Capturing Reality

Images contain more than just colors

3D Reconstruction in the Wild

Martin Bujnak / martin@capturingreality.com

3D reconstructions

What are the challenges today?

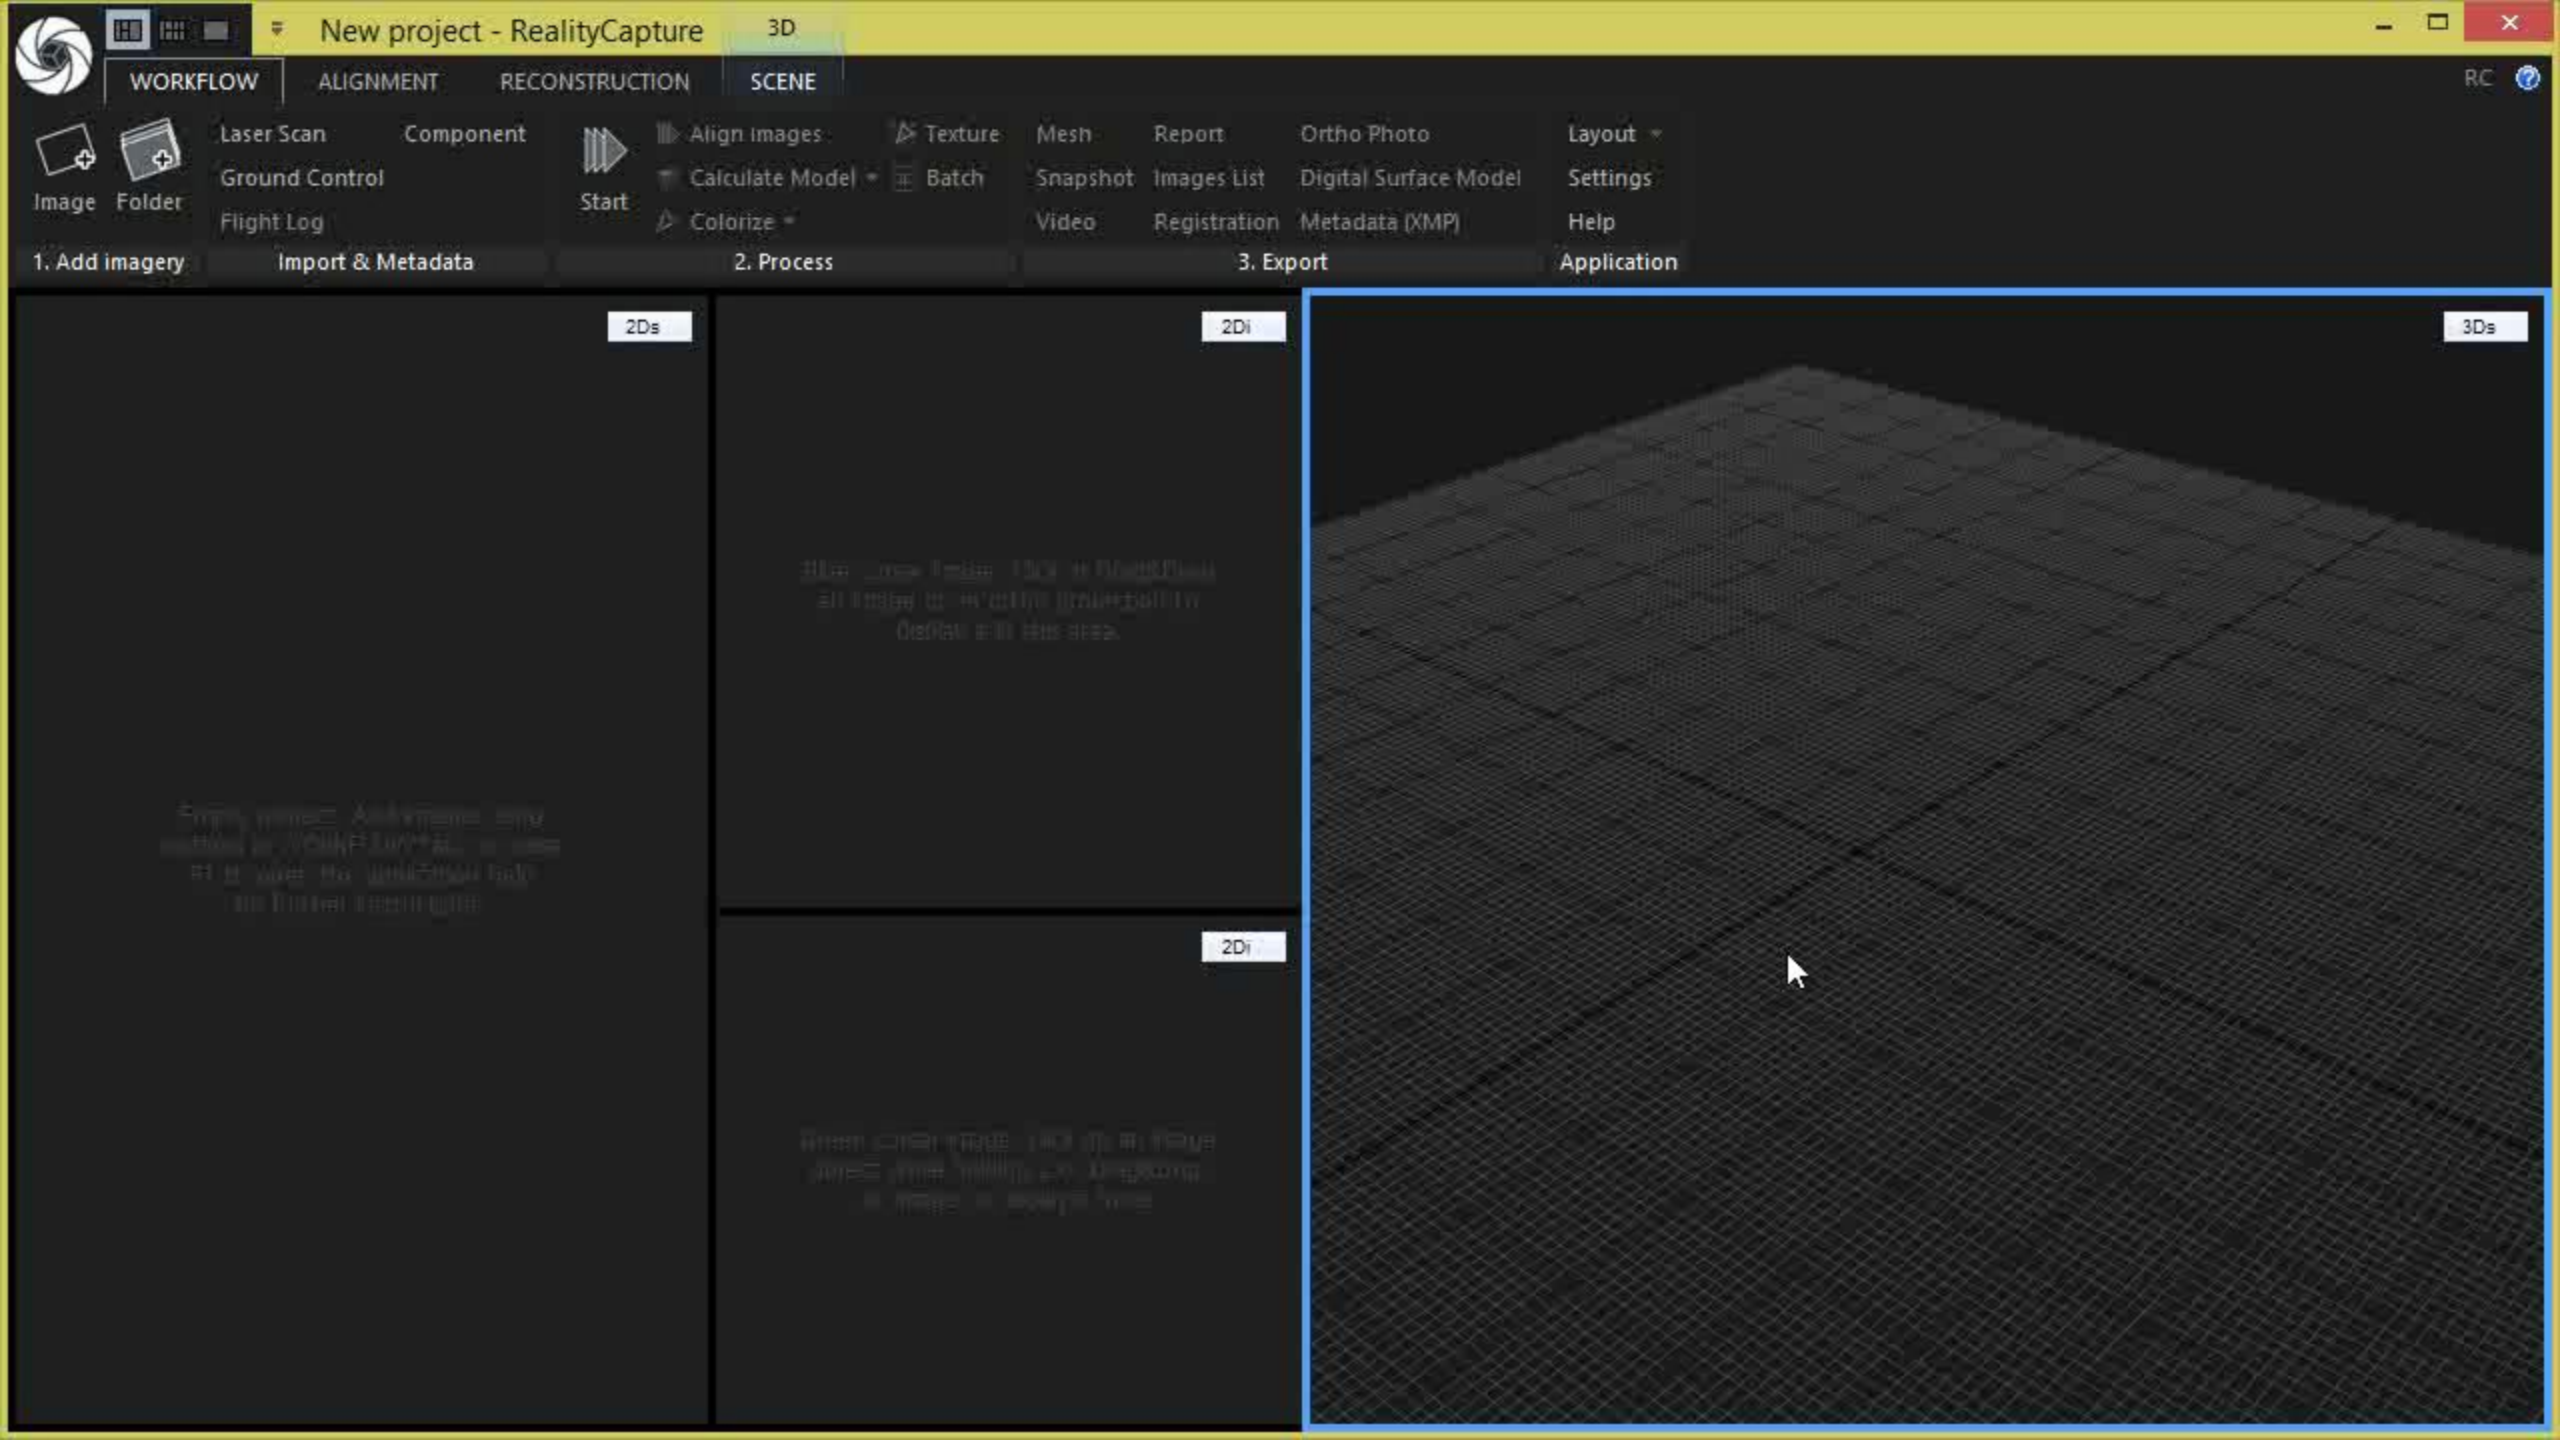
New project - RealityCapture
3D
WORKFLOW
ALIGNMENT
RECONSTRUCTION
SCENE
RC
Image
Folder
Laser Scan
Ground Control
Flight Log
Component
Start
Align images
Calculate Model
Colorize
Texture
Batch
Mesh
Snapshot
Video
Report
Images List
Registration
Ortho Photo
Digital Surface Model
Metadata (XMP)
Layout
Settings
Help
1. Add imagery
Import & Metadata
2. Process
3. Export
Application
2Ds
2Di
2Di
3Ds

Wow! it's a turn table, critical motion, un-calibrated, hmm, yeah

Even your grandma will not say it is cool today

Solved 20yrs ago

Much more detail, more complicated, shiny material, …

General camera setup/configuration/autocalibration

Solved 10yrs ago *

* for a general, small, collection of photos

High resolution cameras / 80mpx / can see small scratches and clever algorithm can hypothesize and recover the correct surface

Not to many people have such cams today

Solved for industry apps*

People do not want to take 3000 images to scan such scenes

* special sprays, structured light

3D Maps…

Is it a challenging problem today?

