


**A Practical Guide to Drone Mapping
Using Free and Open Source Software**



OpenDroneMap

The Missing Guide

Second Edition

Piero Toffanin

Contents

Copyright	i
Epigraph	ii
Preface	iii
Acknowledgement	vi
I. Introduction	1
Why OpenDroneMap?	2
What You Can Do With OpenDroneMap	3
Becoming a Successful User	5
II. Getting Started	7
1. The OpenDroneMap Ecosystem	8
2. Installing The Software	10
Hardware Requirements	11
Installing on Windows	12
Installing on macOS	19
Installing on Linux	23

Contents

Basic Commands and Troubleshooting	25
Hello, WebODM!	26
3. Processing Datasets	28
Dataset Size	28
File Requirements	29
Process Tasks	30
Output Results	34
Testing Different Task Options	35
Share With Others	36
Export To Another WebODM	37
Manage Plugins	37
Change The Look & Feel	37
Create New Users & Groups	37
Project Permissions	38
Managing Tags	40
How Does WebODM Process Images?	40
4. The Processing Pipeline	41
Load Dataset	42
Structure From Motion	42
Multi View Stereo	46
Point Filtering	46
Meshing	47
Texturing	49
Georeferencing	51
Digital Elevation Model Processing	52
Orthophoto Processing	54
Report Generation	56
Post Processing	56

5. Task Options in Depth	57
3d-tiles	59
align	60
auto-boundary	62
auto-boundary-distance	62
bg-removal	63
boundary	64
build-overviews	64
camera-lens	65
cameras	68
cog	68
copy-to	69
crop	69
dem-decimation	70
dem-euclidean-map	71
dem-gapfill-steps	72
dem-resolution	74
dsm	75
dtm	75
end-with	76
fast-orthophoto	77
feature-quality	81
feature-type	82
force-gps	82
gcp	83
geo	83
gltf	84
gps-accuracy	84
help	85

Contents

ignore-gsd	85
matcher-neighbors	87
matcher-order	89
matcher-type	89
max-concurrency	91
merge	91
mesh-octree-depth	92
mesh-size	94
min-num-features	95
no-gpu	98
optimize-disk-space	98
orthophoto-compression	98
orthophoto-cutline	99
orthophoto-kmz	101
orthophoto-no-tiled	101
orthophoto-png	102
orthophoto-resolution	103
pc-classify	103
pc-copc	109
pc-csv	109
pc-ept	110
pc-filter	110
pc-las	111
pc-quality	111
pc-rectify	114
pc-sample	115
pc-skip-geometric	115
primary-band	115
project-path	116

Contents

radiometric-calibration	116
rerun	117
rerun-all	117
rerun-from	117
rolling-shutter	118
rolling-shutter-readout	119
sfm-algorithm	119
sfm-no-partial	120
skip-3dmodel	120
skip-band-alignment	122
skip-orthophoto	122
skip-report	122
sky-removal	122
sm-cluster	124
sm-no-align	124
smrf-scalar	124
smrf-slope	124
smrf-threshold	125
smrf-window	125
split	125
split-image-groups	125
split-overlap	125
texturing-keep-unseen-faces	126
texturing-single-material	128
texturing-skip-global-seam-leveling	128
texturing-skip-local-seam-leveling	130
tiles	130
use-3dmesh	131
use-exif	131

Contents

use-fixed-camera-params	132
use-hybrid-bundle-adjustment	132
version	133
video-limit	133
video-resolution	133
Changing Options and Restarting	134
6. Ground Control Points	139
Creating a GCP file using POSM GCPi	143
Creating a GCP file using GCP Editor Pro	147
Using GCP files	147
How GCP files work	149
7. Geolocation Files	151
Using GEO files	154
8. Multispectral Datasets	155
Supported Images	155
Processing	157
Radiometric Calibration Basics	160
Viewing Results in WebODM	162
Thermal Datasets	164
9. Image Masks	166
Manually Creating Image Masks	167
10. Rolling Shutter Correction	171
Usage	172
Limitations	175
11. Camera Calibration	176

