M.A., Kranenburg, C., and Over, J.-S., 2021, Ground control point data from the Outer Banks, North Carolina, post-Hurricane Dorian, September 2019: U.S. Geological Survey data release, https://doi.org/10.5066/P9DVZC23.

Buscombe, D., and Ritchie, A.C., 2018, Landscape classification with deep neural networks: Geosciences, v. 8, no. 7, paper 244, 23 p., accessed January 22, 2021, at https://doi.org/10.3390/geosciences8070244.

Carrivick, J.L., and Smith, M.W., 2019, Fluvial and aquatic applications of structure from motion photogrammetry and unmanned aerial vehicle/drone technology: WIREs Water, v. 6, no. 1, article e1328, 17 p., accessed August 3, 2020, at https://doi.org/10.1002/wat2.1328.

Carrivick, J.L., Smith, M.W., and Quincey, D.J., 2016, Structure from motion in the geosciences: West Sussex, United Kingdom, John Wiley & Sons, Ltd., 208 p. [Also available at https://doi.org/10.1002/9781118895818.]

Chirico, P.G., Bergstresser, S.E., DeWitt, J.D., and Alessi, M.A., 2020, Geomorphological mapping and anthropogenic landform change in an urbanizing watershed using structure-from-motion photogrammetry and geospatial modeling techniques: Journal of Maps, 12 p., accessed August 3, 2020, at https://doi.org/10.1080/17445647.2020.1746419.

Fonstad, M.A., Dietrich, J.T., Courville, B.C., Jensen, J.L., and Carbonneau, P.E., 2013, Topographic structure from motion—A new development in photogrammetric measurement: Earth Surface Processes and Landforms, v. 38, no. 4, p. 421–430. [Also available at https://doi.org/10.1002/esp.3366.]

Forlani, G., Dall'Asta, E., Diotri, F., di Cella, U.M., Roncella, R., and Santise, M., 2018, Quality assessment of DSMs produced from UAV flights georeferenced with on-board RTK positioning: Remote Sensing, v. 10, no. 2, article 311, 22 p., accessed January 22, 2021, at https://doi.org/10.3390/rs10020311.

Forsmoo, J., Anderson, K., Macleod, C.J.A., Wilkinson, M.E., and Brazier, R., 2018, Drone-based structure-from-motion photogrammetry captures grassland sward height variability: Journal of Applied Ecology, v. 55, no. 6, p. 2587–2599. [Also available at https://doi.org/10.1111/1365-2664.13148.]

Fraser, C.S., 2013, Automatic camera calibration in close range photogrammetry: Photogrammetric Engineering and Remote Sensing, v. 79, no. 4, p. 381–388. [Also available at https://doi.org/10.14358/PERS.79.4.381.]

Hatcher, G.A., Warrick, J.A., Ritchie, A.C., Dailey, E.T., Zawada, D.G., Kranenburg, C., and Yates, K.K., 2020, Accurate bathymetric maps from underwater digital imagery without ground control: Frontiers in Marine Science, v. 7, article 525, 20 p., accessed January 22, 2021, at https://doi.org/10.3389/fmars.2020.00525.

Höhle, J., and Höhle, M., 2009, Accuracy assessment of digital elevation models by means of robust statistical methods: ISPRS Journal of Photogrammetry and Remote Sensing, v. 64, no. 4, p. 398–406. [Also available at https://doi.org/10.1016/j.isprsjprs.2009.02.003.]

James, M.R., and Robson, S., 2012, Straightforward reconstruction of 3D surfaces and topography with a camera—Accuracy and geoscience application: Journal of Geophysical Research: Earth Surface, v. 117, article F03017, 17 p., accessed January 22, 2021, at https://doi.org/10.1029/2011JF002289.