Depends on the expected detail

Use GPS to cut the space into sectors

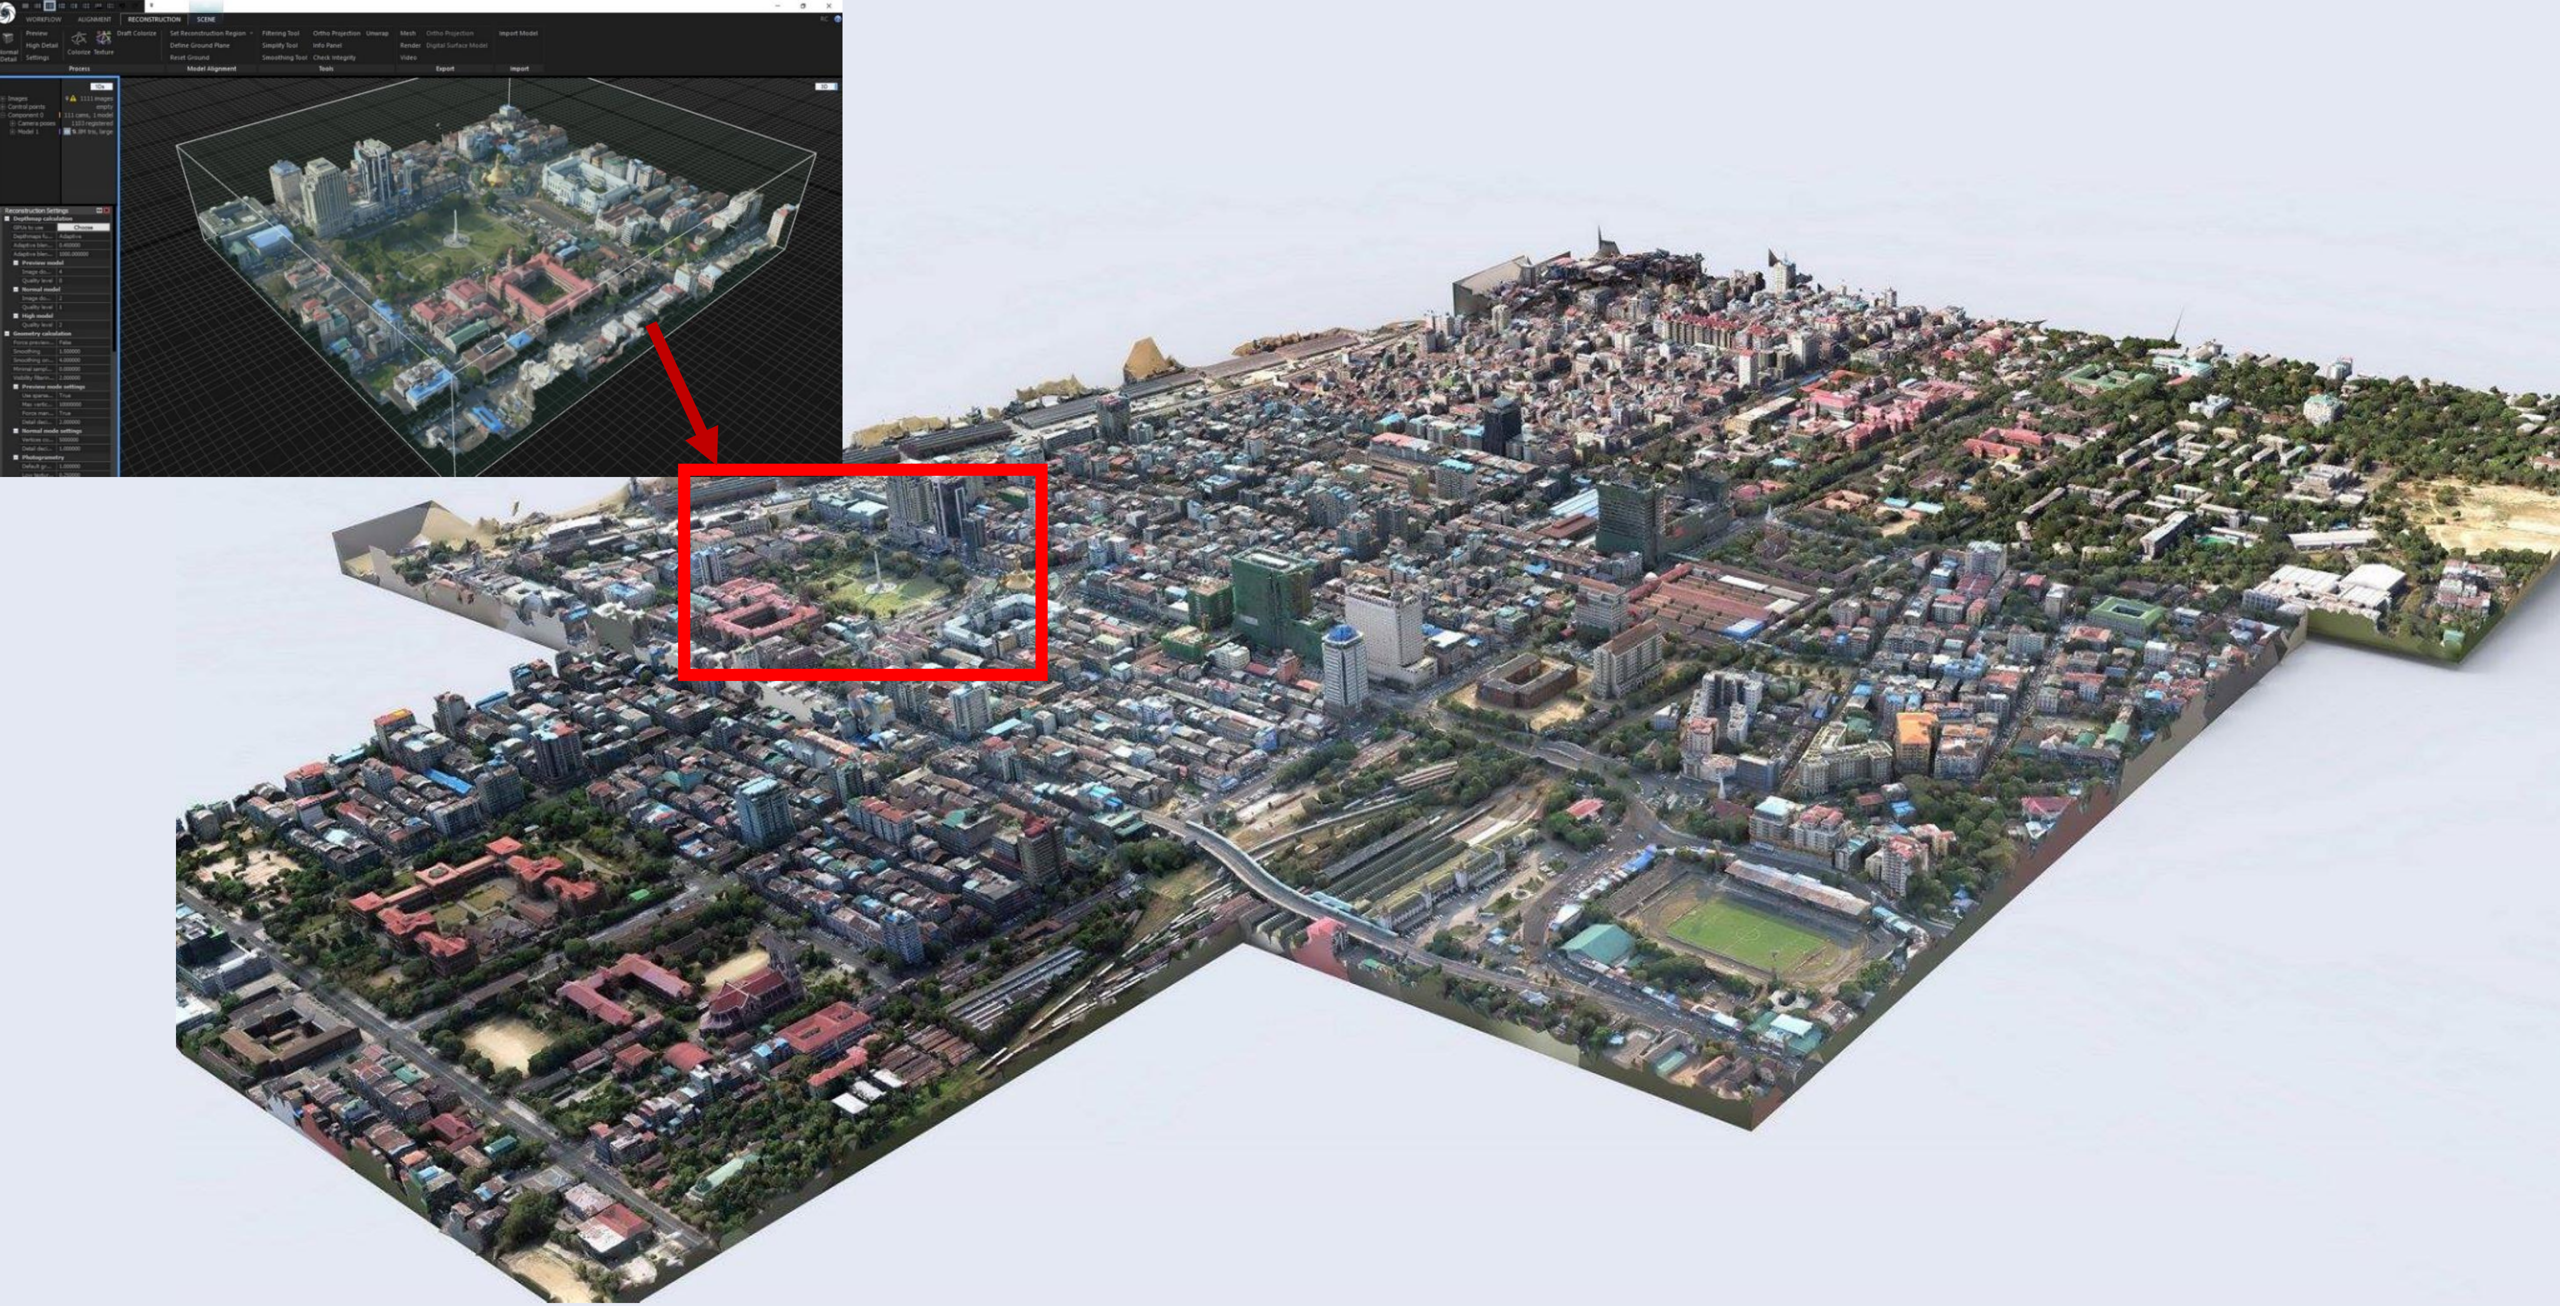

www.capturingreality.com / martin@capturingreality.com

How about - many photos & without GPS?

Much more complicated

Mostly because of a difficult question: Which images see the same physical 3D space? (the matching problem)

… and speed & memory problems during camera registration

“Rome” in a cloudless day (ECCV2010)

2.8Mil photos, low resolution

It is more “how to get some models from a huge set of images”

Wow for the research community

Detail like in computer games from 1995 to “normal people”

So how far we are from “the ultimate goal“ ?

High quality reconstruction

Automatic processing

Processed in a reasonable time

Close to that!

Today we can easily create models from 100K+ images

On a single machine

Limited computer resources (memory, hdd)

100,000-image project

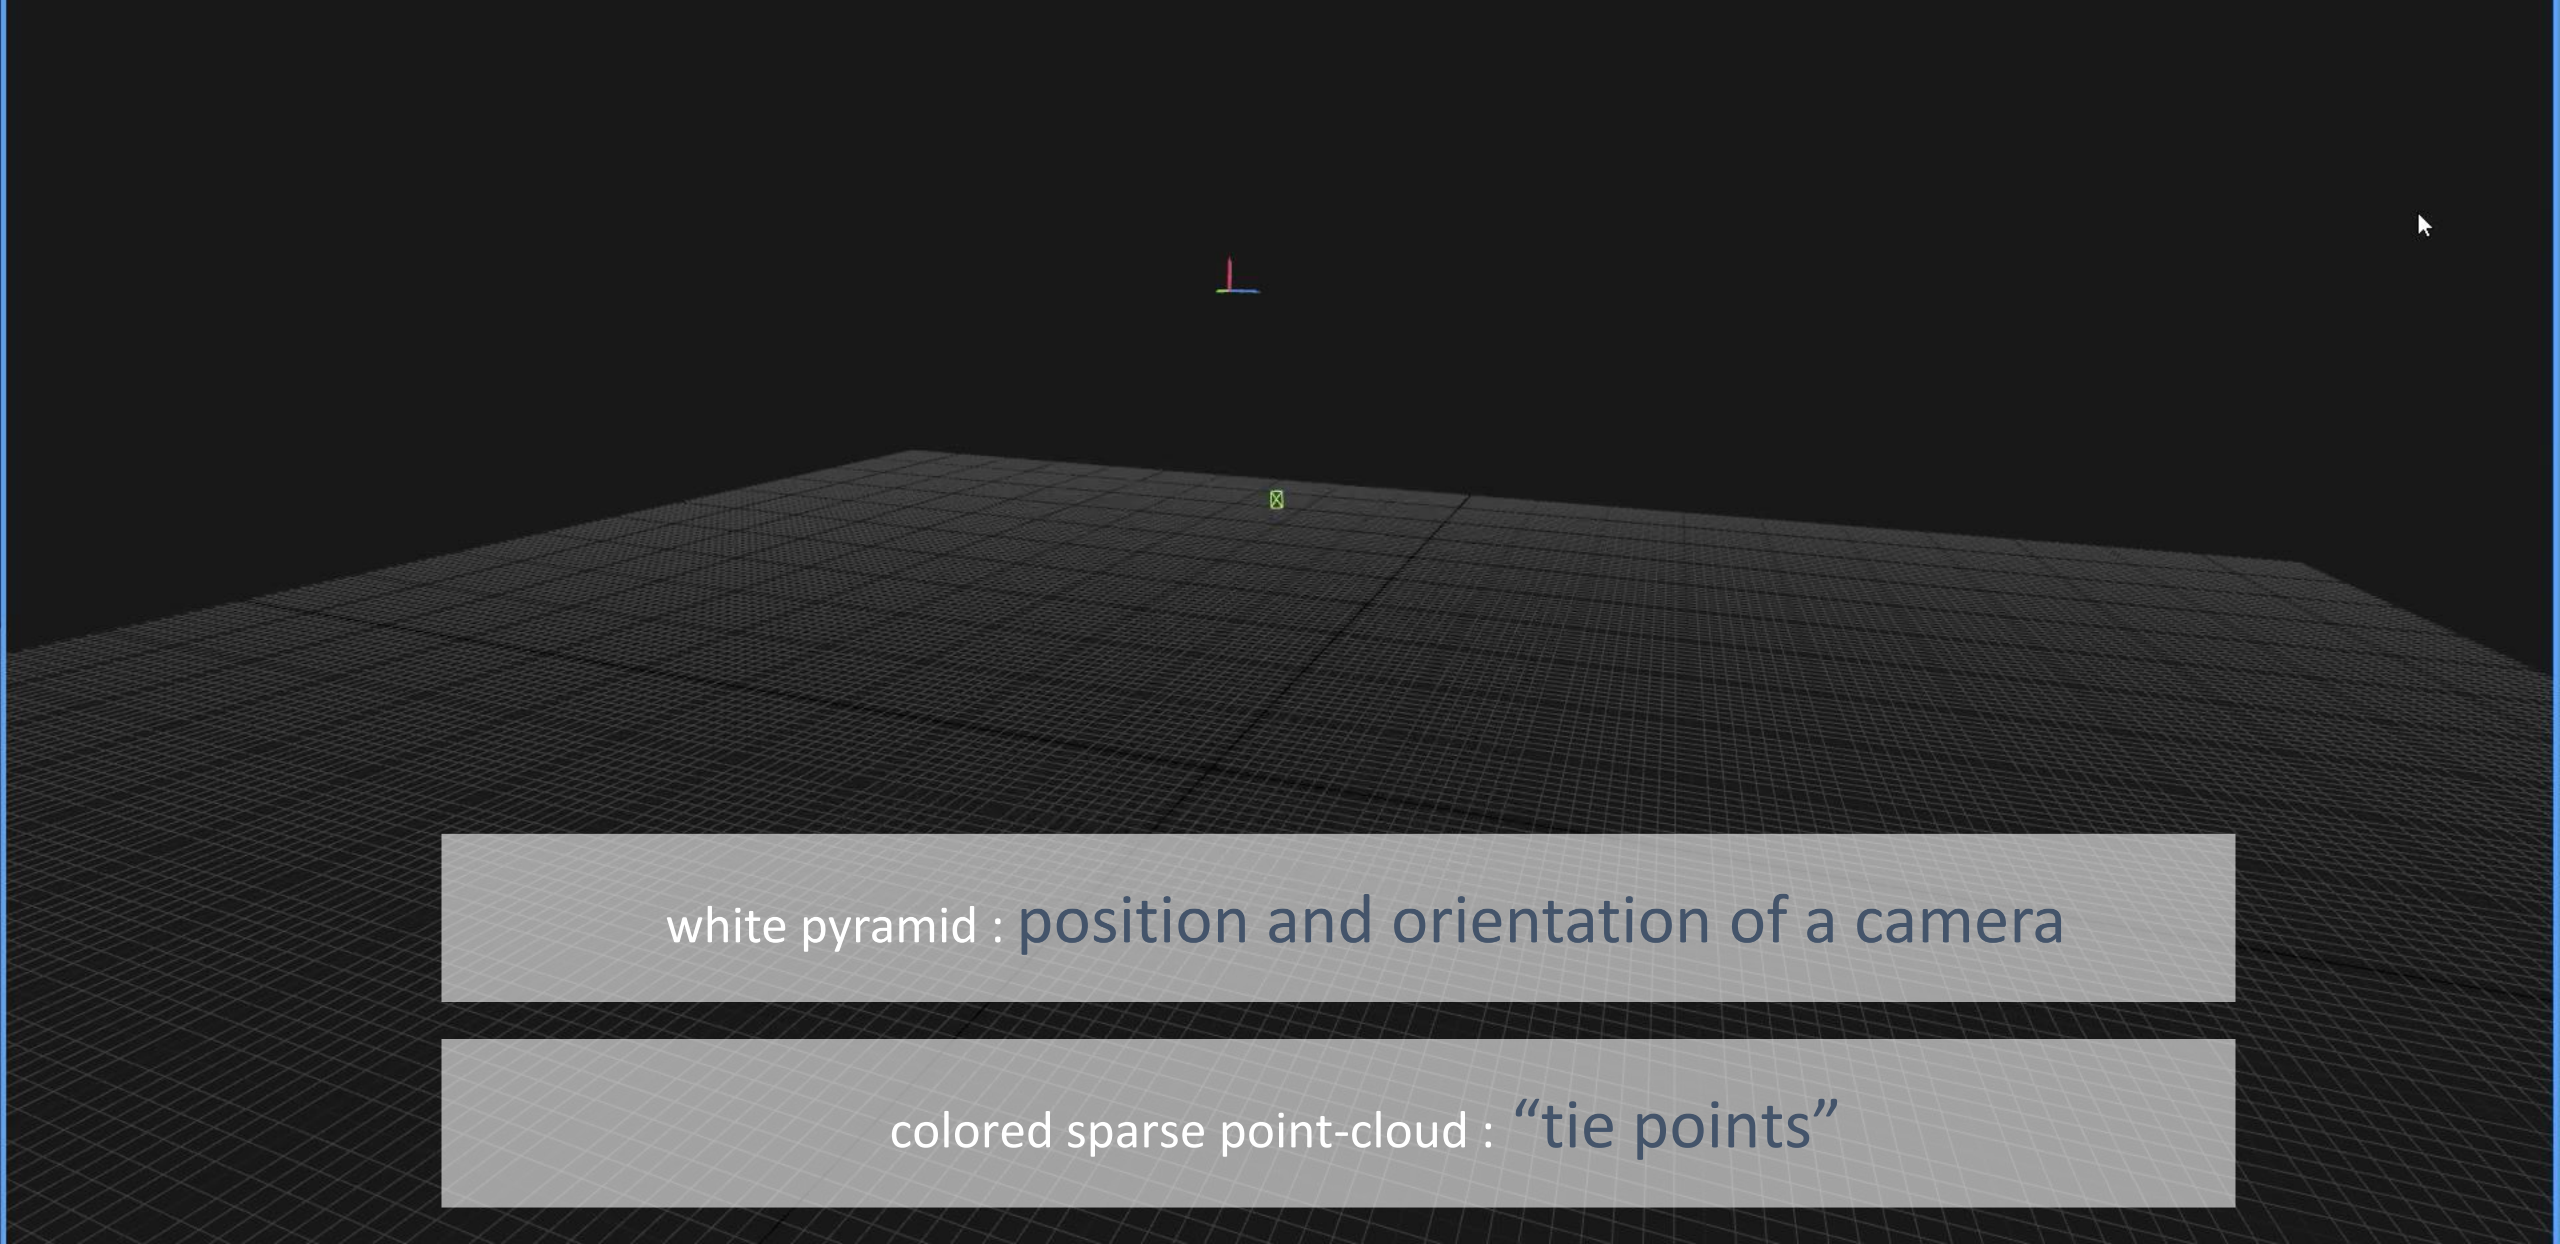

white pyramid : position and orientation of a camera

colored sparse point-cloud : “tie points”