12. Report Analysis	180
Dataset Summary	181
Processing Summary	182
Previews	183
Survey Data	183
GPS/GCP/3D Errors Details	184
Feature Details	189
Reconstruction Details	191
Track Details	194
Camera Models Details	194
JSON output	196
13. Flying Tips	197
Fly Higher	197
Fly on Overcast Days	198
Fly Between 10am and 2pm	198
Fly at Different Elevations and Capture Multiple Angles	198
Fly on Calm Days	199
Increase Overlap	200
Set Drone to Hover While Taking Images	200
Check Camera Settings	201
III. Advanced Usages	202
14. The Command Line	203
Command Line Basics	204
Using ODM	206
Processed Files Owned By Root	207
Add New Processing Nodes to WebODM	207

Contents

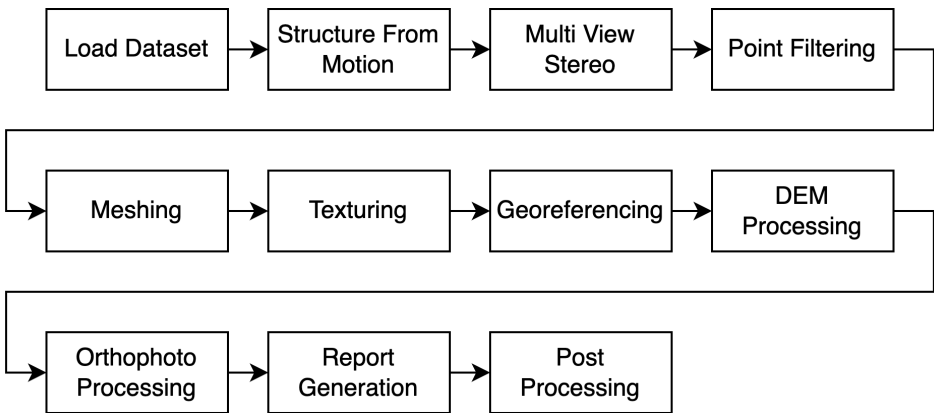
Examine EXIF/XMP Tags	208
Further Readings	209
15. Docker Essentials	210
Docker Basics	210
Managing Containers	212
Managing Images	215
Managing Volumes	216
Docker Compose Basics	218
Managing Disk Space	220
Changing Entrypoint	221
Assigning Names To Containers	221
Jumping Into Existing Containers	222
16. GPU Processing	224
Windows	225
Linux	228
Notes on GPU usage	230
17. Processing Large Datasets	231
Split-Merge Options	233
Local Split-Merge	235
Distributed Split-Merge	237
Using Image Groups and GCPs	241
Limitations	242
18. The NodeODM API	244
Launching a NodeODM Instance	245
NodeODM Configuration	246
Using the API with cURL	247

Contents

Remove a Task	250
API Specification	250
Definitions	271
Exercises	271
19. Automated Processing With Python	273
Getting Started	274
Example 1: Hello NodeODM	275
Example 2: Process Datasets	276
Concluding Remarks	279
API Reference	279
Glossary	287
About The Author	291

4. The Processing Pipeline

Going from images to 3D models and orthophotos is a process best visualized as a series of incremental steps. Each step relies on the work of previous steps.



ODM's processing pipeline

In this chapter we will explore an overview of the pipeline. We will not cover too many details, as each step's behavior can be tweaked by changing the task options. We will discuss in depth of how task options affect the inner workings of the pipeline in the next chapter.

5. Task Options in Depth

There are several steps involved in the data processing pipeline. Each step has several adjustable settings that influence the output. The software exposes a subset of these available knobs through various options. When creating a task, a user can choose to tweak one or more options to change the behavior of the pipeline.

5. Task Options in Depth

Edit Task Options ✕

3d-tiles ⓘ ⌵

☐ Enable

auto-boundary ⓘ ↻

☒ Enable

auto-boundary-distance (positive float) ⓘ

bg-removal ⓘ

☐ Enable

Cancel Save

Options as shown in WebODM when creating a task

If the list seems overwhelming, just remember that this is a subset of all possible options that could be available from the various steps of the data pipeline. Hidden features and processing capabilities could be hiding in the source code of ODM, in the form of an option not yet exposed! The software exposes only those options that have shown the biggest impact on results, or those necessary to handle different workflows. But many, many more options, under the hood, remain unexposed in order to keep their number somewhat manageable.