James, M.R., Chandler, J.H., Eltner, A., Fraser, C., Miller, P.E., Mills, J.P., Noble, T., Robson, S., and Lane, S.N., 2019, Guidelines on the use of structure-from-motion photogrammetry in geomorphic research: Earth Surface Processes and Landforms, v. 44, no. 10, p. 2081–2084. [Also available at https://doi.org/10.1002/esp.4637.]

James, M.R., Robson, S., and Smith, M.W., 2017, 3-D uncertainty-based topographic change detection with structure-from-motion photogrammetry—Precision maps for ground control and directly georeferenced surveys: Earth Surface Processes and Landforms, v. 42, no. 12, p. 1769–1788. [Also available at https://doi.org/10.1002/esp.4125.]

Jones, C.A., and Church, E., 2020, Photogrammetry is for everyone—Structure-from-motion software user experiences in archaeology: Journal of Archaeological Science: Reports, v. 30, article 102261, 10 p., accessed January 22, 2021, at https://doi.org/10.1016/j.jasrep.2020.102261.

Kranenburg, C., Ritchie, A.C., Brown, J.A., Over, J.R., Buscombe, D., Sherwood, C.R., Warrick, J.A., and Wernette, P.A., 2020, Post-Hurricane Florence aerial imagery—Cape Fear to Duck, North Carolina, October 6–8, 2018: U.S. Geological Survey data release, accessed March 12, 2021, at https://doi.org/10.5066/P91KB9SF.

Logan, J.B., and Ritchie, A.C., 2020, Metashape (version 1.0): U.S. Geological Survey code repository, accessed January 22, 2021, at https://code.usgs.gov/PhotoscanScripts/PhotoscanAlignmentErrorReduction/-/releases/Metashape.

Matthews, N.A., Noble, T.A., and Breithaupt, B., 2006, The application of photogrammetry, remote sensing and geographic information systems (GIS) to fossil resource management, *in* Lucas, S.G., Spielmann, J.A., Hester, P.M., Kenworthy, J.P., and Santucci, V.L., eds., America's Antiquities—100 Years of Managing Fossils on Federal Lands, 7th Federal Fossils Conference, Albuquerque, May 2006, [abstracts and papers]: New Mexico Museum of Natural History Bulletin 34, p. 119–131. [Also available at https://econtent.unm.edu/digital/collection/bulletins/id/510.]

National Oceanic and Atmospheric Administration, 2020, Emergency response imagery: National Oceanic and Atmospheric Administration web page, accessed January 25, 2021, at https://storms.ngs.noaa.gov/.

Nouwakpo, S.K., Weltz, M.A., and McGwire, K., 2016, Assessing the performance of structure-from-motion photogrammetry and terrestrial lidar for reconstructing soil surface microtopography of naturally vegetated plots: Earth Surface Processes and Landforms, v. 41, no. 3, p. 308–322. [Also available at https://doi.org/10.1002/esp.3787.]

NovAtel, 2020, GrafNav GNSS post-processing software: NovAtel web page, accessed January 25, 2021, at https://novatel.com/products/waypoint-software/grafnav.

Puget Systems, 2020, Technology articles about Metashape: Puget Systems web page, accessed January 25, 2021, at https://www.pugetsystems.com/all_articles.php?filter[]=Metashape.

rapidlasso GmbH, 2020, LAStools: rapidlasso GmbH software, accessed January 27, 2021, at https://rapidlasso.com/lastools/.

Sanz-Ablanedo, E., Chandler, J.H., Rodríguez-Pérez, J.R., and Ordóñez, C., 2018, Accuracy of unmanned aerial vehicle (UAV) and SFM photogrammetry survey as a function of the number and location of ground control points used: Remote Sensing, v. 10, no. 10, article 1606, 19 p., accessed January 22, 2021, at https://doi.org/10.3390/rs10101606.

Schenk, T., 2005, Introduction to photogrammetry: Ohio State University course GS400.02, 95 p., accessed January 22, 2021, at http://www.mat.uc.pt/~gil/downloads/IntroPhoto.pdf.