100,000-images project : image samples

100,000-image project : final alignment

100,000-images project : just tower selection

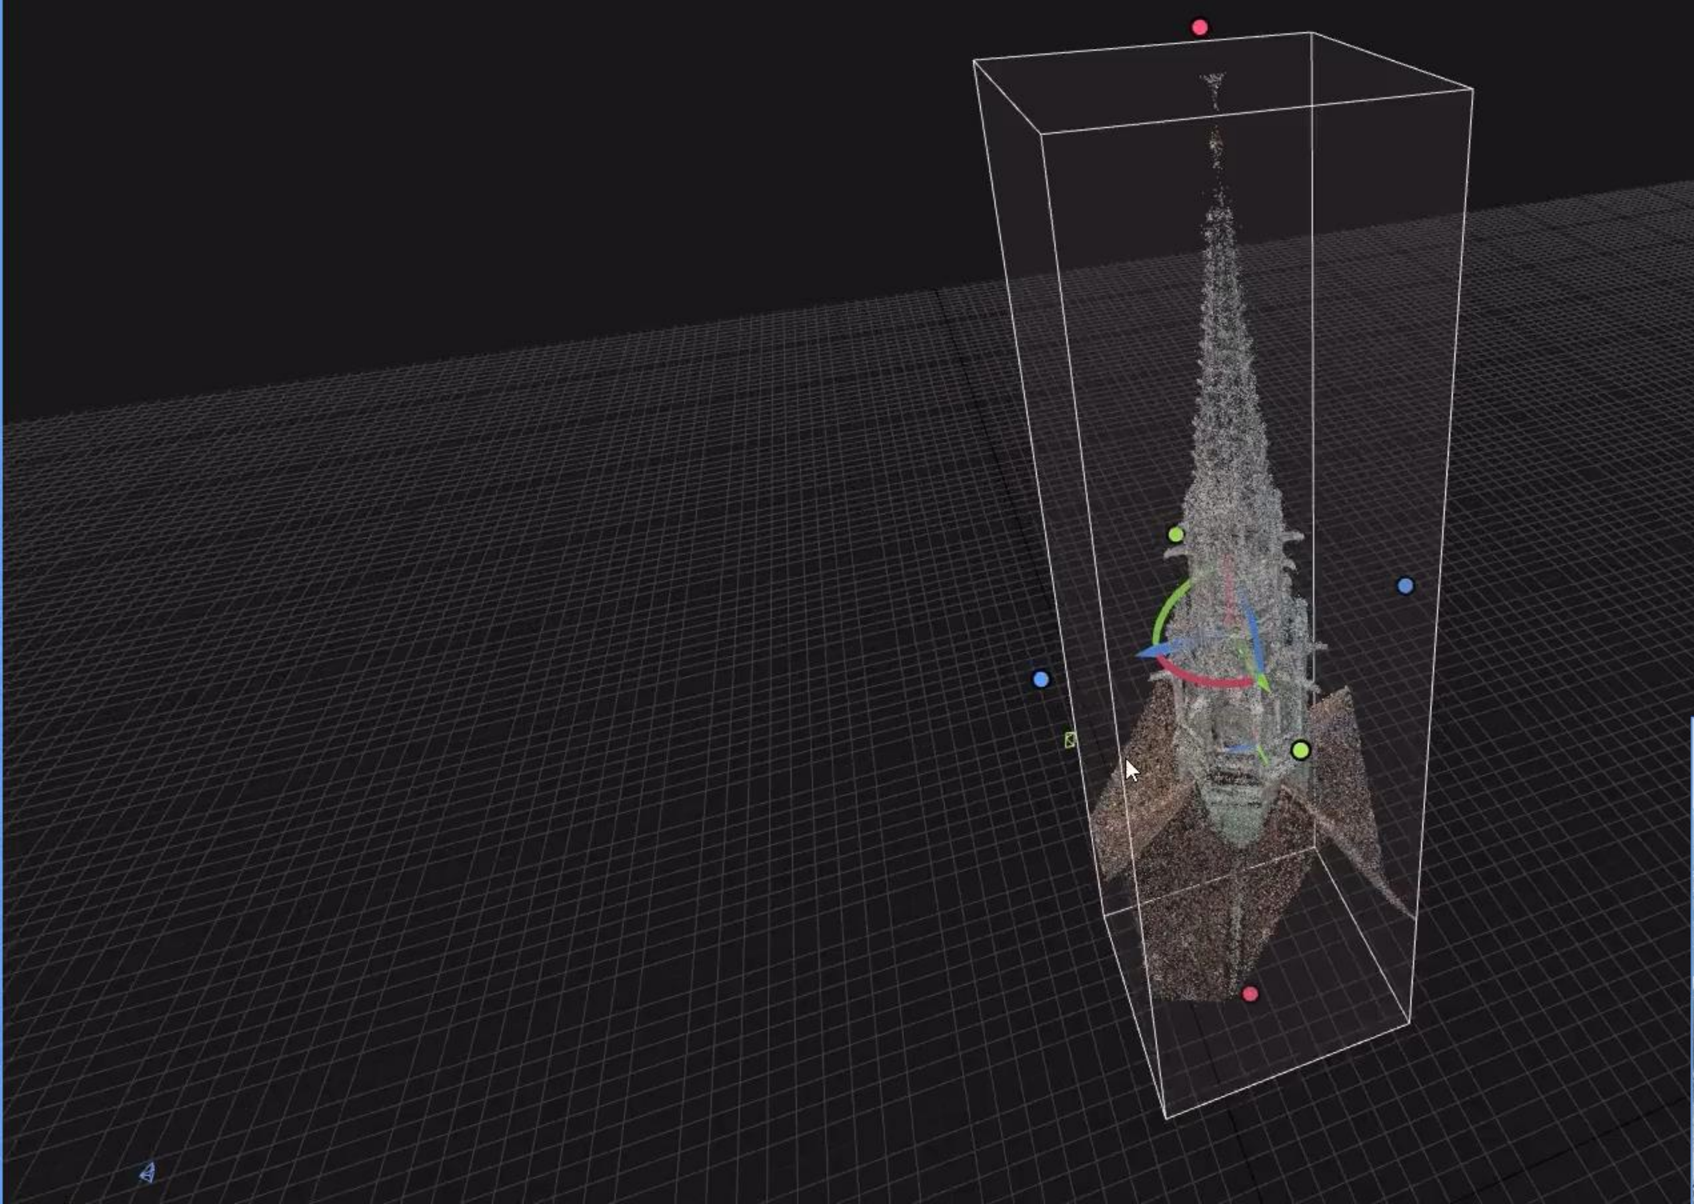

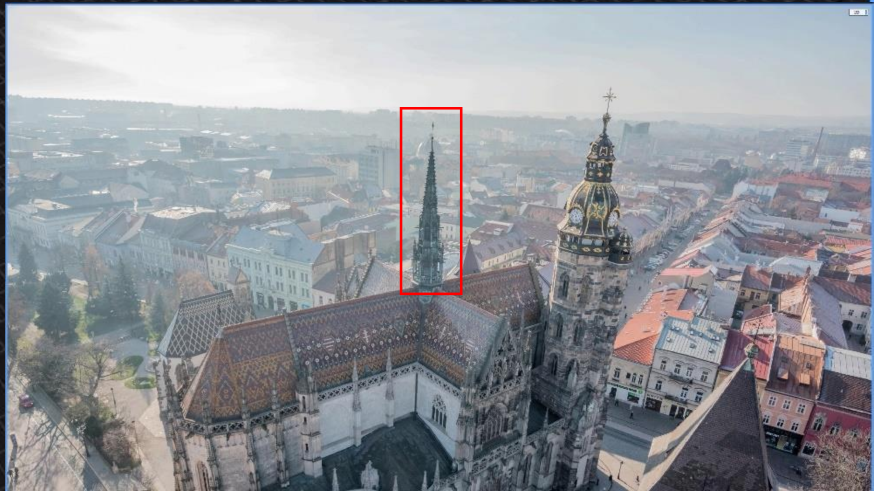

100,000-images project : model of the tower

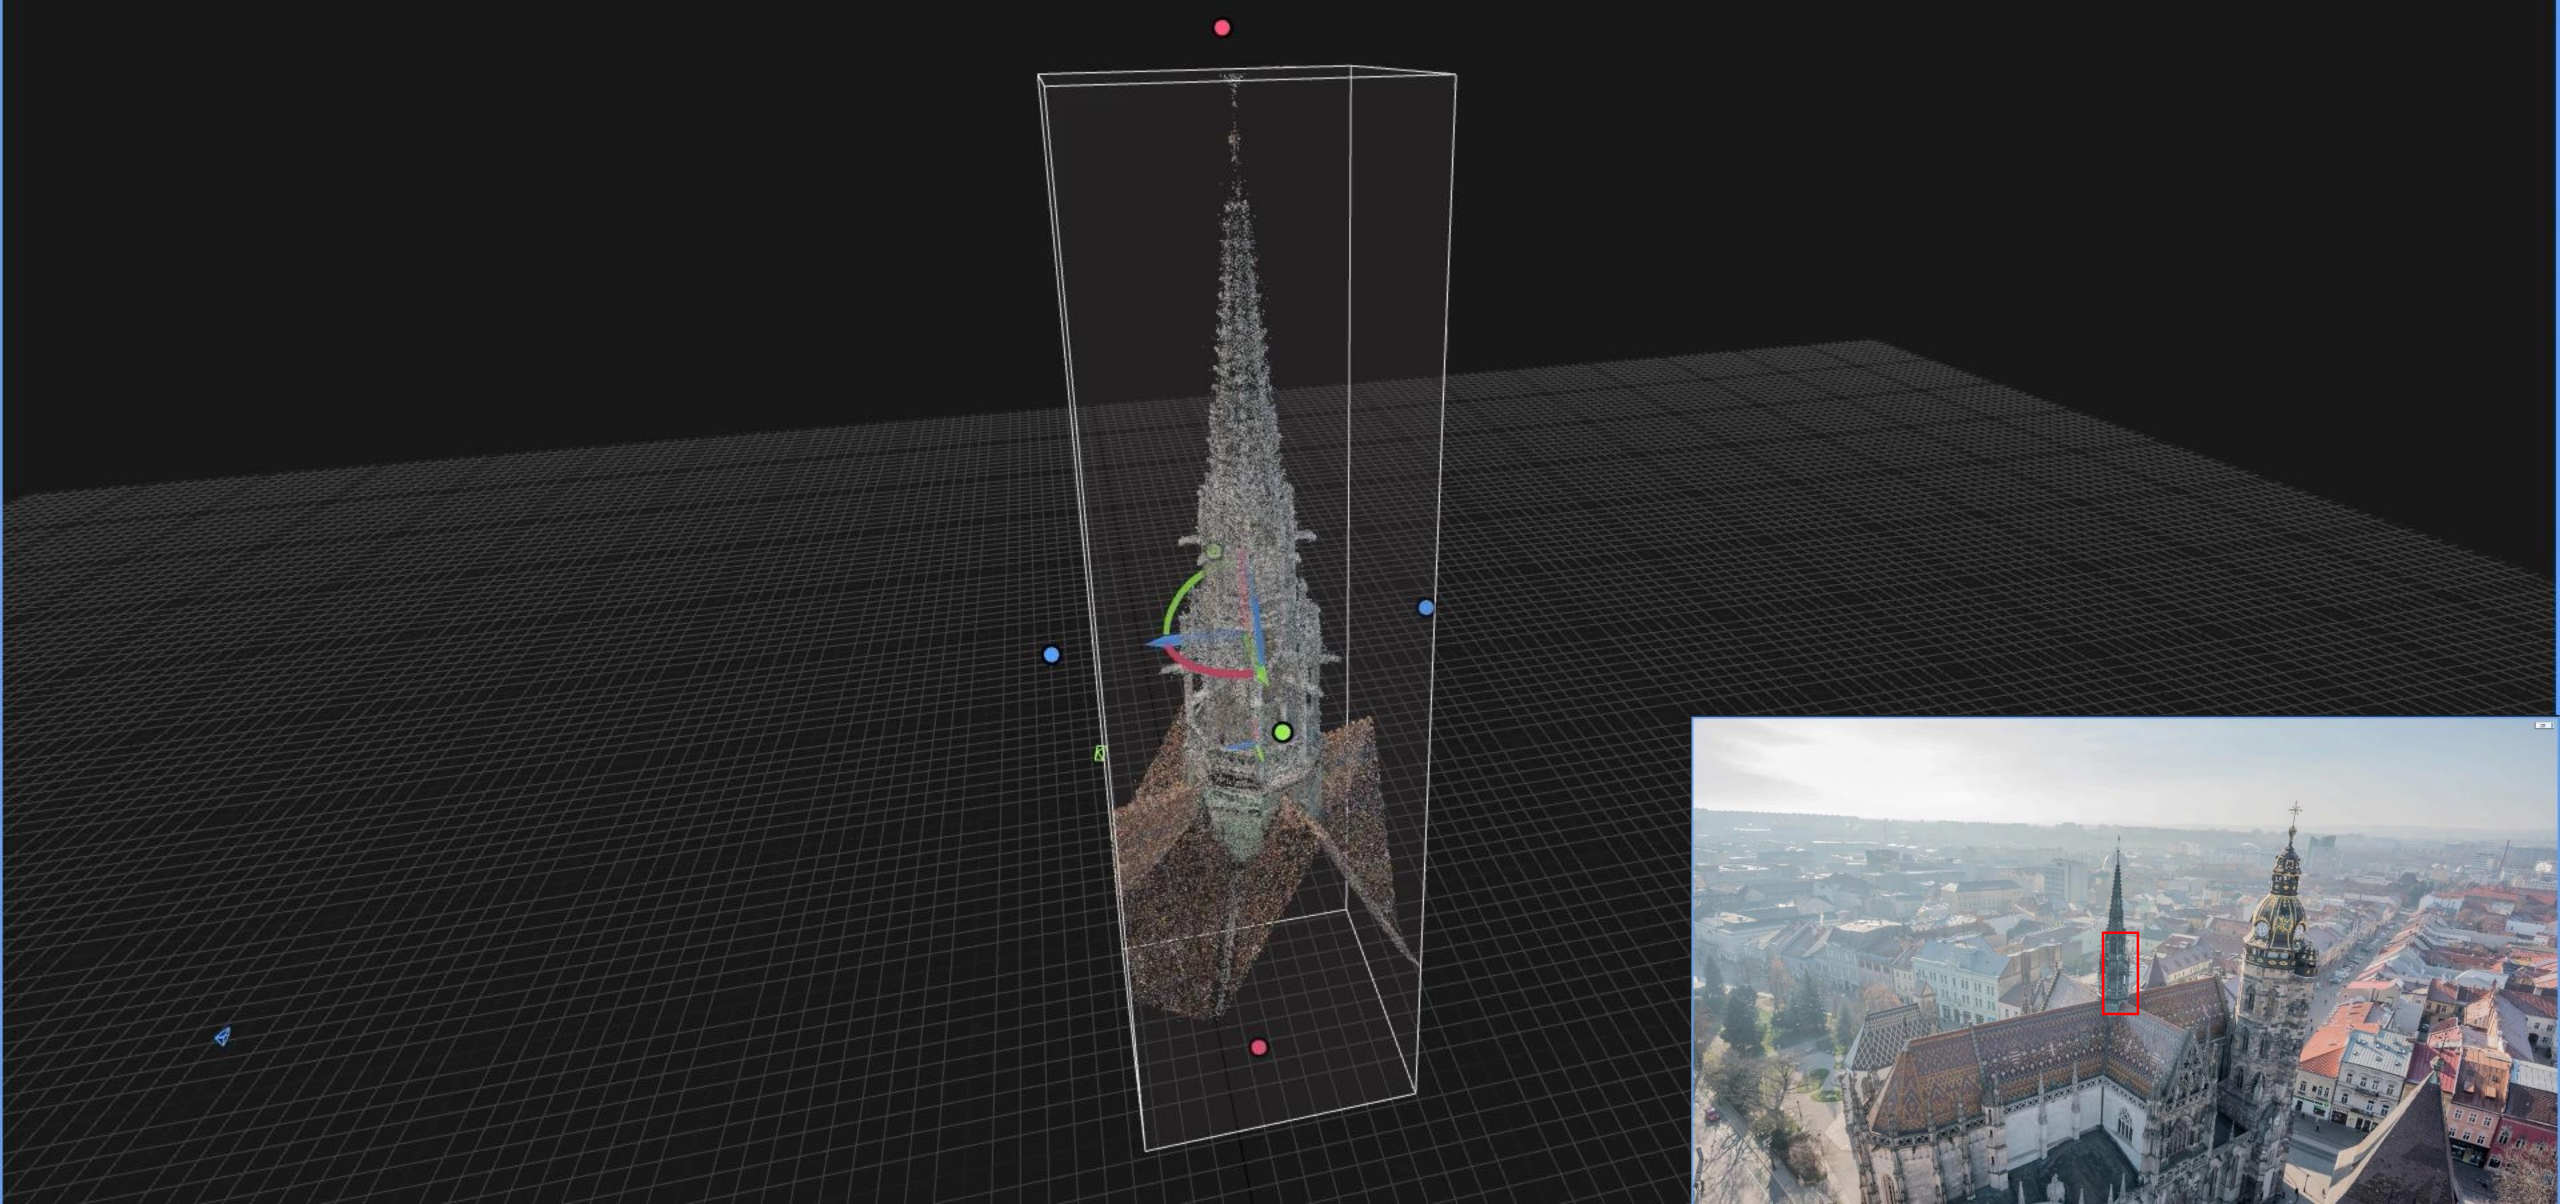

100,000-images project : model of the pillar

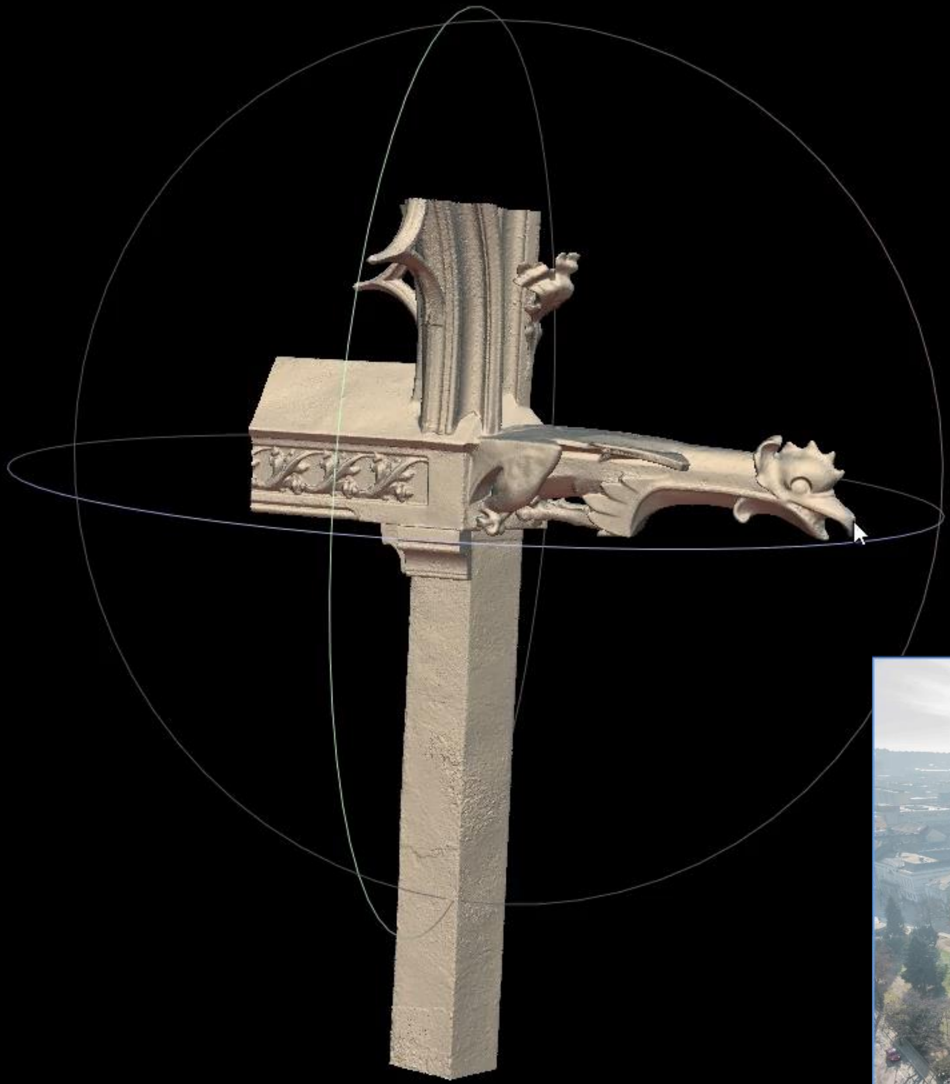

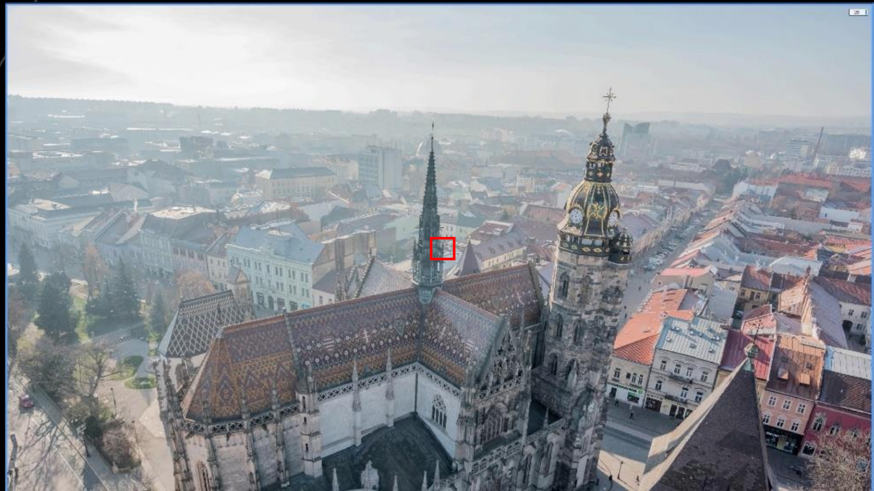

This is RealityCapture which we are developing

Let’s return back to the beginning

How it works?