Tuning options is more art than science. That's mostly because the best options for

5. Task Options in Depth

certain datasets do not automatically transfer to others. As a general rule, one should start with the defaults, which work fairly well for most datasets and apply tweaks as needed.

This chapter is about understanding in detail what each option does. By the end of the chapter you'll be able to quickly improve your results, explain why certain models turn out the way they do and know what to tweak if the results don't turn out the way you want.

A few of these options might be missing from WebODM and might be available only from ODM. This is because sometimes the option does not make sense in the context of the graphic interface workflow, or it's simply not supported.

When there is some math to explain, I write the formulas using Python because it's easier than math notation and can be typed on a computer. You can copy/paste the code on a website such as online-python.com and follow along even if you don't know Python.

Feel free to jump around and use this chapter as a reference. As the software gets better, some of these options might disappear from future versions and new ones might be introduced. The list below is taken from the software as of June 10th 2023. In alphabetical order:

3d-tiles

OGC 3D Tiles¹ are a format specification for visualization and interaction with 3D geospatial content. These files can be displayed with software such as the open source virtual globe engine Cesium². ODM has support for generating point clouds

¹OGC 3D Tiles: ogc.org/standard/3dtiles/


²Cesium: github.com/CesiumGS/cesium

5. Task Options in Depth

Euclidean map results are stored in the *odm_dem* directory.

dem-gapfill-steps


The process of going from point cloud to DEMs is not as straightforward as it may seem. Since DEMs are *rasters* (images), they have *cells* (pixels). Each cell, should have a value. Depending on the resolution of the raster, certain cells may have zero, one or more points that fall into it. Every cell needs a value, even if no points fall directly into it, otherwise there will be empty areas (gaps) in the DEM! One way to overcome this is to use a radius around each cell. Every point that falls within the radius is considered part of the cell.



Pixels and points (left), radius of 0.5 (middle) and radius of 1 (right)

But how big should the radius be? If too small, as in the 0.5 radius example above, some cells might remain empty. If too big, there will be too much smoothing and accuracy will suffer. Since different point clouds have varying degrees of density, one solution is to compute multiple DEMs with different radiuses and stack them.


5. Task Options in Depth



Sparse (top) vs. dense (bottom) point cloud outputs


Both point clouds can be used to generate a mesh. However, it's better to have more points, as meshes can be created with more details. In the dense point cloud screenshot above, the building on the right side of the scene is well defined, but it's poorly represented in the sparse point cloud. Buildings are especially difficult to model without a dense point cloud, so this option tends to yield poor results in

5. Task Options in Depth



Initial graph (left) and graph with randomly moved positions and new edges (right).
Every edge indicates an image pair

ODM also supports a different method to perform preemptive matching by considering only the nearest neighbors of each image instead of using a graph. It is enabled by setting this option. The illustration below shows the result of setting this option to 8:




Dots represent approximate image locations, extracted from EXIF tags. When the `matcher-neighbors` is set to 8, only the 8 nearest neighbors (highlighted in gray) are considered for matching with image p1

This option can sometimes be beneficial for speeding up processing by reducing the number of matching pairs. If no GPS information is available, this option is disabled and all image pairs are considered, unless `matcher-order` is specified.

mesh-octree-depth

When it comes to generating 3D models, this is probably the most important option. It specifies a key variable for the Screened Poisson Reconstruction²¹ algorithm, which is responsible for generating a mesh from the point cloud. The details of the algorithm are fascinating, but probably outside the scope of this book.

To understand how this option affects the output, it helps to visually understand the concept of an octree. First, octree means *eight-tree* (okta is *eight* in Greek). Why eight? Because at each level (or *depth*) of the tree, each box (or *node* or *branch*) of the tree is divided in eight parts. At the first level there's only one branch. At the second level there's 8. At the third there's 64 and so forth.




An octree with depth 1, 2 and 3

Lower depths in an octree allow finer details to be captured.

²¹Screened Poisson Reconstruction: [watertight surfaces from oriented point sets.](http://watertight-surfaces.cs.jhu.edu/~misha/MyPapers/ToG13.pdf)
cs.jhu.edu/~misha/MyPapers/ToG13.pdf

6. Ground Control Points



A GCP marker is photographed by camera A to produce Image A. In a second step, the pixel location of the marker (1500,1000) from Image A can be manually tagged with its real world coordinates (latitude 40, longitude -85)

Using ground control points can increase the georeferencing accuracy of a reconstruction, since measurements of static (non-moving) objects using a high precision GPS are often better than those obtained from the GPS of moving UAVs.