Shervais, K., 2015, Structure from motion introductory guide: UNAVCO, Inc., 18 p., accessed January 22, 2021, at https://www.unavco.org/education/resources/modules-and-activities/field-geodesy/module-materials/sfm-intro-guide.pdf.

Sherwood, C.R., Warrick, J.A., Hill, A.D., Ritchie, A.C., Andrews, B.D., and Plant, N.G., 2018, Rapid, remote assessment of Hurricane Matthew impacts using four-dimensional structure-from-motion photogrammetry: Journal of Coastal Research, v. 34, no. 6, p. 1303–1316. [Also available at https://doi.org/10.2112/JCOASTRES-D-18-00016.1.]

Smith, M.W., Carrivick, J.L., and Quincey, D.J., 2016, Structure from motion photogrammetry in physical geography: Progress in Physical Geography, v. 40, no. 2, p. 247–275. [Also available at https://doi.org/10.1177/0309133315615805.]

Storlazzi, C.D., Dartnell, P., Hatcher, G.A., and Gibbs, A.E., 2016, End of the chain? Rugosity and fine-scale bathymetry from existing underwater digital imagery using structure-from-motion (SFM) technology: Coral Reefs, v. 35, no. 3, p. 889–894. [Also available at https://doi.org/10.1007/s00338-016-1462-8.]

Triggs, B., McLauchlan, P.F., Hartley, R.I., and Fitzgibbon, A.W., 2000, Bundle adjustment—A modern synthesis, *in* Triggs, B., Zisserman, A., and Szeliski, R., eds., Vision Algorithms—Theory and Practice—Seventh International Conference on Computer Vision, Kanoni, Corfu, Greece, September 21–22, 1999, [Proceedings]: New York, Springer-Verlag, p. 298–372. [Also available at https://doi.org/10.1007/3-540-44480-7_21.]

Turner, D., Lucieer, A., and Wallace, L., 2014, Direct georeferencing of ultrahigh-resolution UAV imagery: IEEE Transactions on Geoscience and Remote Sensing, v. 52, no. 5, p. 2738–2745. [Also available at https://doi.org/10.1109/TGRS.2013.2265295.]

U.S. Geological Survey, 2017, Digital single-lens reflex (DSLR) imagery, sec. 1 *of* Unmanned aircraft systems data post-processing—Structure-from-motion photogrammetry: U.S. Geological Survey class book, 21 p., accessed February 10, 2021, at https://uas.usgs.gov/nupo/pdf/PhotoScanProcessingDSLRMar2017.pdf.

Warrick, J.A., Ritchie, A.C., Adelman, G., Adelman, K., and Limber, P.W., 2017, New techniques to measure cliff change from historical oblique aerial photographs and structure-from-motion photogrammetry: Journal of Coastal Research, v. 33, no. 1, p. 39–55. [Also available at https://doi.org/10.2112/JCOASTRES-D-16-00095.1.]

Westoby, M.J., Brasington, J., Glasser, N.F., Hambrey, M.J., and Reynolds, J.M., 2012, 'Structure-from-motion' photogrammetry—A low-cost, effective tool for geoscience applications: Geomorphology, v. 179, p. 300–314. [Also available at https://doi.org/10.1016/j.geomorph.2012.08.021.]

Woodget, A.S., Carbonneau, P.E., Visser, F., and Maddock, I.P., 2015, Quantifying submerged fluvial topography using hyperspatial resolution UAS imagery and structure from motion photogrammetry: Earth Surface Processes and Landforms, v. 40, no. 1, p. 47–64. [Also available at https://doi.org/10.1002/esp.3613.]

Yilmaz, H.M., Yakar, M., Gulec, S.A., and Dulgerler, O.N., 2007, Importance of digital close-range photogrammetry in documentation of cultural heritage: Journal of Cultural Heritage, v. 8, no. 4, p. 428–433. [Also available at https://doi.org/10.1016/j.culher.2007.07.004.]