Small dataset

200 images

captured in 39 minutes

Processing times

alignment in 4 minutes

model in 42 minutes

PC used

ordinary gaming PC
16 GB RAM, i7 CPU, GTX 960

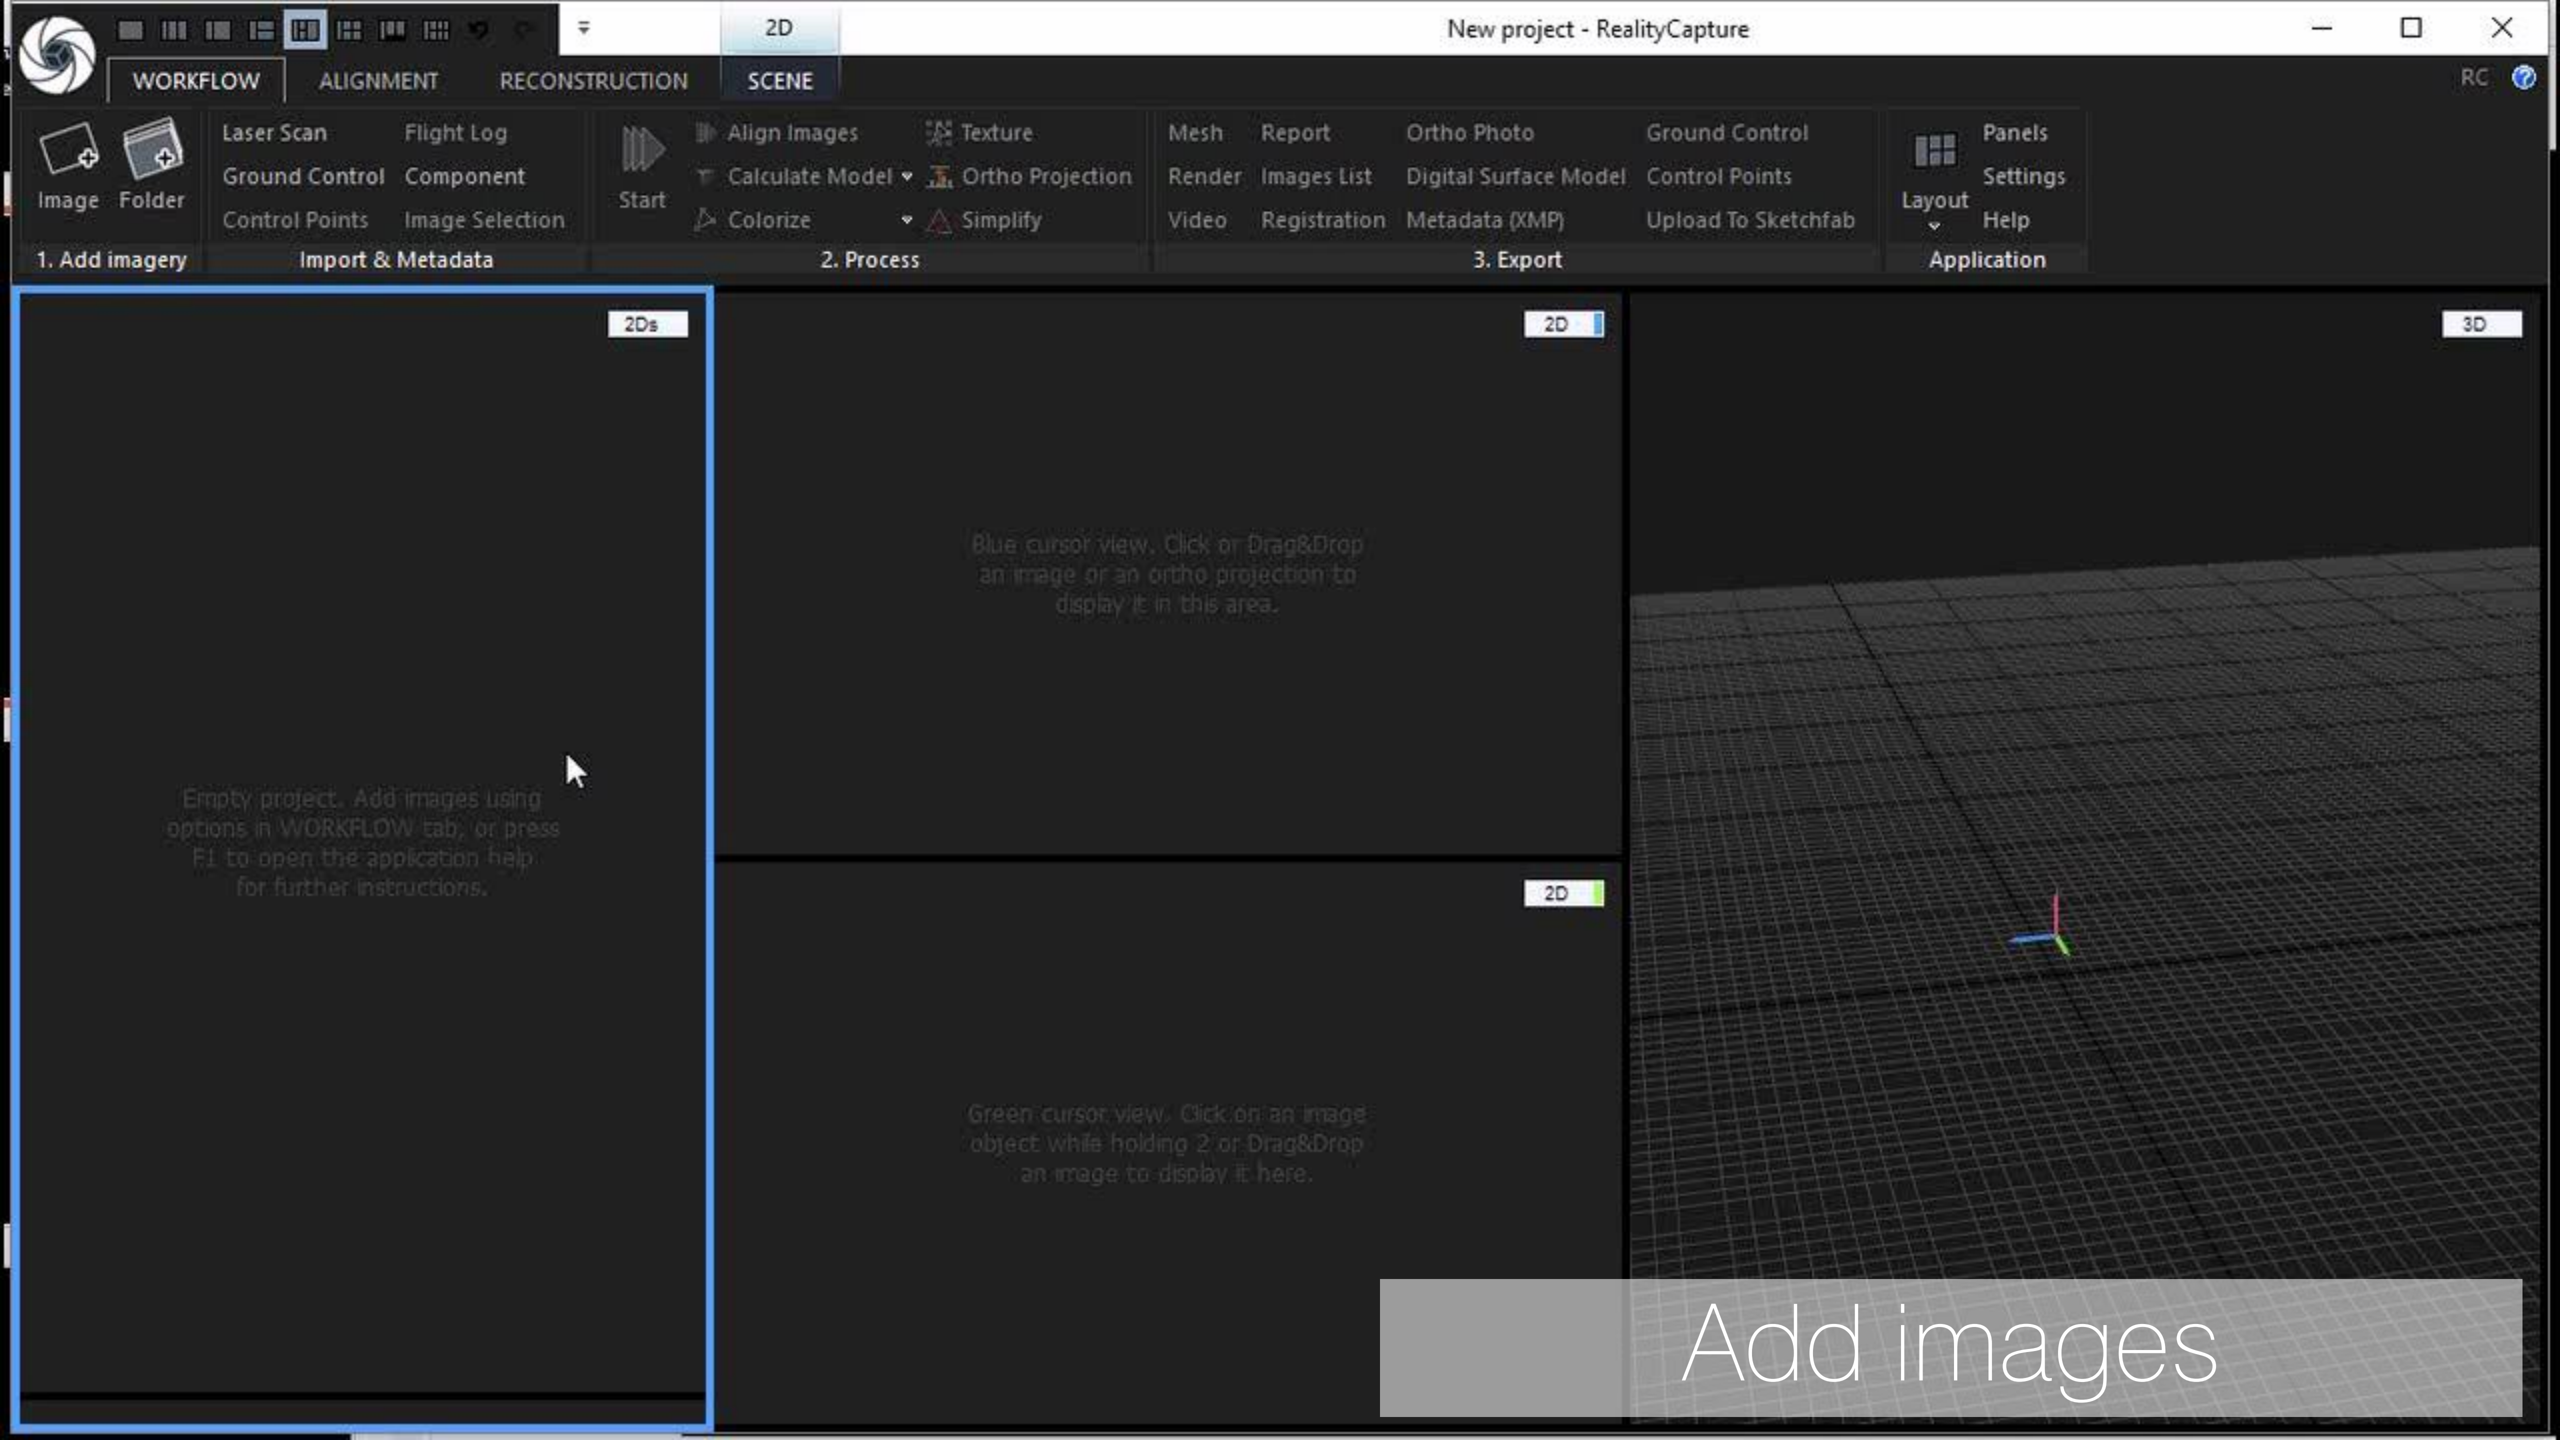
2D
New project - RealityCapture
WORKFLOW
ALIGNMENT
RECONSTRUCTION
SCENE
RC
Image
Folder
Laser Scan
Ground Control
Control Points
Flight Log
Component
Image Selection
Start
Align Images
Calculate Model
Colorize
Texture
Ortho Projection
Simplify
Mesh
Render
Video
Report
Images List
Registration
Ortho Photo
Digital Surface Model
Metadata (XMP)
Ground Control
Control Points
Upload To Sketchfab
Layout
Panels
Settings
Help
1. Add imagery
Import & Metadata
2. Process
3. Export
Application
2Ds
2D
3D
Blue cursor view. Click or Drag&Drop an image or an ortho projection to display it in this area.
Empty project. Add images using options in WORKFLOW tab, or press F1 to open the application help for further instructions.
2D
Green cursor view. Click on an image object while holding 2 or Drag&Drop an image to display it here.
Add images

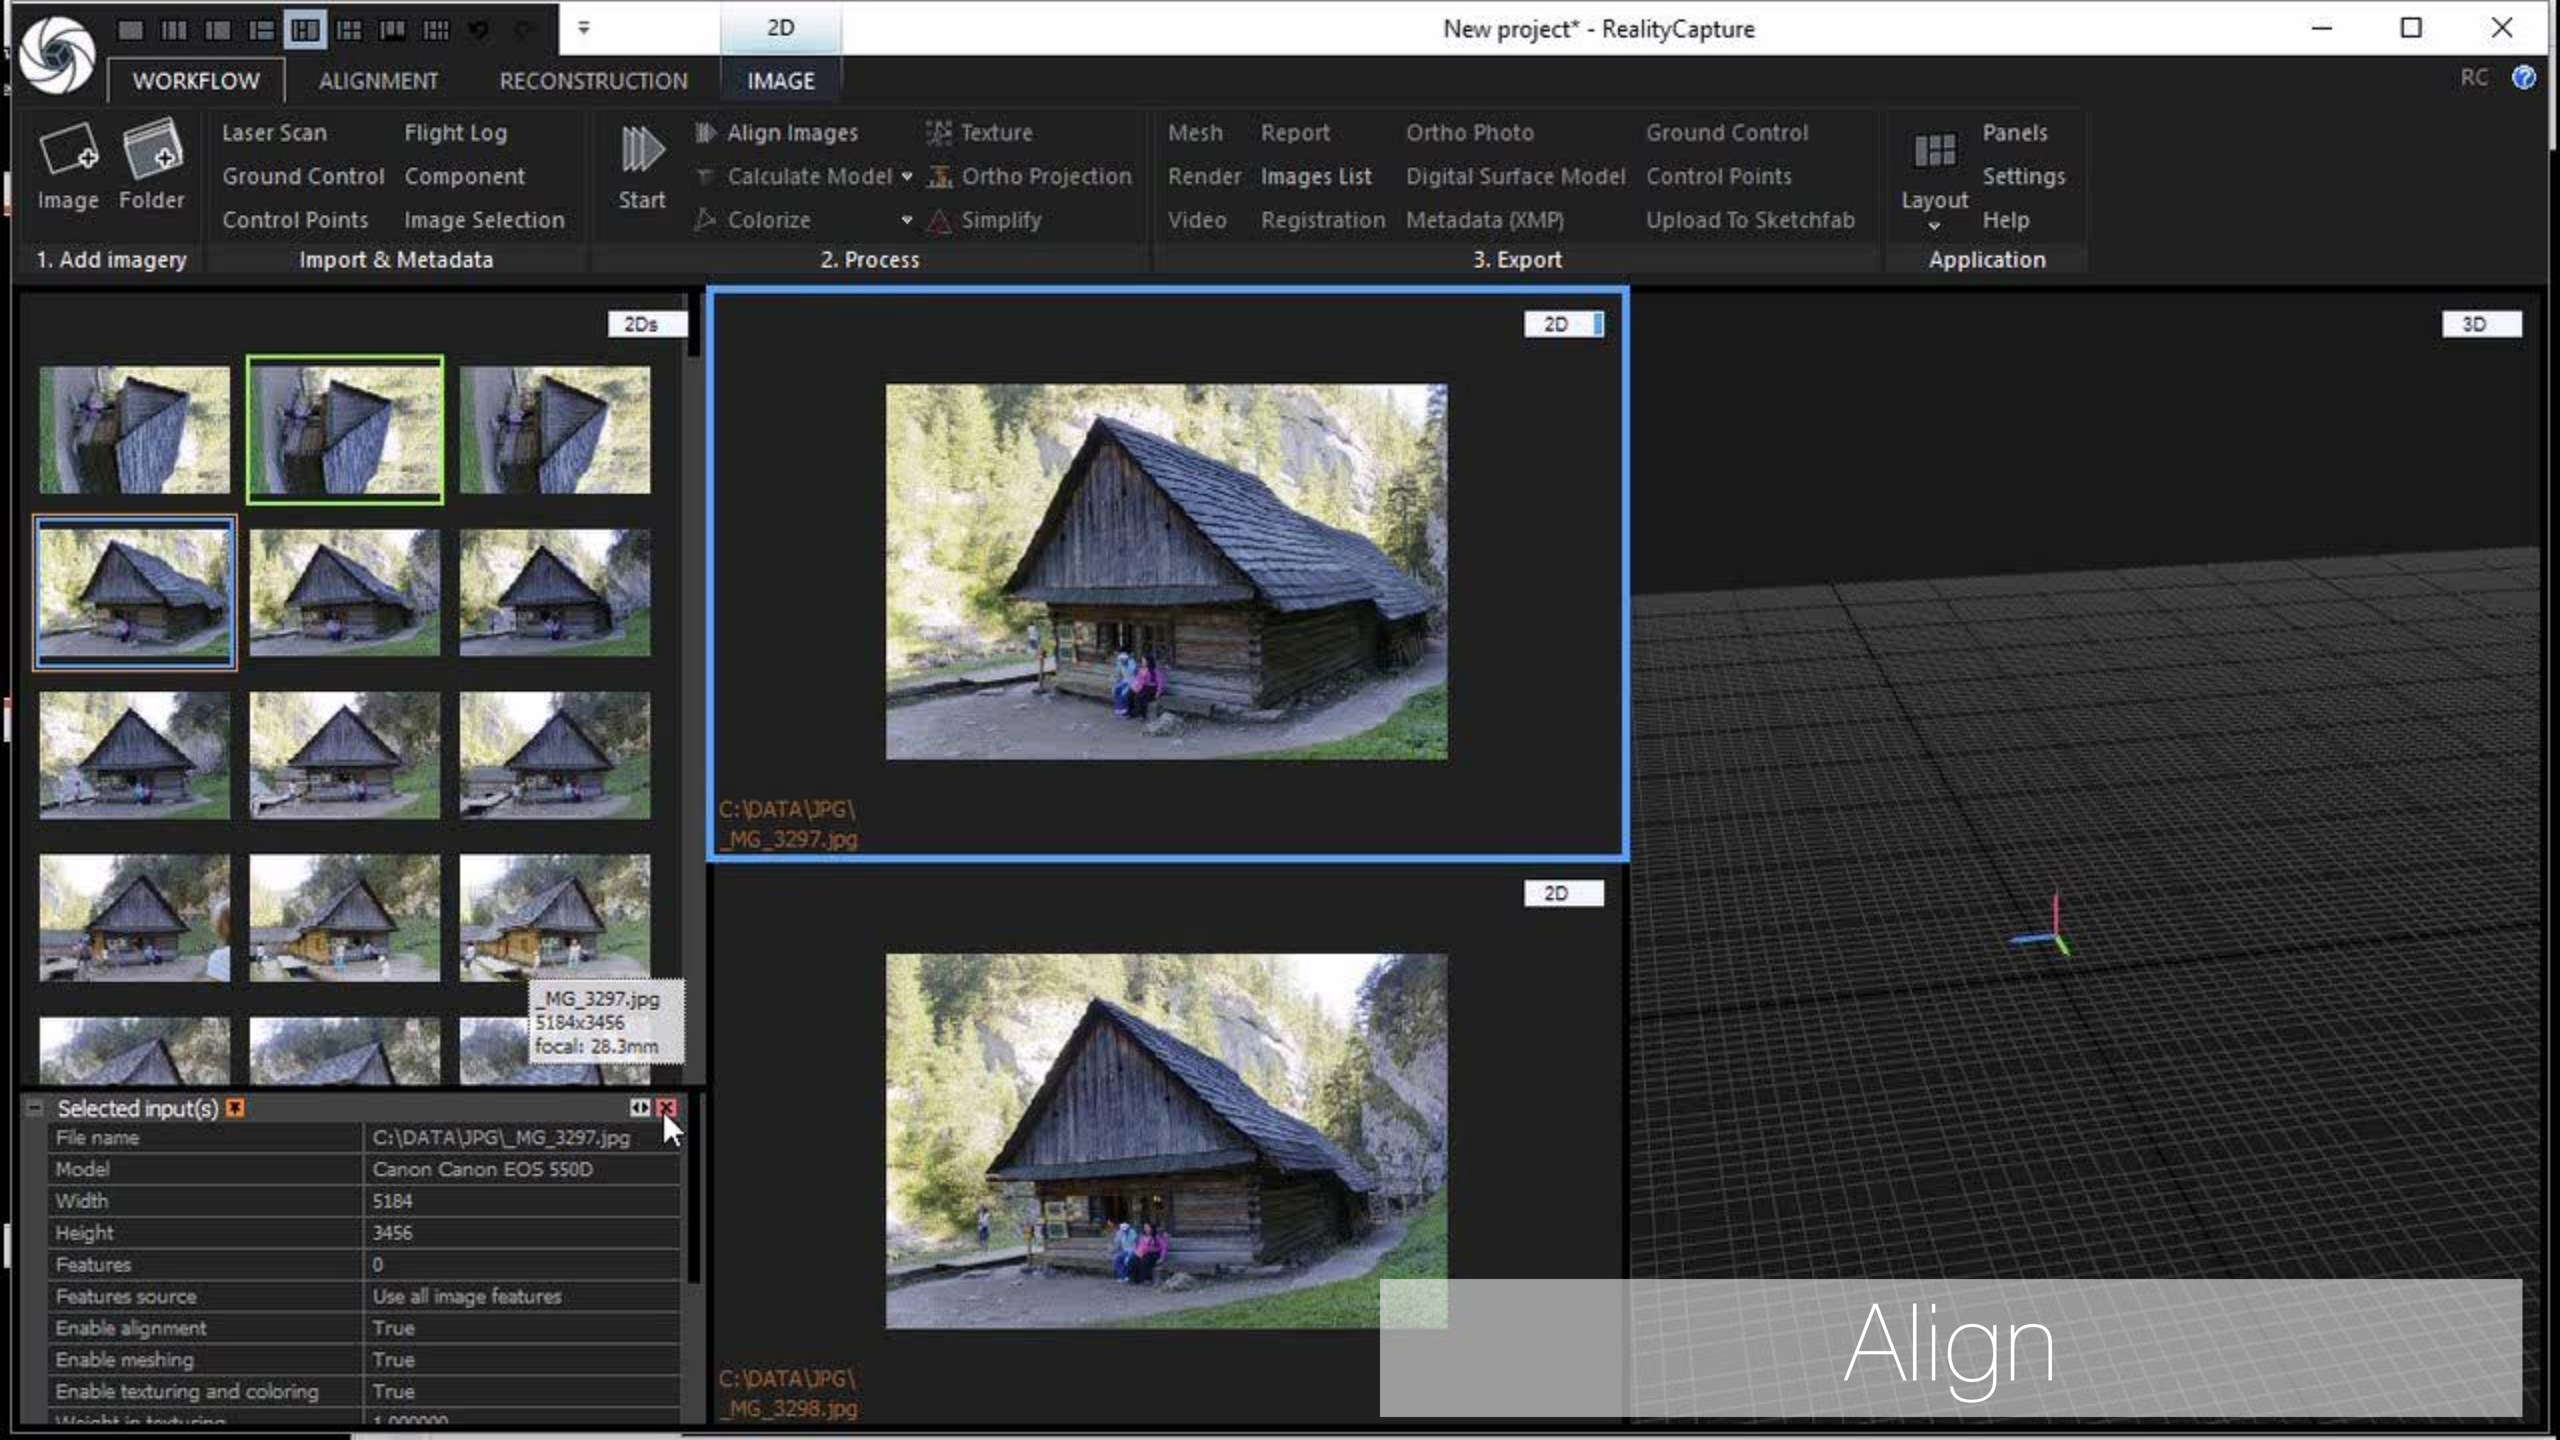
2D
New project* - RealityCapture
WORKFLOW
ALIGNMENT
RECONSTRUCTION
IMAGE
RC
Image
Folder
Laser Scan
Ground Control
Control Points
Flight Log
Component
Image Selection
Start
Align Images
Calculate Model
Colorize
Texture
Ortho Projection
Simplify
Mesh
Render
Video
Report
Images List
Registration
Ortho Photo
Digital Surface Model
Metadata (XMP)
Ground Control
Control Points
Upload To Sketchfab
Layout
Panels
Settings
Help
1. Add imagery
Import & Metadata
2. Process
3. Export
Application
2Ds
2D
3D
C:\DATA\JPG_MG_3297.jpg
_MG_3297.jpg
5184x3456
focal: 28.3mm
Selected input(s)
File name
C:\DATA\JPG_MG_3297.jpg
Model
Canon Canon EOS 550D
Width
5184
Height
3456
Features
0
Features source
Use all image features
Enable alignment
True
Enable meshing
True
Enable texturing and coloring
True
C:\DATA\JPG_MG_3298.jpg
Align