The ideal number of ground control points ranges between 5 to 8, placed evenly across the area to be flown. Adding more than 8 ground control points does not necessarily result in increased accuracy.

If the same marker is visible from multiple images, it should be tagged multiple times for each image. Ideally each marker should be tagged at least 3 times. Another way to think of it is to capture each marker on at least 3 images. This is so that the marker's location can be triangulated during computation.

Ground control points can be used by providing an additional text file along with the input images. The file follows a simple format:


The first line indicates the spatial reference system (SRS) of the world coordinates. There are no restrictions on the type of SRS you can use. Internally the program will

11. Camera Calibration



Typical flight path from mission planning software. Not great for self-calibration

This doesn't mean a person should never fly this pattern. It just means that when flying this pattern, people need to be aware that the internal camera parameters' estimates will not be as good. Inaccurate parameters lead to an improper camera lens model, which over large areas typically results in a *doming* effect.



Point cloud exhibiting doming. The terrain appears arched instead of straight

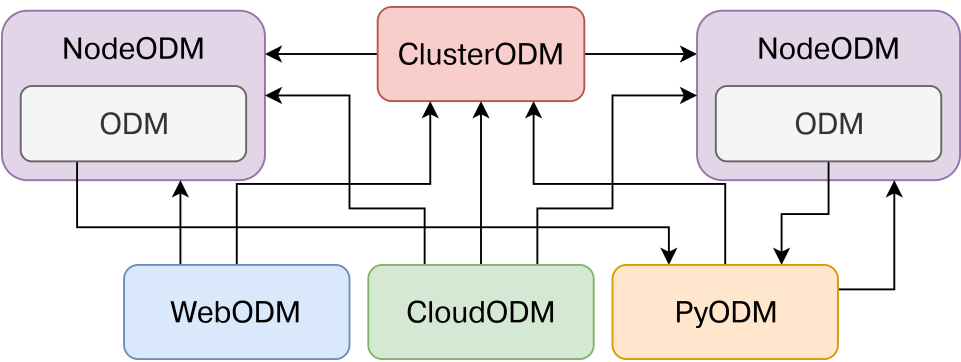
Doming is best cured by following best practices while collecting aerial imagery: flying at different elevations (maximize scale variation) and varying angles.

Unfortunately the luxury of capturing perfect images is not always available. Perhaps a dataset has already been captured and there's no opportunity for a retake,

18. The NodeODM API

ODM is a processing engine and WebODM is a friendly user interface. NodeODM¹ was historically built to allow WebODM to communicate with ODM over a network. Today NodeODM has expanded its role and is the glue that binds together many of OpenDroneMap projects. Each project understands the API that NodeODM defines. When we say *NodeODM* we are referring to the reference implementation of the NodeODM API available at github.com/OpenDroneMap/NodeODM.

At its core, the API defines ways to easily create new tasks, manage such tasks (cancel, delete, restart), download results and query status information.



Many OpenDroneMap projects use the NodeODM API to communicate with each other

¹Node is a reference to Node.js, the language NodeODM is written in

Glossary

2.5D Model: A model where elevation is simply *extruded* from the ground plane and thus is not a true 3D model.

Artifacts: undesired alterations generated as the result of a digital process.

API: Application Programming Interface. A set of functions allowing the creation of applications that access the features or data of another application.

Bundle Adjustment: a refinement step during the Structure From Motion process that improves the location of cameras, the 3D points of the scene and the camera parameters.

CloudODM: A command line tool to process aerial imagery in the cloud.

ClusterODM: A NodeODM API compatible autoscalable load balancer and task tracker for connecting multiple NodeODM nodes under a single network address.

CRS: Coordinate Reference System. A CRS is a coordinate-based system used to locate geographical entities.

CSV: Comma Separated Value is a textual file format where fields are typically separated by commas or some other character such as a space or a tab.

cURL: a software providing a library and command-line tool for transferring data using many protocols.

DEM: Digital Elevation Model (either a DSM or a DTM).