Glossary

aerial triangulation The iterative process that uses collinearity equations and a bundle adjustment to solve for the three-dimensional coordinates of cameras and key points based on markers, tie points, and initial approximations of exterior orientations.

alignment The process of creating key points and tie points between images to determine their relation and finding the position of the camera for each image.

benchmark The act of running a computer operation to assess the relative performance.

bundle adjustment Refining a visual reconstruction to produce jointly optimal three-dimensional structure and viewing parameter (camera pose or calibration) estimates. See Triggs and others (2000) for an in-depth discussion.

camera location The Cartesian coordinates (x,y,z) of each image based on a Global Positioning System or other global navigation satellite system unit that can be paired with a positional correction methodology, such as postprocessing kinematic and real-time kinematic techniques.

camera orientation Determines the attitude, position, and intrinsic geometric characteristics of the camera and includes interior and exterior orientations. See also “exterior orientation” and “interior orientation”.

check point A type of marker in Agisoft Metashape Professional Edition used to validate the accuracy of the results of camera alignment and optimization procedures rather than to reference the model.

control point A type of marker in Agisoft Metashape Professional Edition, equivalent to ground control point. Used to reference the model. See also “ground control point”.

exterior orientation The position and orientation, or attitude, of the camera when the image was taken or the relation between the ground and the image. The exterior orientation is usually collected with an onboard Global Positioning System or inertial navigation system and computed from ground control points and tie points. Position refers to the locations of ground-space coordinates x,y,z of the camera focal point and the orientation is given by coordinates ω,φ,κ, which are the angles of rotation around the x, y, and z axes, respectively; ω, φ, and κ are also referred to as yaw, pitch, and roll, respectively, when defining the camera with respect to a navigation coordinate system.

ground control point A point on the surface of the Earth with a known location that is used to georeference imagery. See “control point”.

interior orientation Defines the internal geometry of a camera or sensor as it existed at the time an image was taken. It defines the image’s space coordinates based on the focal length, principle point coordinates, and lens distortion coefficients.

key point A point of interest on a two-dimensional image that the program determines to be an important feature that is often a point of high contrast or texture.

key point size The scale of the key point for an image; measured as the standard deviation of the algorithm to reduce image noise and detail.

optimization See bundle adjustment.

photogrammetry The science and technology of obtaining reliable information about the properties of surfaces and objects through the combination and analysis of photographic images. In this report, photogrammetry is performed by extracting three-dimensional measurements from two-dimensional data using geometry and a known scale.

postprocessing kinematic While images are being collected with an onboard global navigation satellite system unit, a separate network with more accurate triangulation is also recording positional information. After images are collected, the two sets of data are matched, and the onboard data is corrected, or postprocessed. The technique is an alternative to real-time kinematic processing, which requires a connection of the global navigation satellite system to a base receiver.

structure from motion A method or application of photogrammetry that relies on extracting three-dimensional structure from overlapping offset two-dimensional images without the need for a previously defined specification of camera positions and orientations. In order to place the three-dimensional structure in real world coordinates, camera positions and orientations are solved for in an iterative bundle adjustment procedure.

tie points A set of key points that have been matched as projections of the same three-dimensional point on different images. Metashape matches images on different scales; the accuracy of tie point projections depends on the map scale at which they are located.

For more information, contact:
Director, Woods Hole Coastal and Marine Science Center
U.S. Geological Survey
384 Woods Hole Road
Quissett Campus
Woods Hole, MA 02543–1598
508–548–8700 or 508–457–2200
whsc_science_director@usgs.gov
or visit our website at
https://www.usgs.gov/centers/whcmsc

Publishing support provided by the
Pembroke Publishing Service Center

Over and others—**Processing Coastal Imagery With Agisoft Metashape—Structure From Motion Workflow Documentation**—OFR 2021–1039

ISSN 2331-1258 (online)
https://doi.org/10.3133/ofr20211039