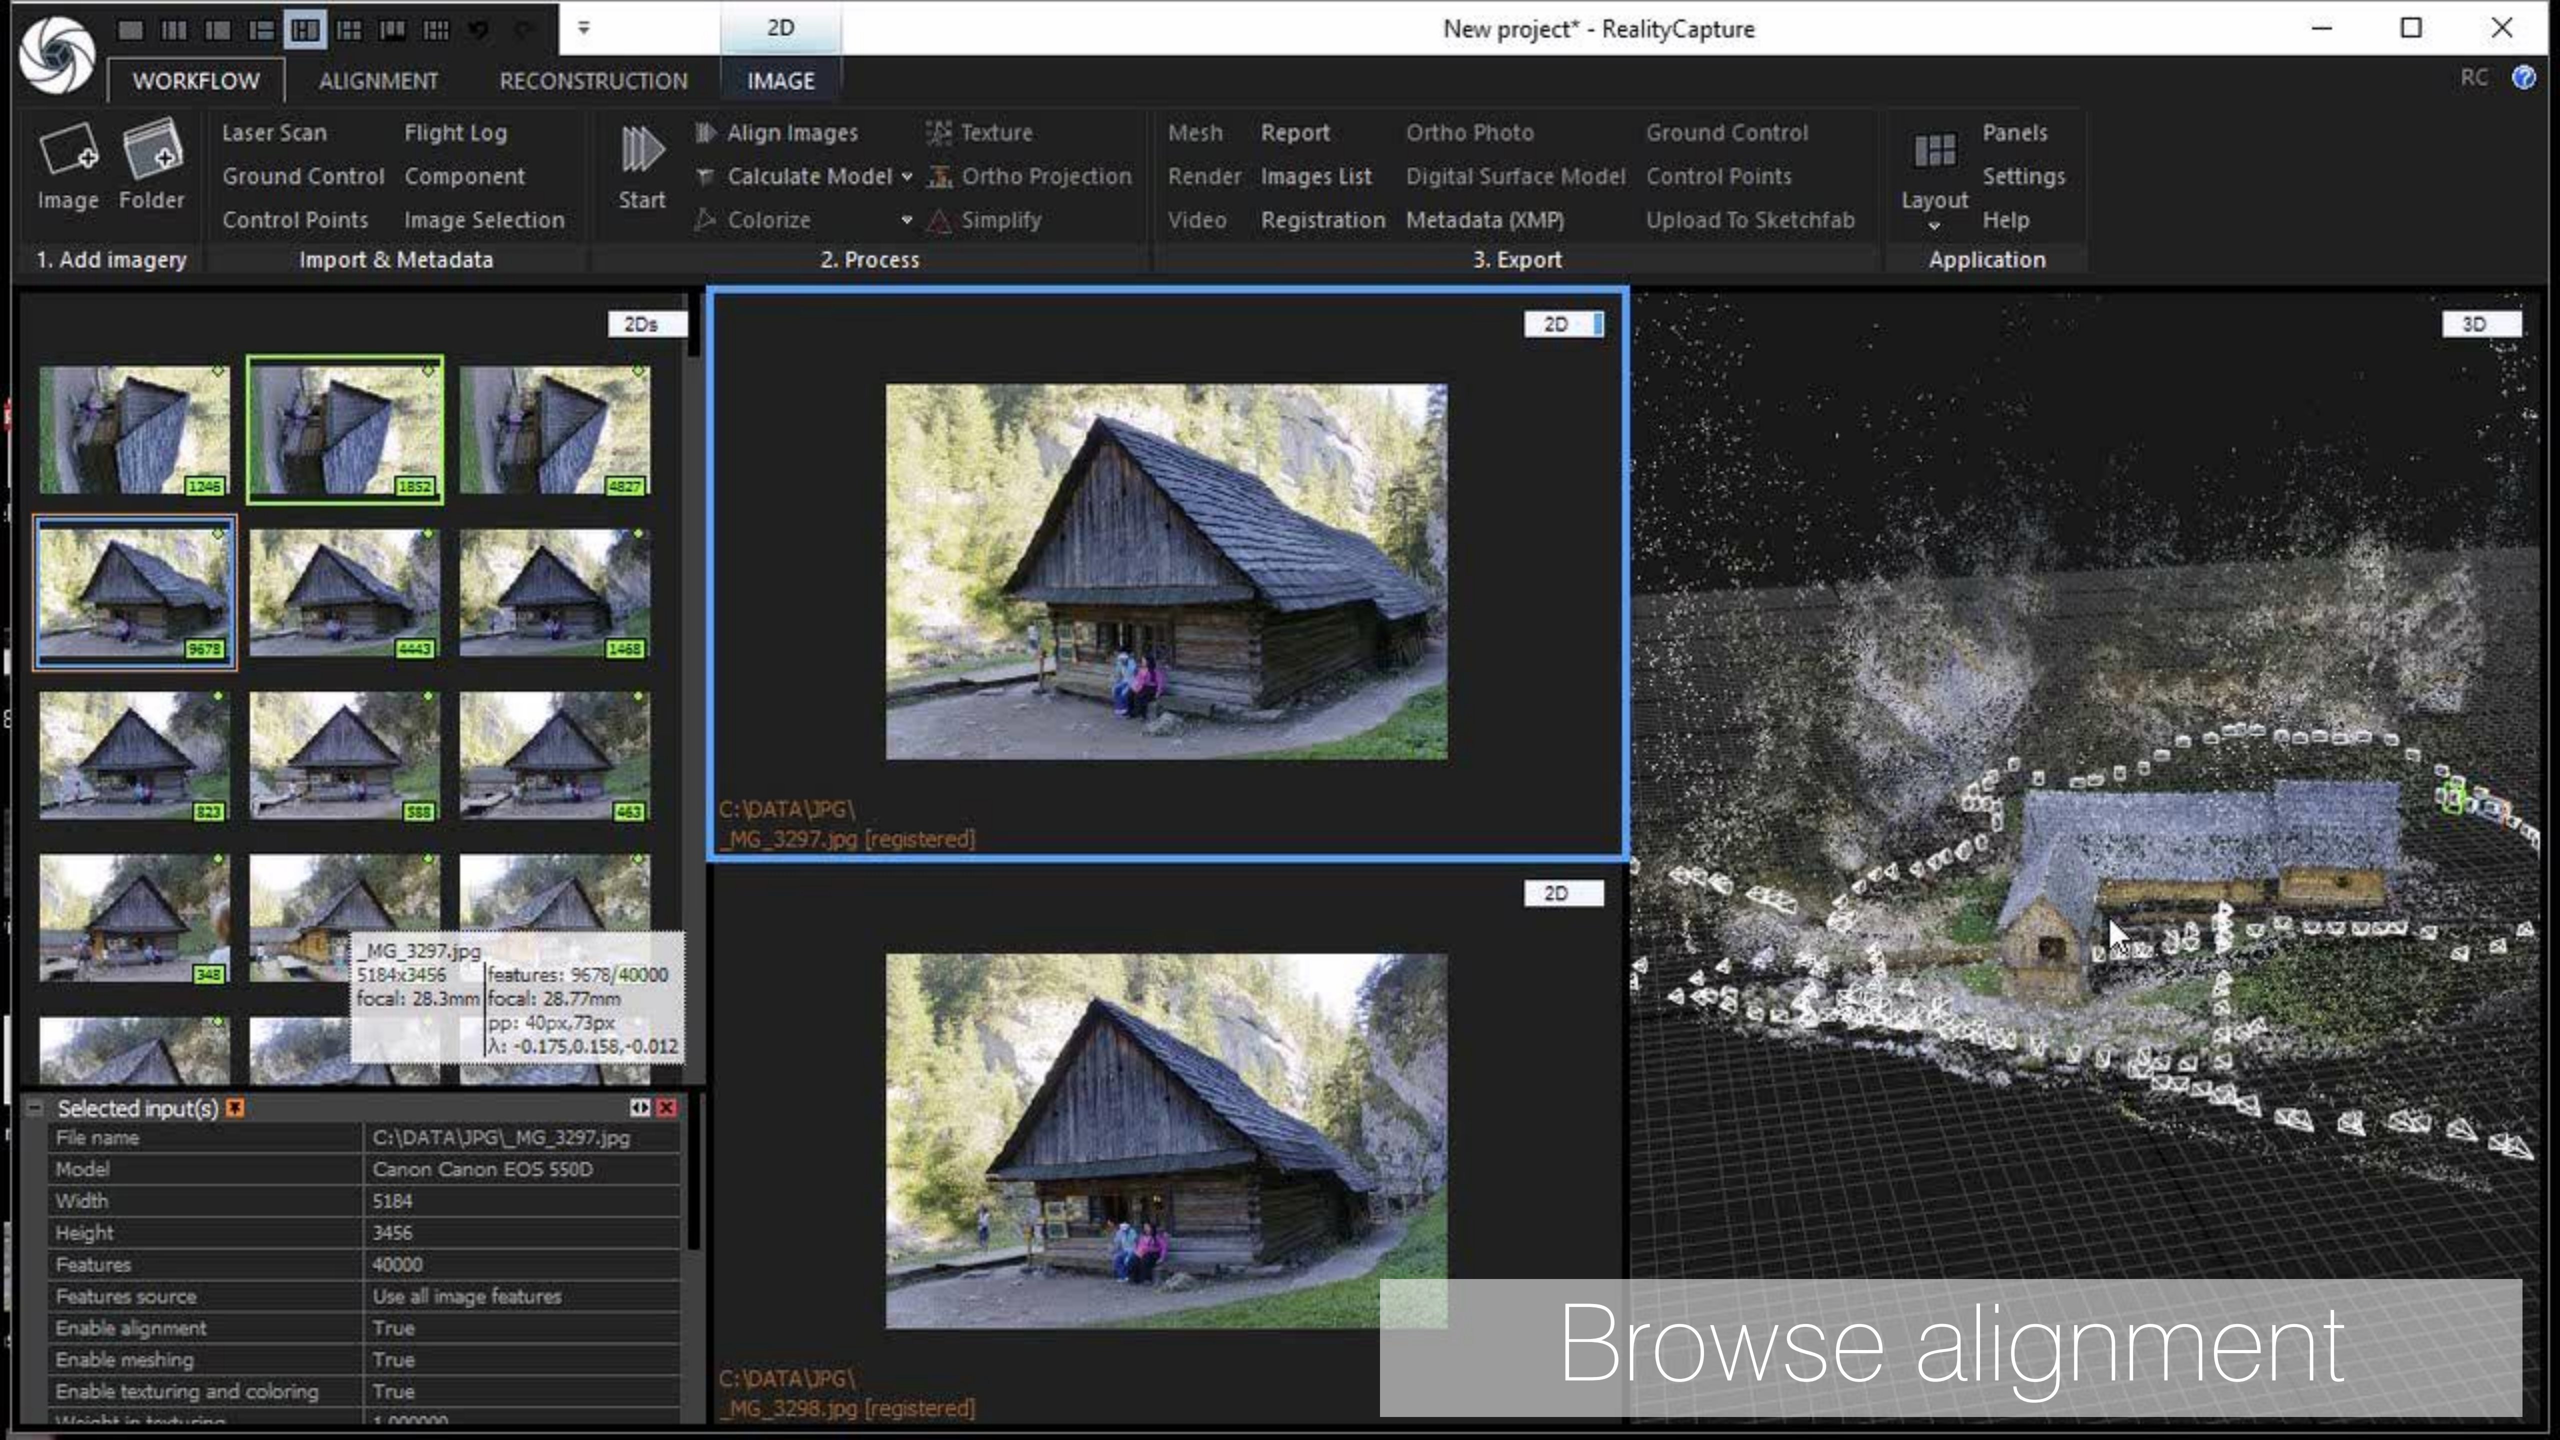
New project* - RealityCapture
2D
WORKFLOW
ALIGNMENT
RECONSTRUCTION
IMAGE
RC
Image
Folder
1. Add imagery
Laser Scan
Ground Control
Control Points
Flight Log
Component
Image Selection
Import & Metadata
Start
Align Images
Calculate Model
Colorize
Texture
Ortho Projection
Simplify
2. Process
Mesh
Render
Video
Report
Images List
Registration
Ortho Photo
Digital Surface Model
Metadata (XMP)
Ground Control
Control Points
Upload To Sketchfab
3. Export
Layout
Panels
Settings
Help
Application
2Ds
2D
3D
C:\DATA\JPG\
_MG_3297.jpg [registered]
_MG_3297.jpg
5184x3456
focal: 28.3mm
features: 9678/40000
focal: 28.77mm
pp: 40px,73px
λ: -0.175,0.158,-0.012
Selected input(s)
File name C:\DATA\JPG_MG_3297.jpg
Model Canon Canon EOS 550D
Width 5184
Height 3456
Features 40000
Features source Use all image features
Enable alignment True
Enable meshing True
Enable texturing and coloring True
C:\DATA\JPG\
_MG_3298.jpg [registered]
Browse alignment

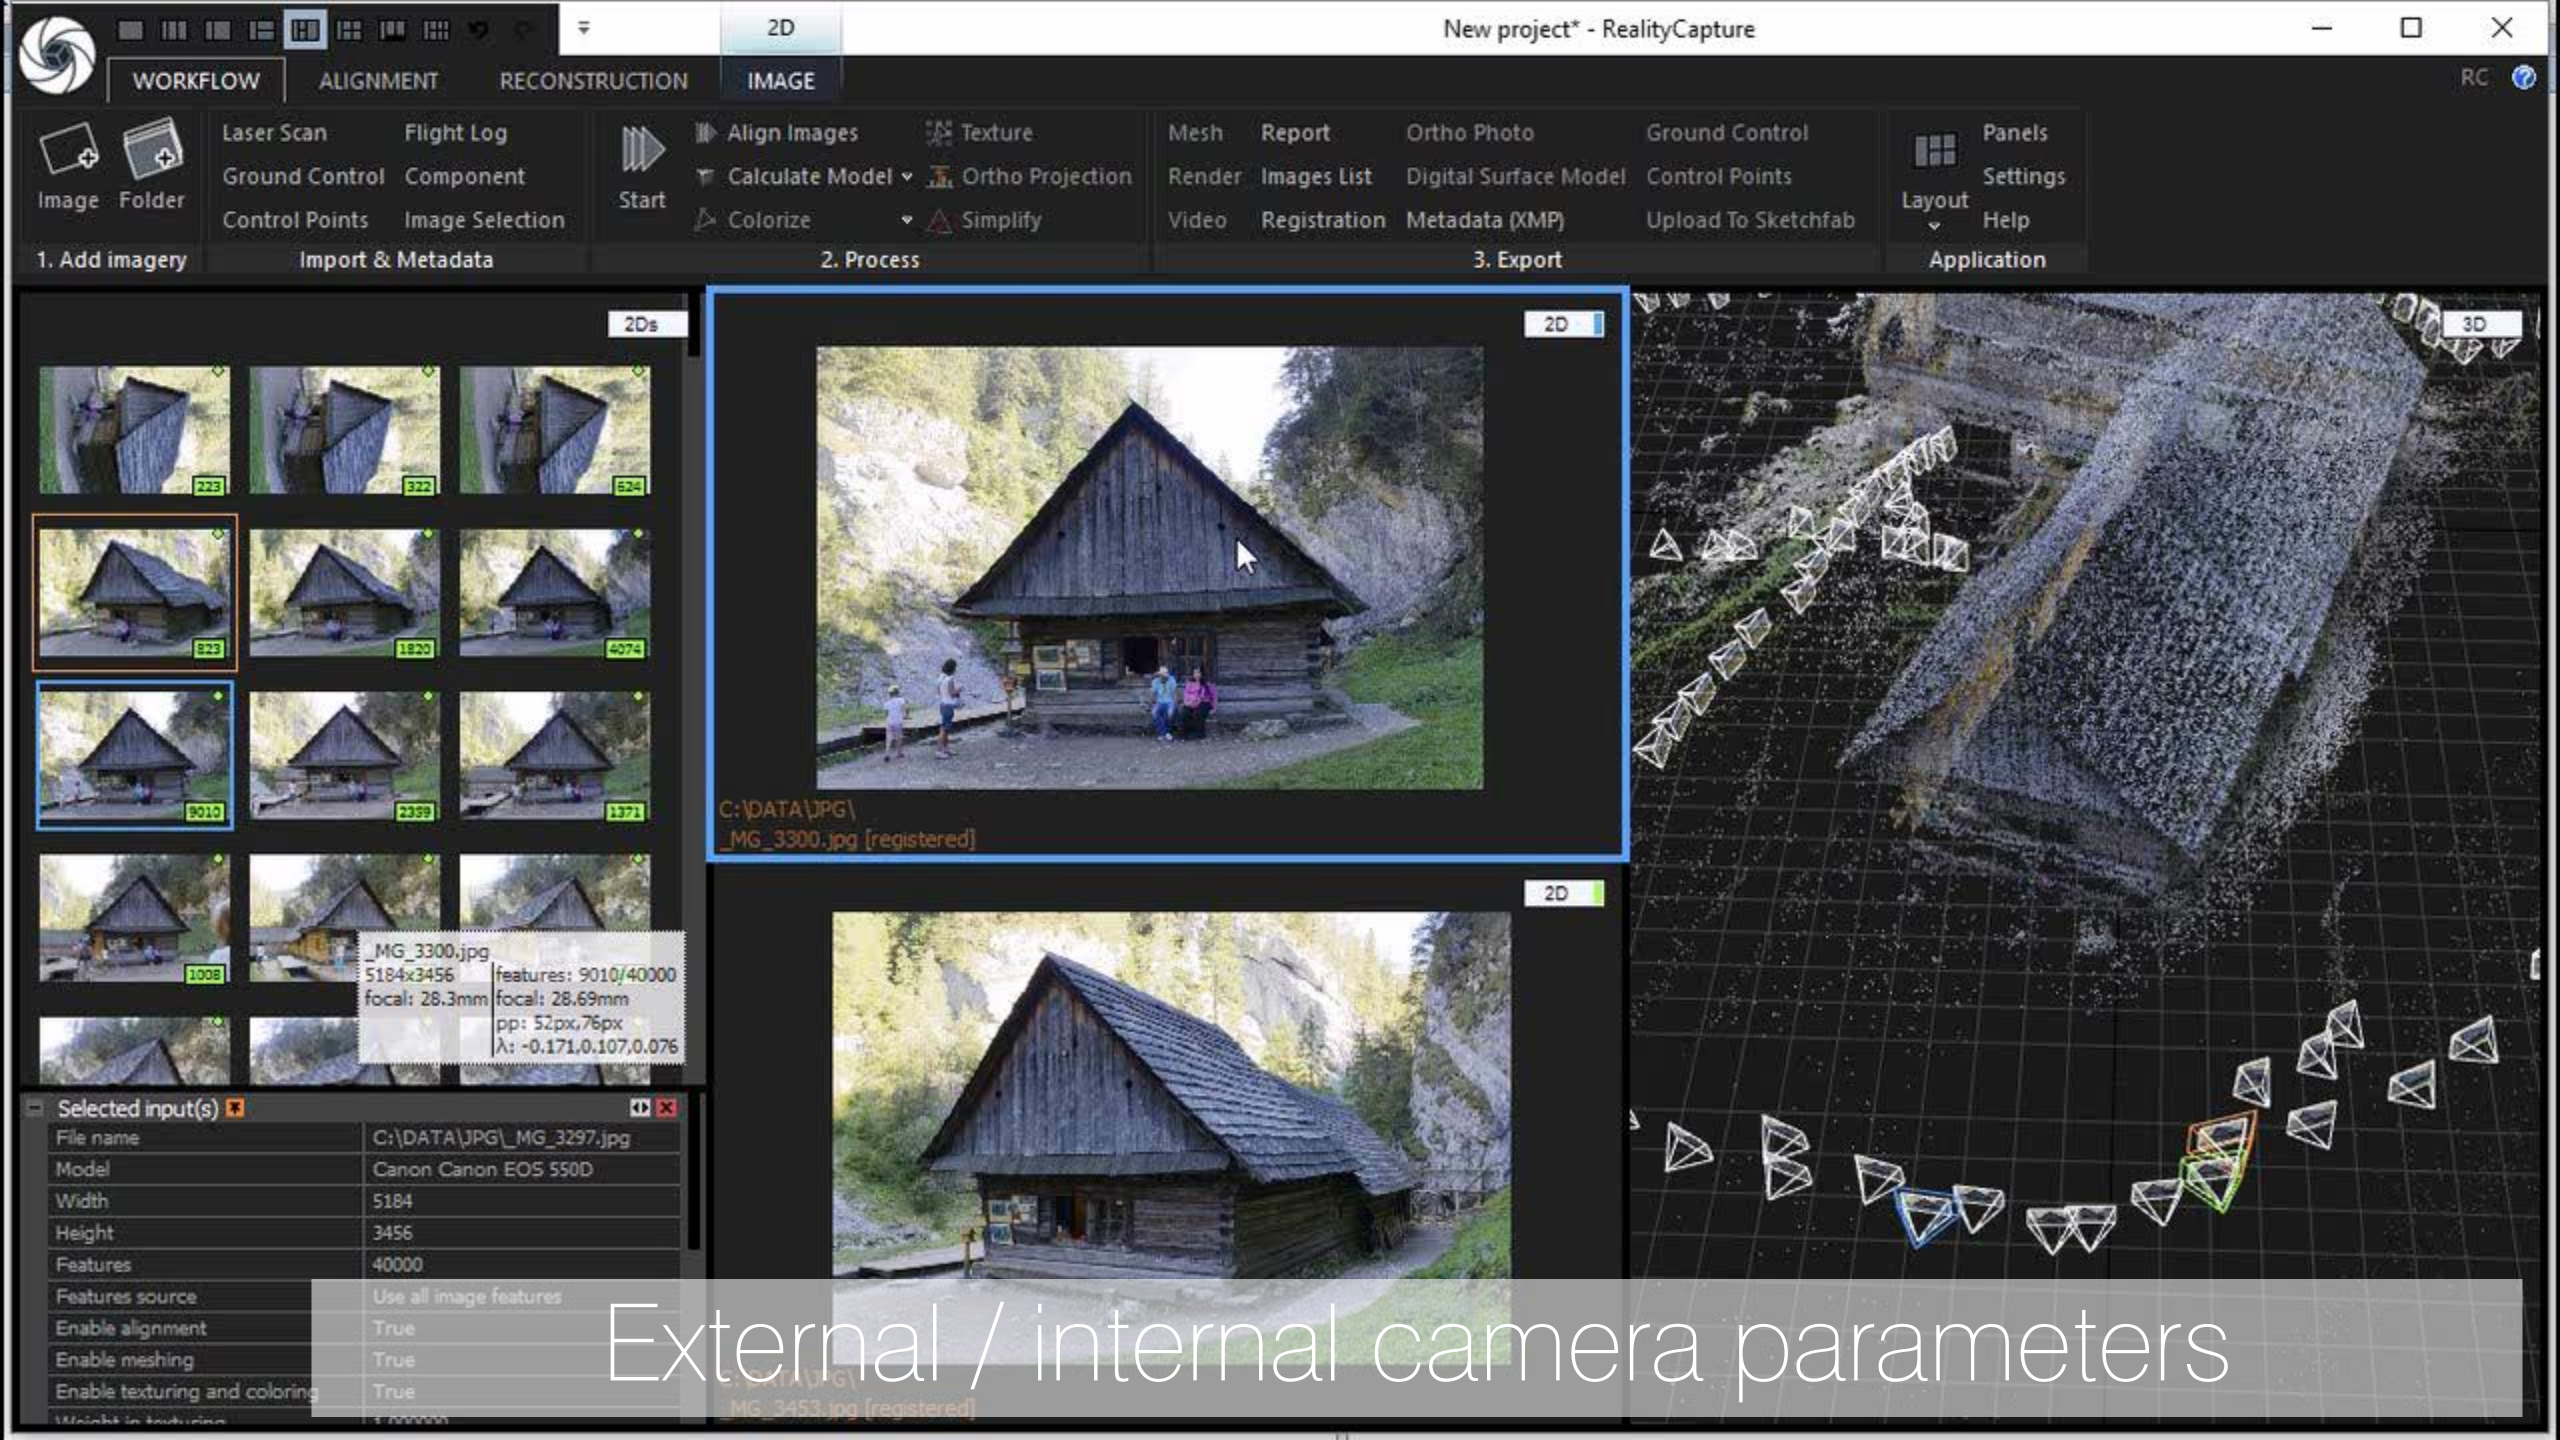
New project* - RealityCapture
2D
WORKFLOW
ALIGNMENT
RECONSTRUCTION
IMAGE
RC
Image
Folder
1. Add imagery
Laser Scan
Ground Control
Control Points
Flight Log
Component
Image Selection
Import & Metadata
Start
Align Images
Calculate Model
Colorize
Texture
Ortho Projection
Simplify
2. Process
Mesh
Render
Video
Report
Images List
Registration
Ortho Photo
Digital Surface Model
Metadata (XMP)
Ground Control
Control Points
Upload To Sketchfab
3. Export
Layout
Panels
Settings
Help
Application
2Ds
2D
3D
C:\DATA\JPG\
_MG_3300.jpg [registered]
_MG_3300.jpg
5184x3456
focal: 28.3mm
features: 9010/40000
focal: 28.69mm
pp: 52px,76px
λ: -0.171,0.107,0.076
Selected input(s)
File name
C:\DATA\JPG_MG_3297.jpg
Model
Canon Canon EOS 550D
Width
5184
Height
3456
Features
40000
Features source
Use all image features
Enable alignment
True
Enable meshing
True
Enable texturing and coloring
True
External / internal camera parameters

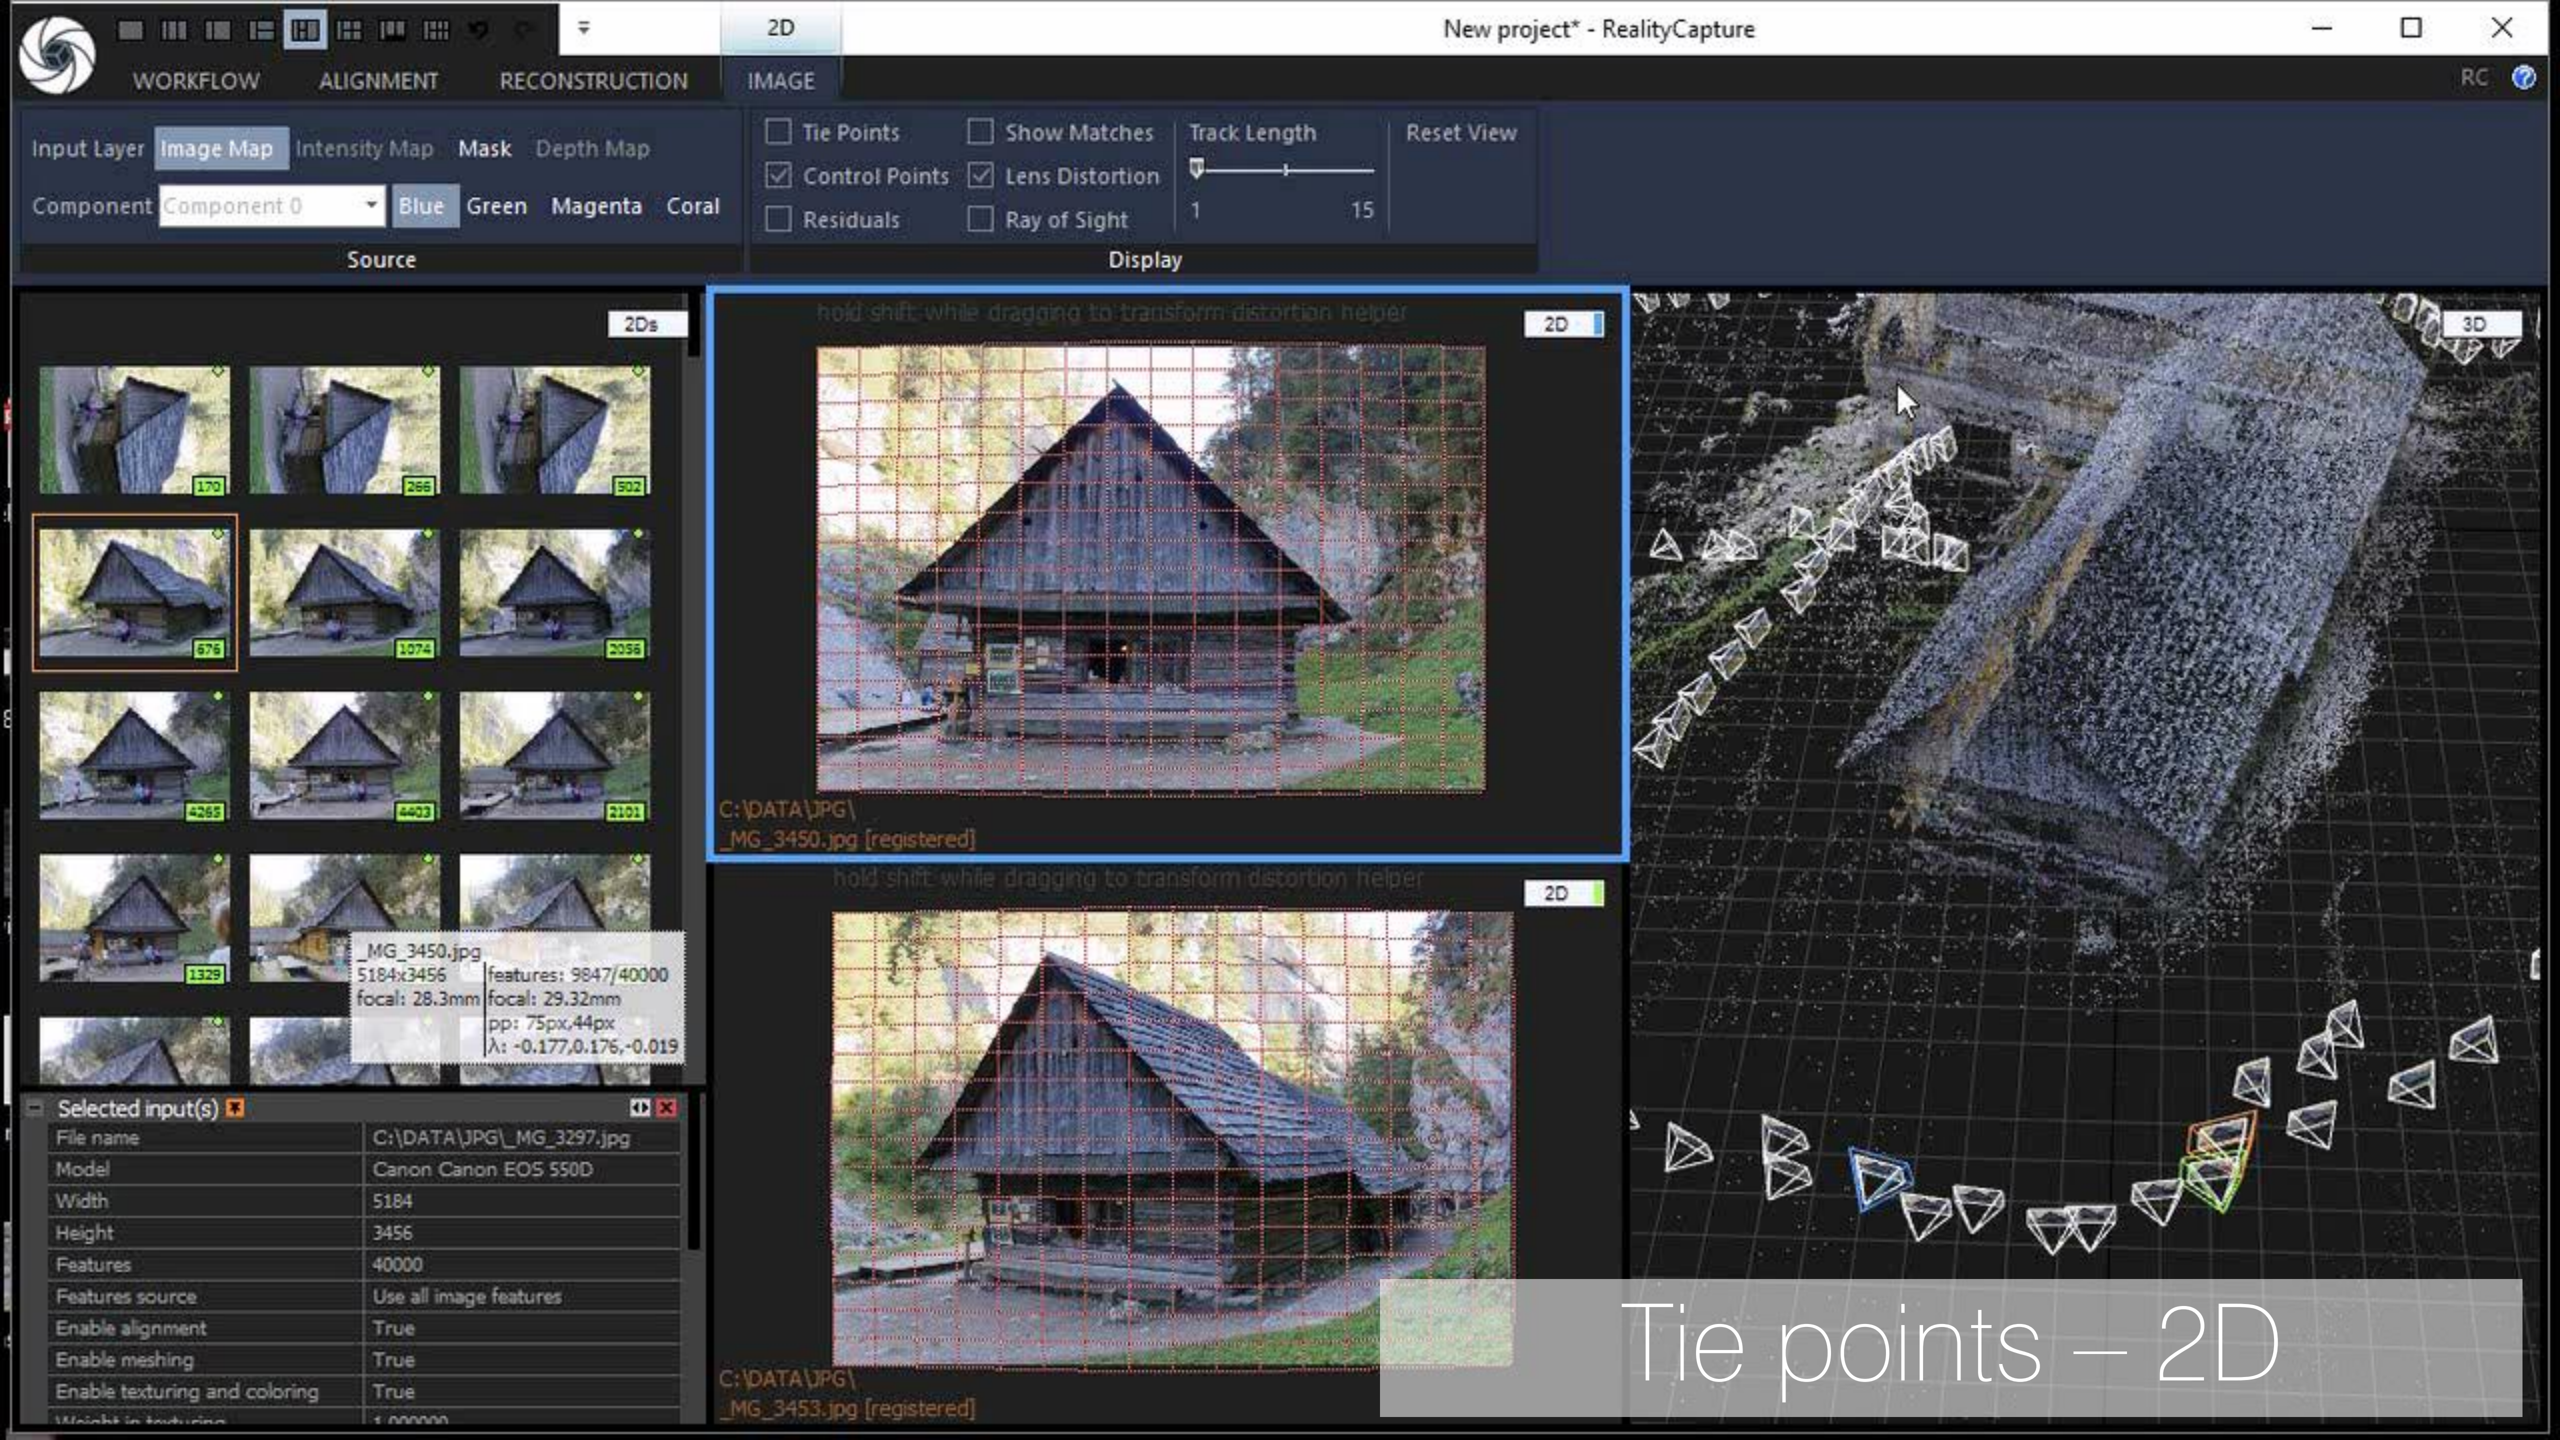

New project* - RealityCapture
WORKFLOW
ALIGNMENT
RECONSTRUCTION
IMAGE
2D
Input Layer
Image Map
Intensity Map
Mask
Depth Map
Component
Component 0
Blue
Green
Magenta
Coral
Source
Tie Points
Control Points
Residuals
Show Matches
Lens Distortion
Ray of Sight
Track Length
1
15
Reset View
Display
2Ds
2D
3D
hold shift while dragging to transform distortion helper
C:\DATA\JPG\
_MG_3450.jpg [registered]
C:\DATA\JPG\
_MG_3453.jpg [registered]
_MG_3450.jpg
5184x3456
focal: 28.3mm
features: 9847/40000
focal: 29.32mm
pp: 75px,44px
λ: -0.177,0.176,-0.019
Selected input(s)
File name C:\DATA\JPG_MG_3297.jpg
Model Canon Canon EOS 550D
Width 5184
Height 3456
Features 40000
Features source Use all image features
Enable alignment True
Enable meshing True
Enable texturing and coloring True
Tie points – 2D

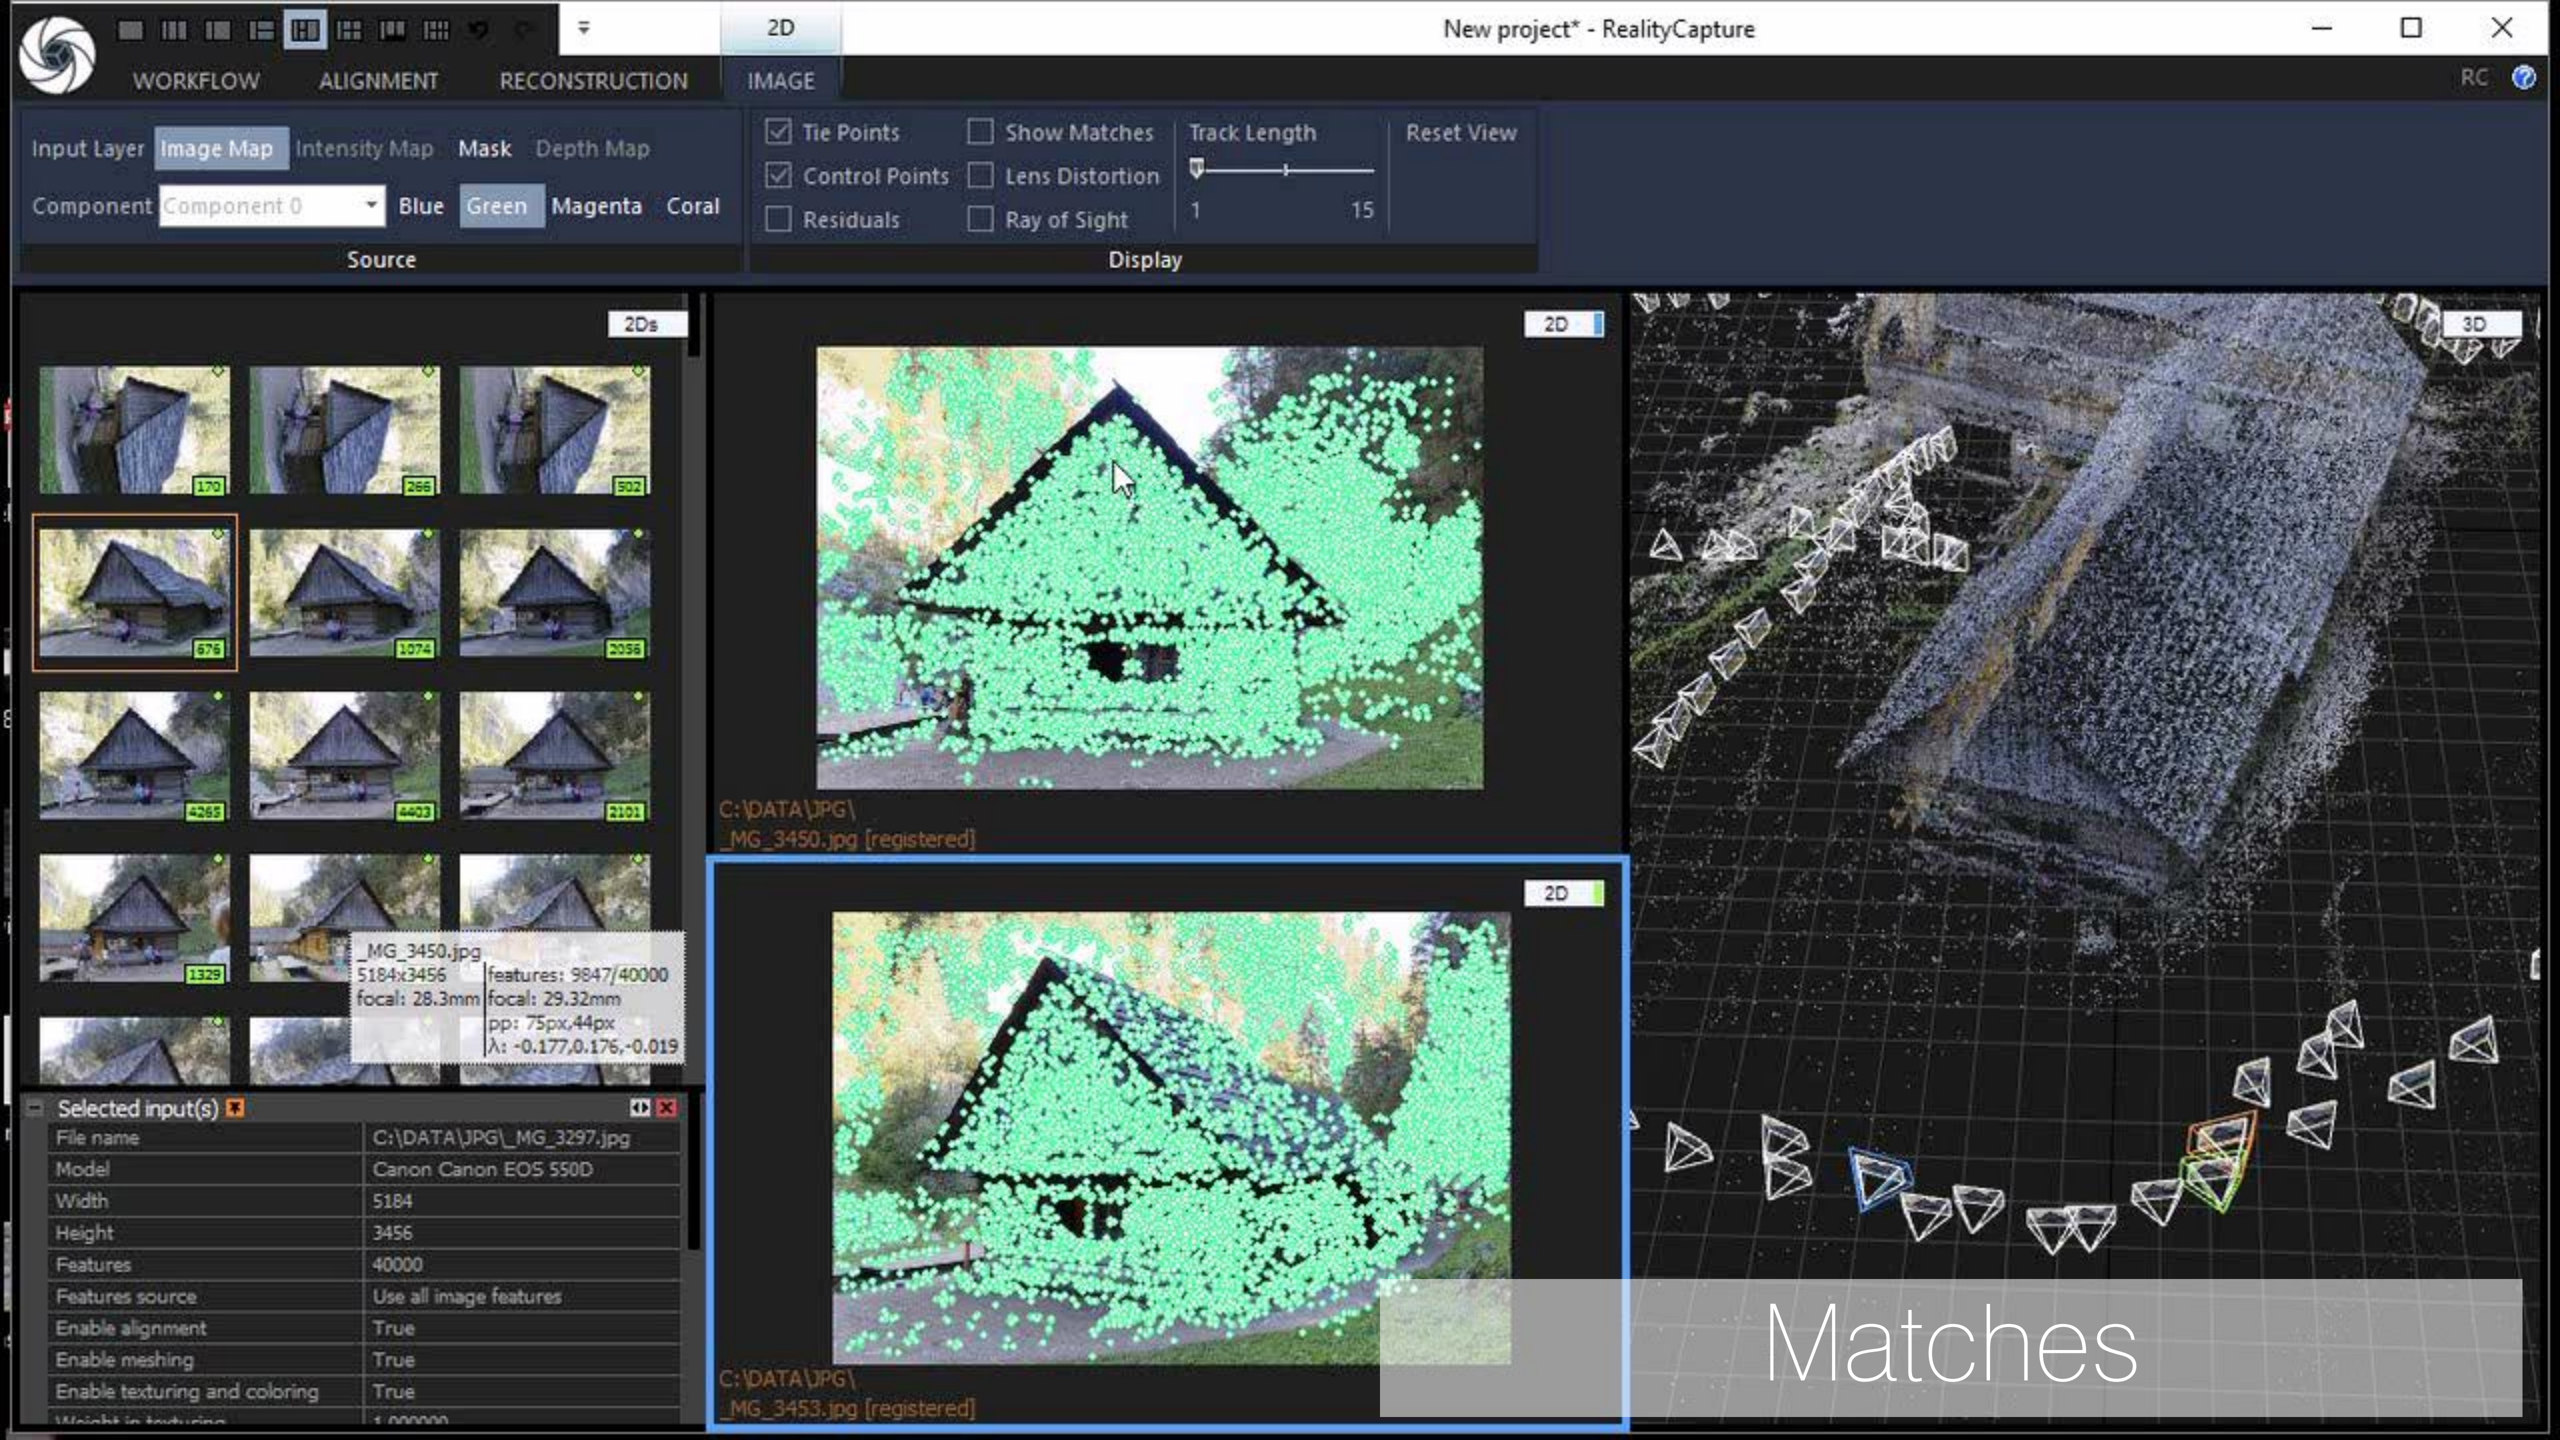

2D
New project* - RealityCapture
WORKFLOW
ALIGNMENT
RECONSTRUCTION
IMAGE
RC
Input Layer
Image Map
Intensity Map
Mask
Depth Map
Component
Component 0
Blue
Green
Magenta
Coral
Source
Tie Points
Control Points
Residuals
Show Matches
Lens Distortion
Ray of Sight
Track Length
1
15
Reset View
Display
2Ds
2D
3D
C:\DATA\JPG\
_MG_3450.jpg [registered]
_MG_3450.jpg
5184x3456
focal: 28.3mm
features: 9847/40000
focal: 29.32mm
pp: 75px,44px
λ: -0.177,0.176,-0.019
C:\DATA\JPG\
_MG_3453.jpg [registered]
Selected input(s)
File name
C:\DATA\JPG_MG_3297.jpg
Model
Canon Canon EOS 550D
Width
5184
Height
3456
Features
40000
Features source
Use all image features
Enable alignment
True
Enable meshing
True
Enable texturing and coloring
True
Matches

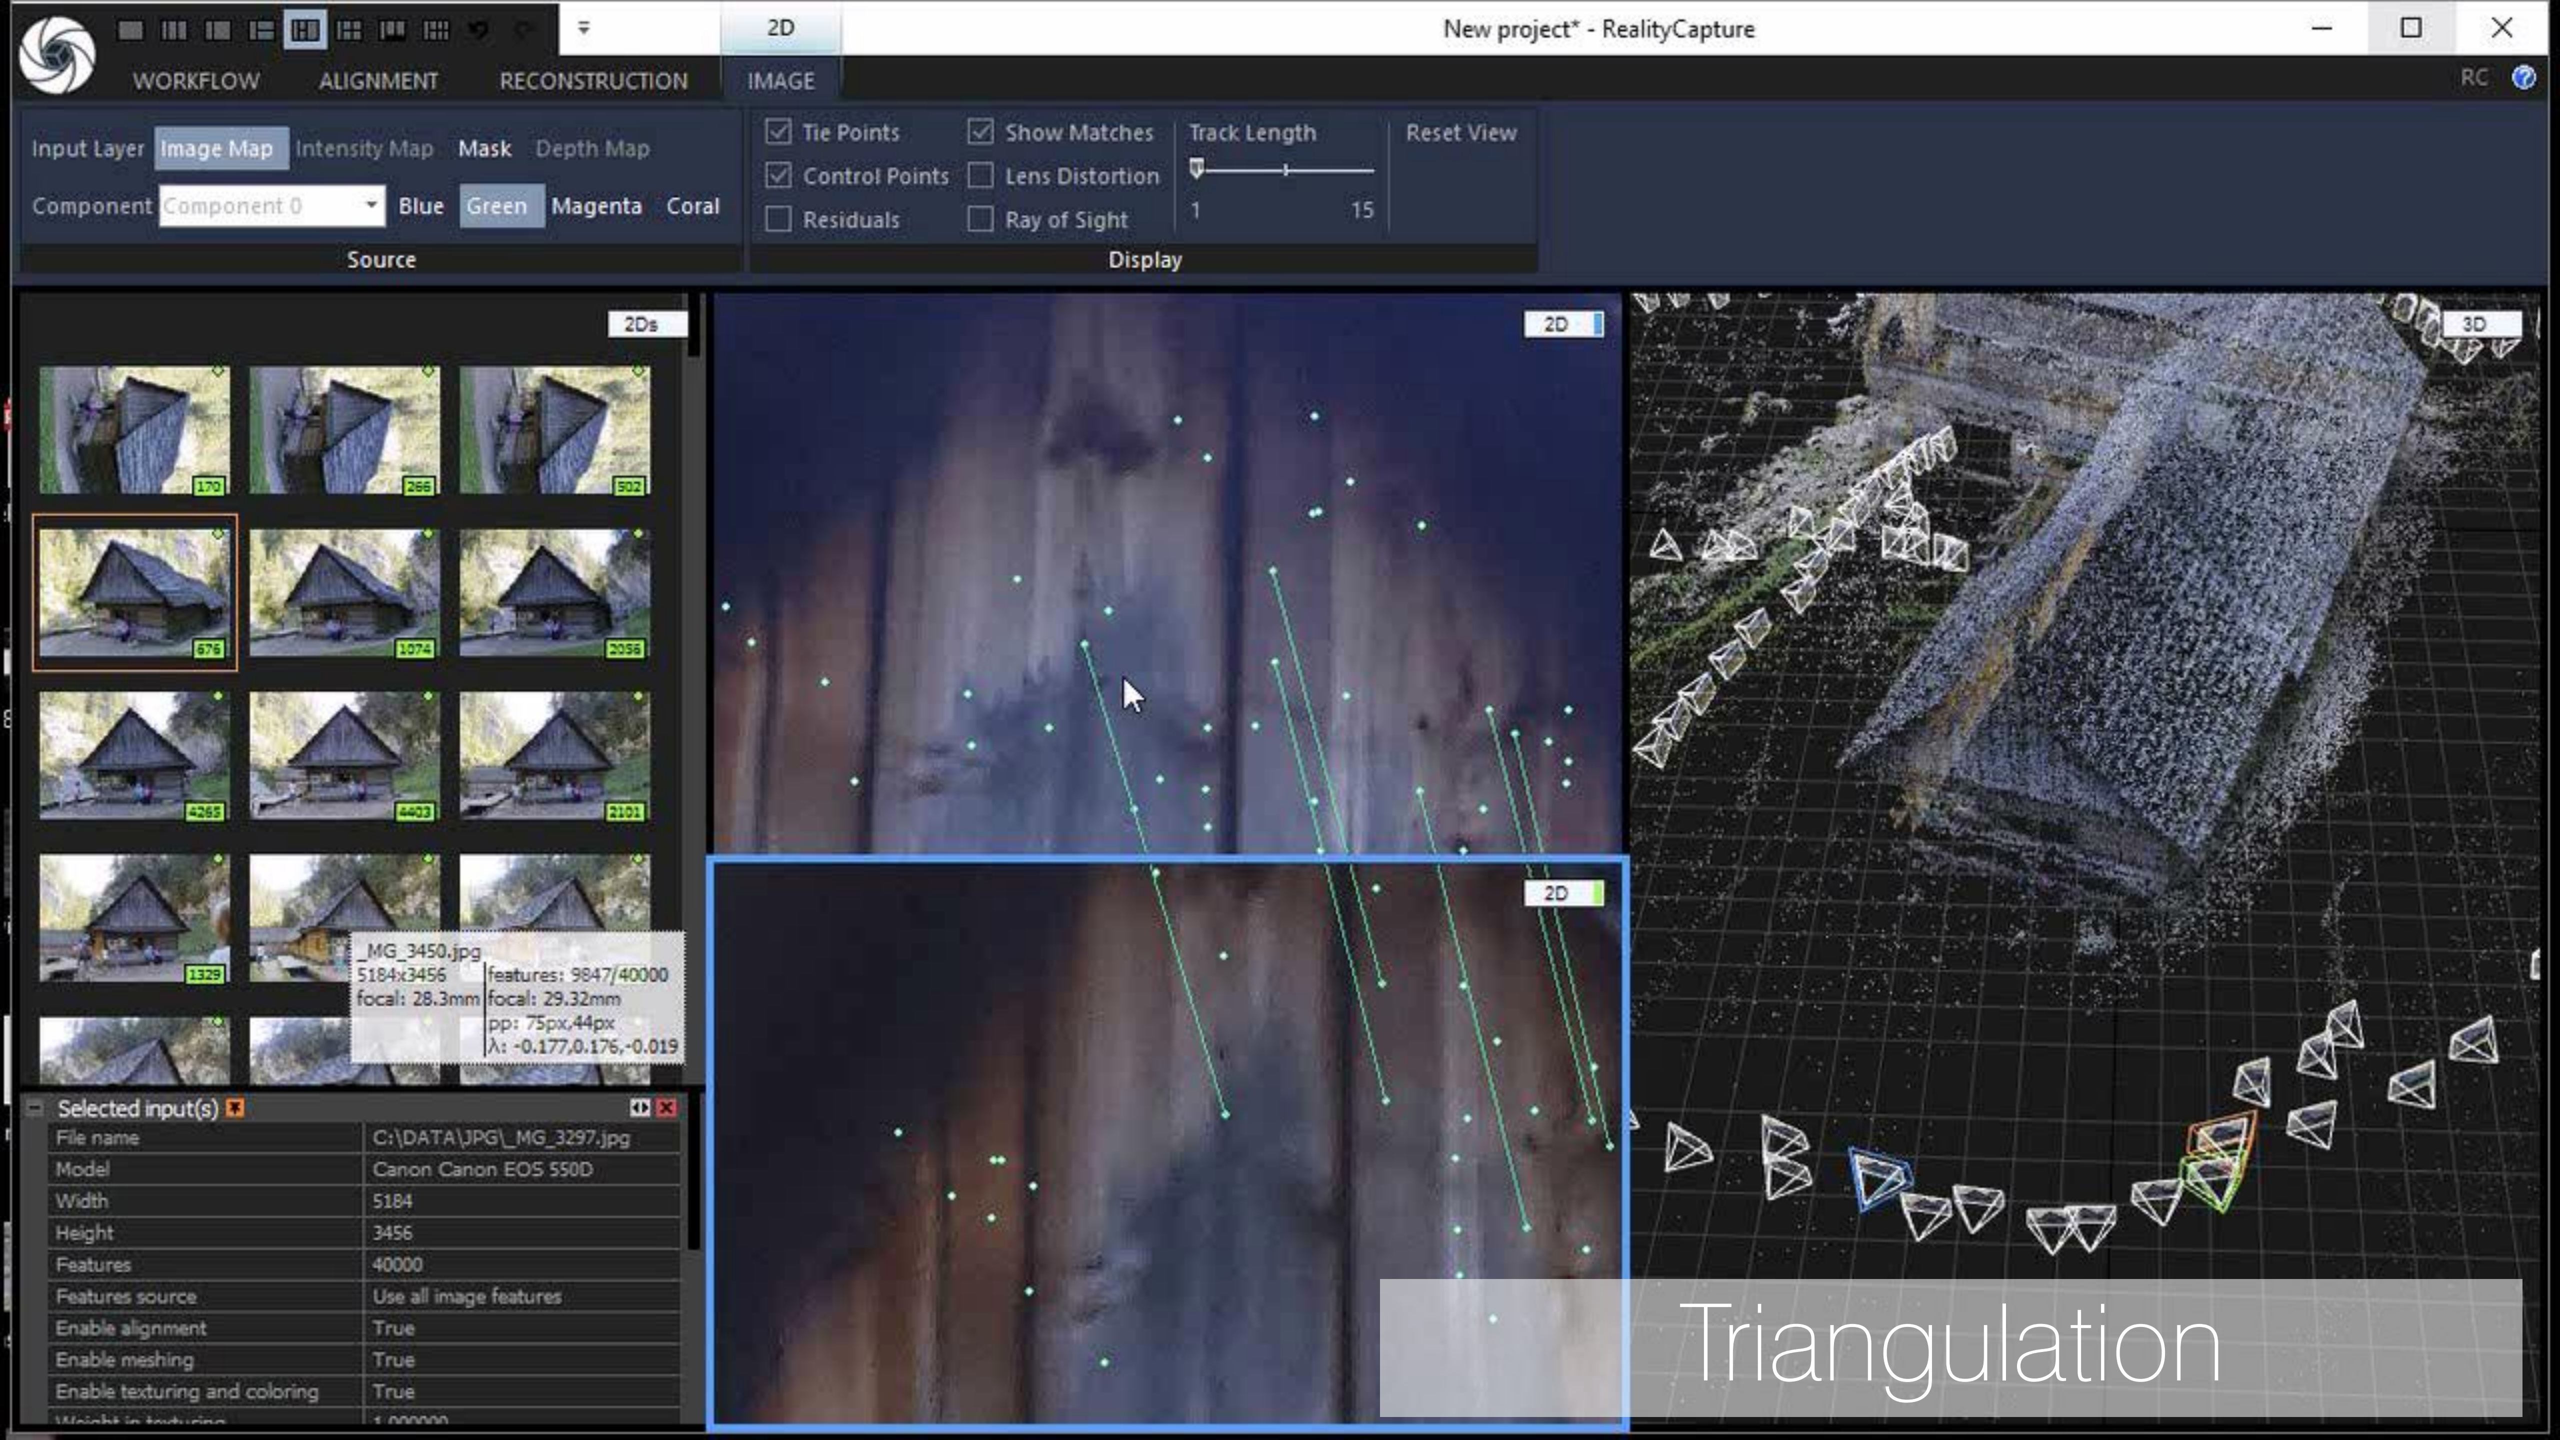

2D
New project* - RealityCapture
WORKFLOW
ALIGNMENT
RECONSTRUCTION
IMAGE
RC
Input Layer
Image Map
Intensity Map
Mask
Depth Map
Component
Component 0
Blue
Green
Magenta
Coral
Source
Tie Points
Control Points
Residuals
Show Matches
Lens Distortion
Ray of Sight
Track Length
1
15
Reset View
Display
2Ds
2D
3D
_MG_3450.jpg
5184x3456
focal: 28.3mm
features: 9847/40000
focal: 29.32mm
pp: 75px,44px
λ: -0.177,0.176,-0.019
Selected input(s)
File name
C:\DATA\JPG_MG_3297.jpg
Model
Canon Canon EOS 550D
Width
5184
Height
3456
Features
40000
Features source
Use all image features
Enable alignment
True
Enable meshing
True
Enable texturing and coloring
True
Triangulation

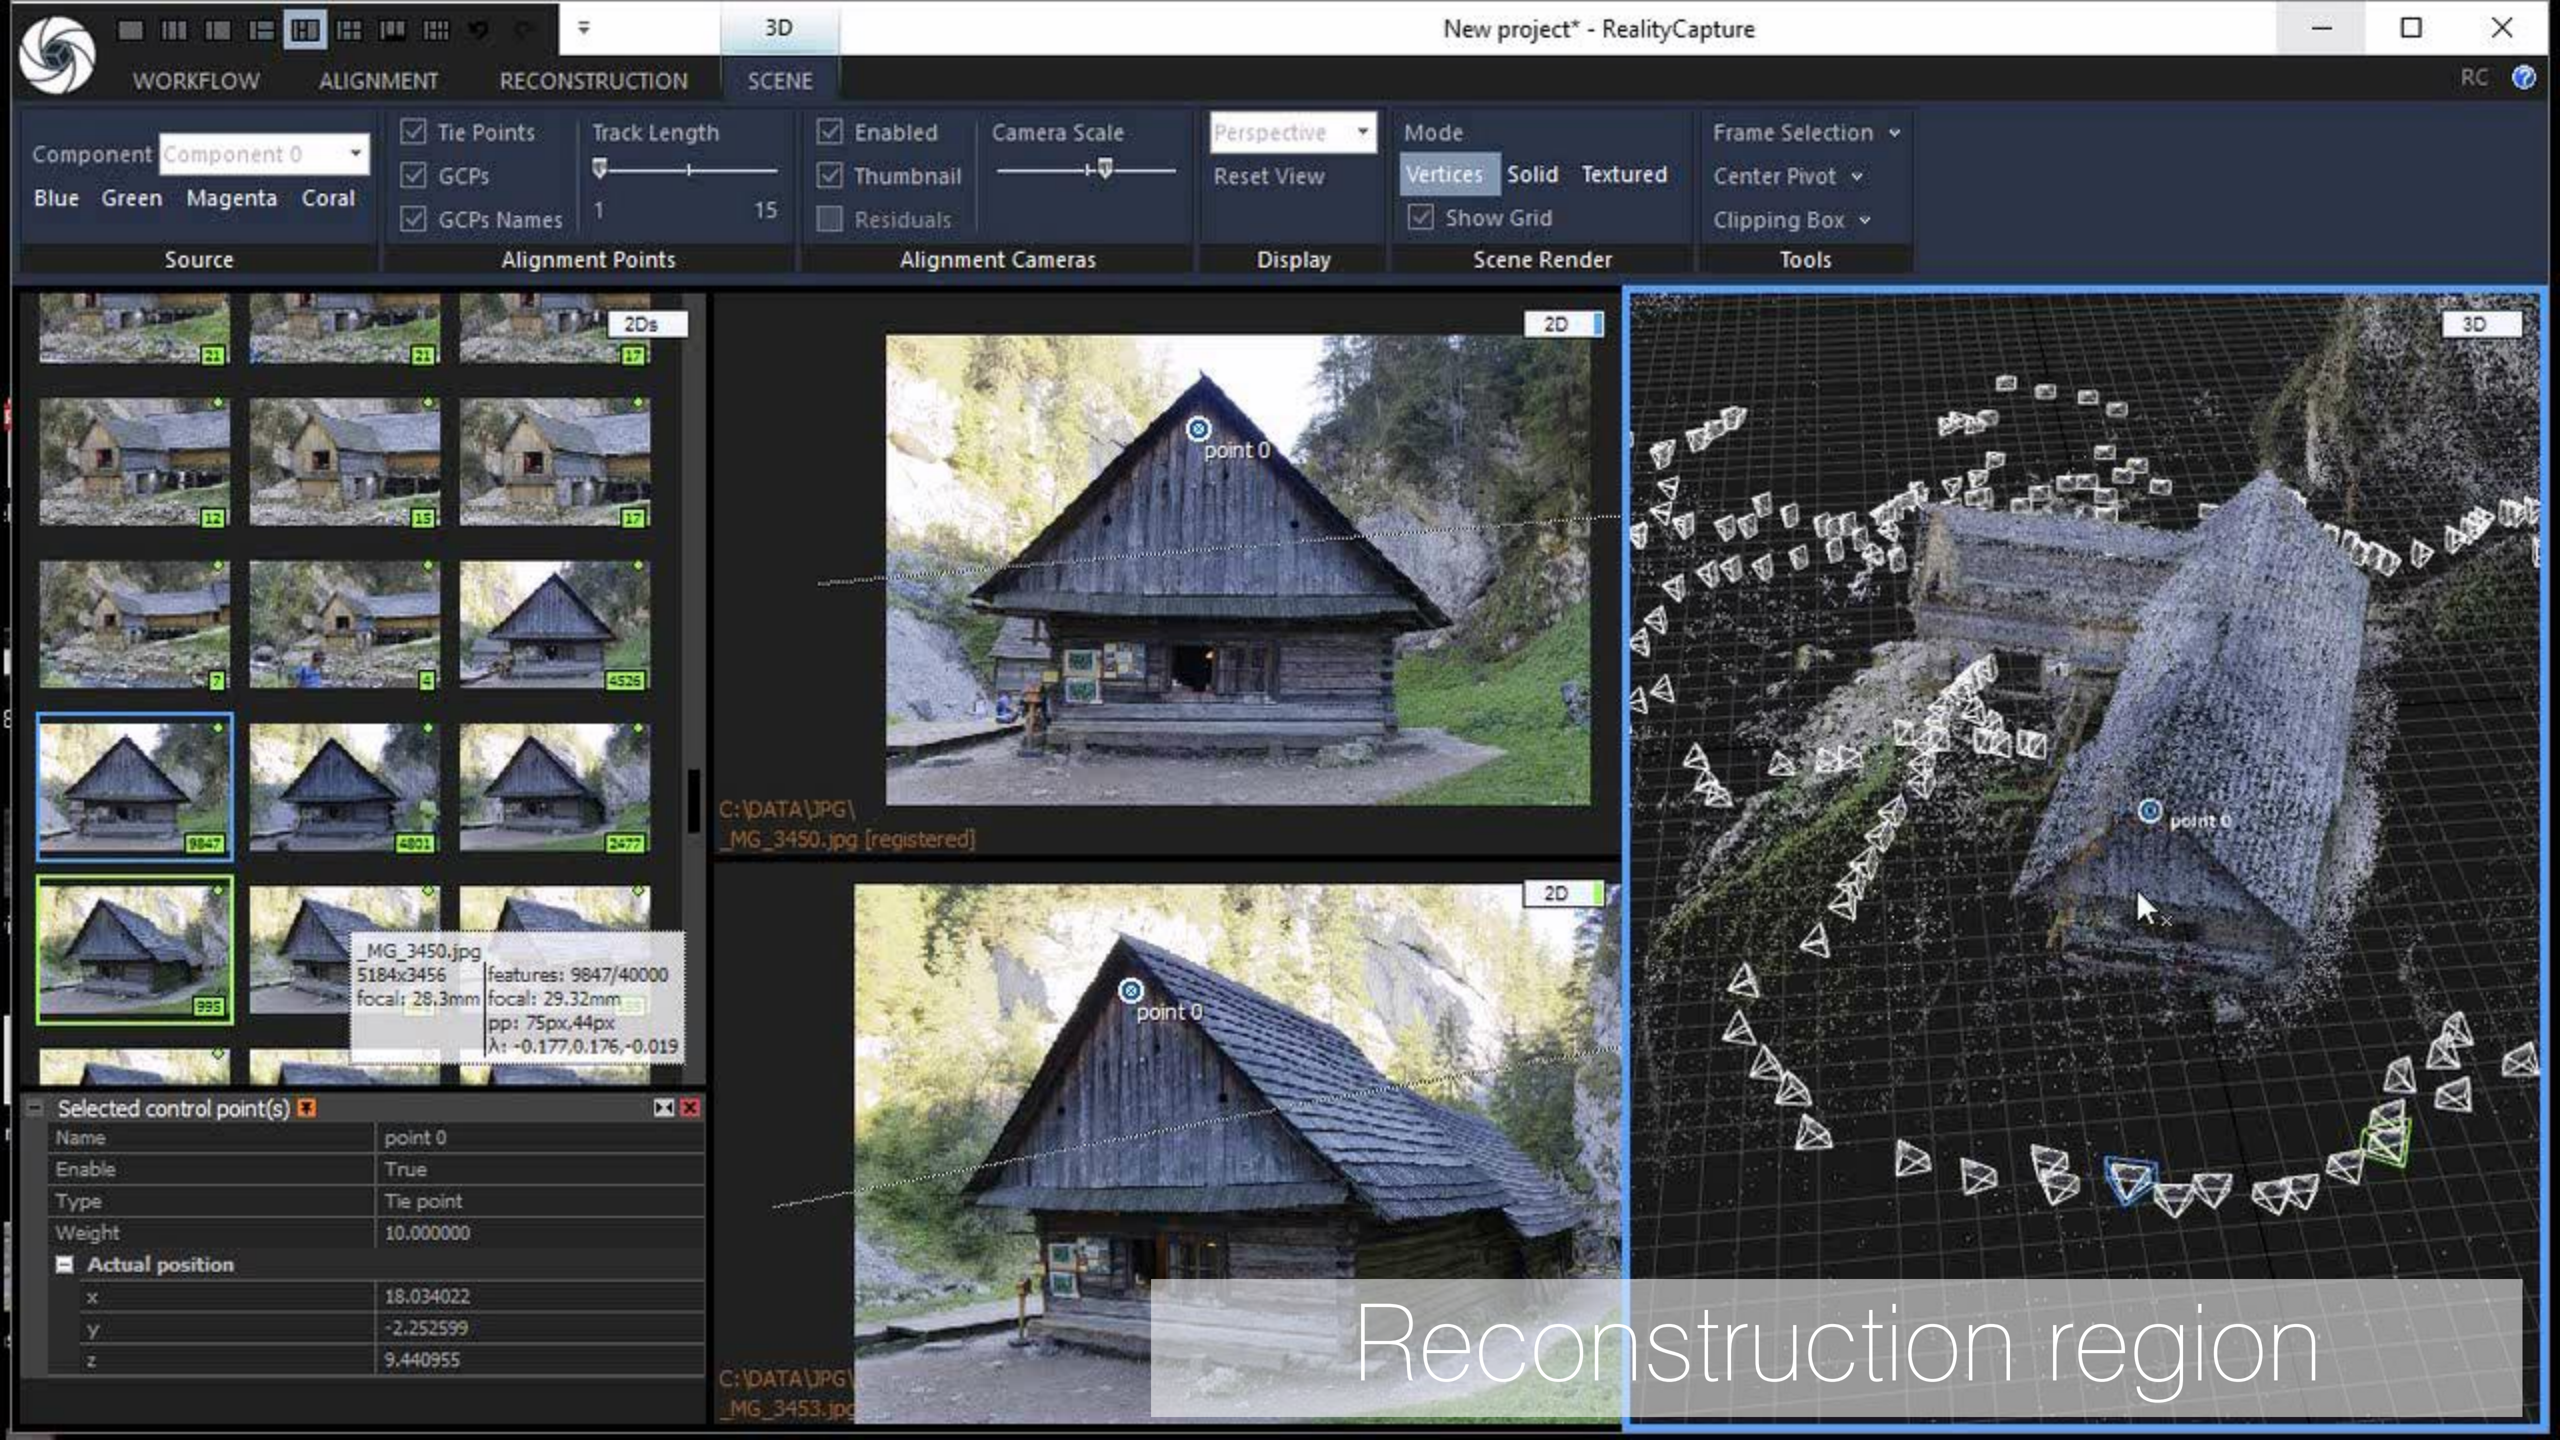
New project* - RealityCapture
3D
WORKFLOW
ALIGNMENT
RECONSTRUCTION
SCENE
RC
Component
Component 0
Blue
Green
Magenta
Coral
Source
Tie Points
GCPs
GCPs Names
Track Length
1
15
Alignment Points
Enabled
Thumbnail
Residuals
Camera Scale
Alignment Cameras
Perspective
Reset View
Display
Mode
Vertices
Solid
Textured
Show Grid
Scene Render
Frame Selection
Center Pivot
Clipping Box
Tools
2Ds
2D
3D
point 0
C:\DATA\JPG_MG_3450.jpg [registered]
C:\DATA\JPG_MG_3453.jpg
_MG_3450.jpg
5184x3456
focal: 28.3mm
features: 9847/40000
focal: 29.32mm
pp: 75px,44px
λ: -0.177,0.176,-0.019
Selected control point(s)
Name
point 0
Enable
True
Type
Tie point
Weight
10.000000
Actual position
x
18.034022
y
-2.252599
z
9.440955
Reconstruction region

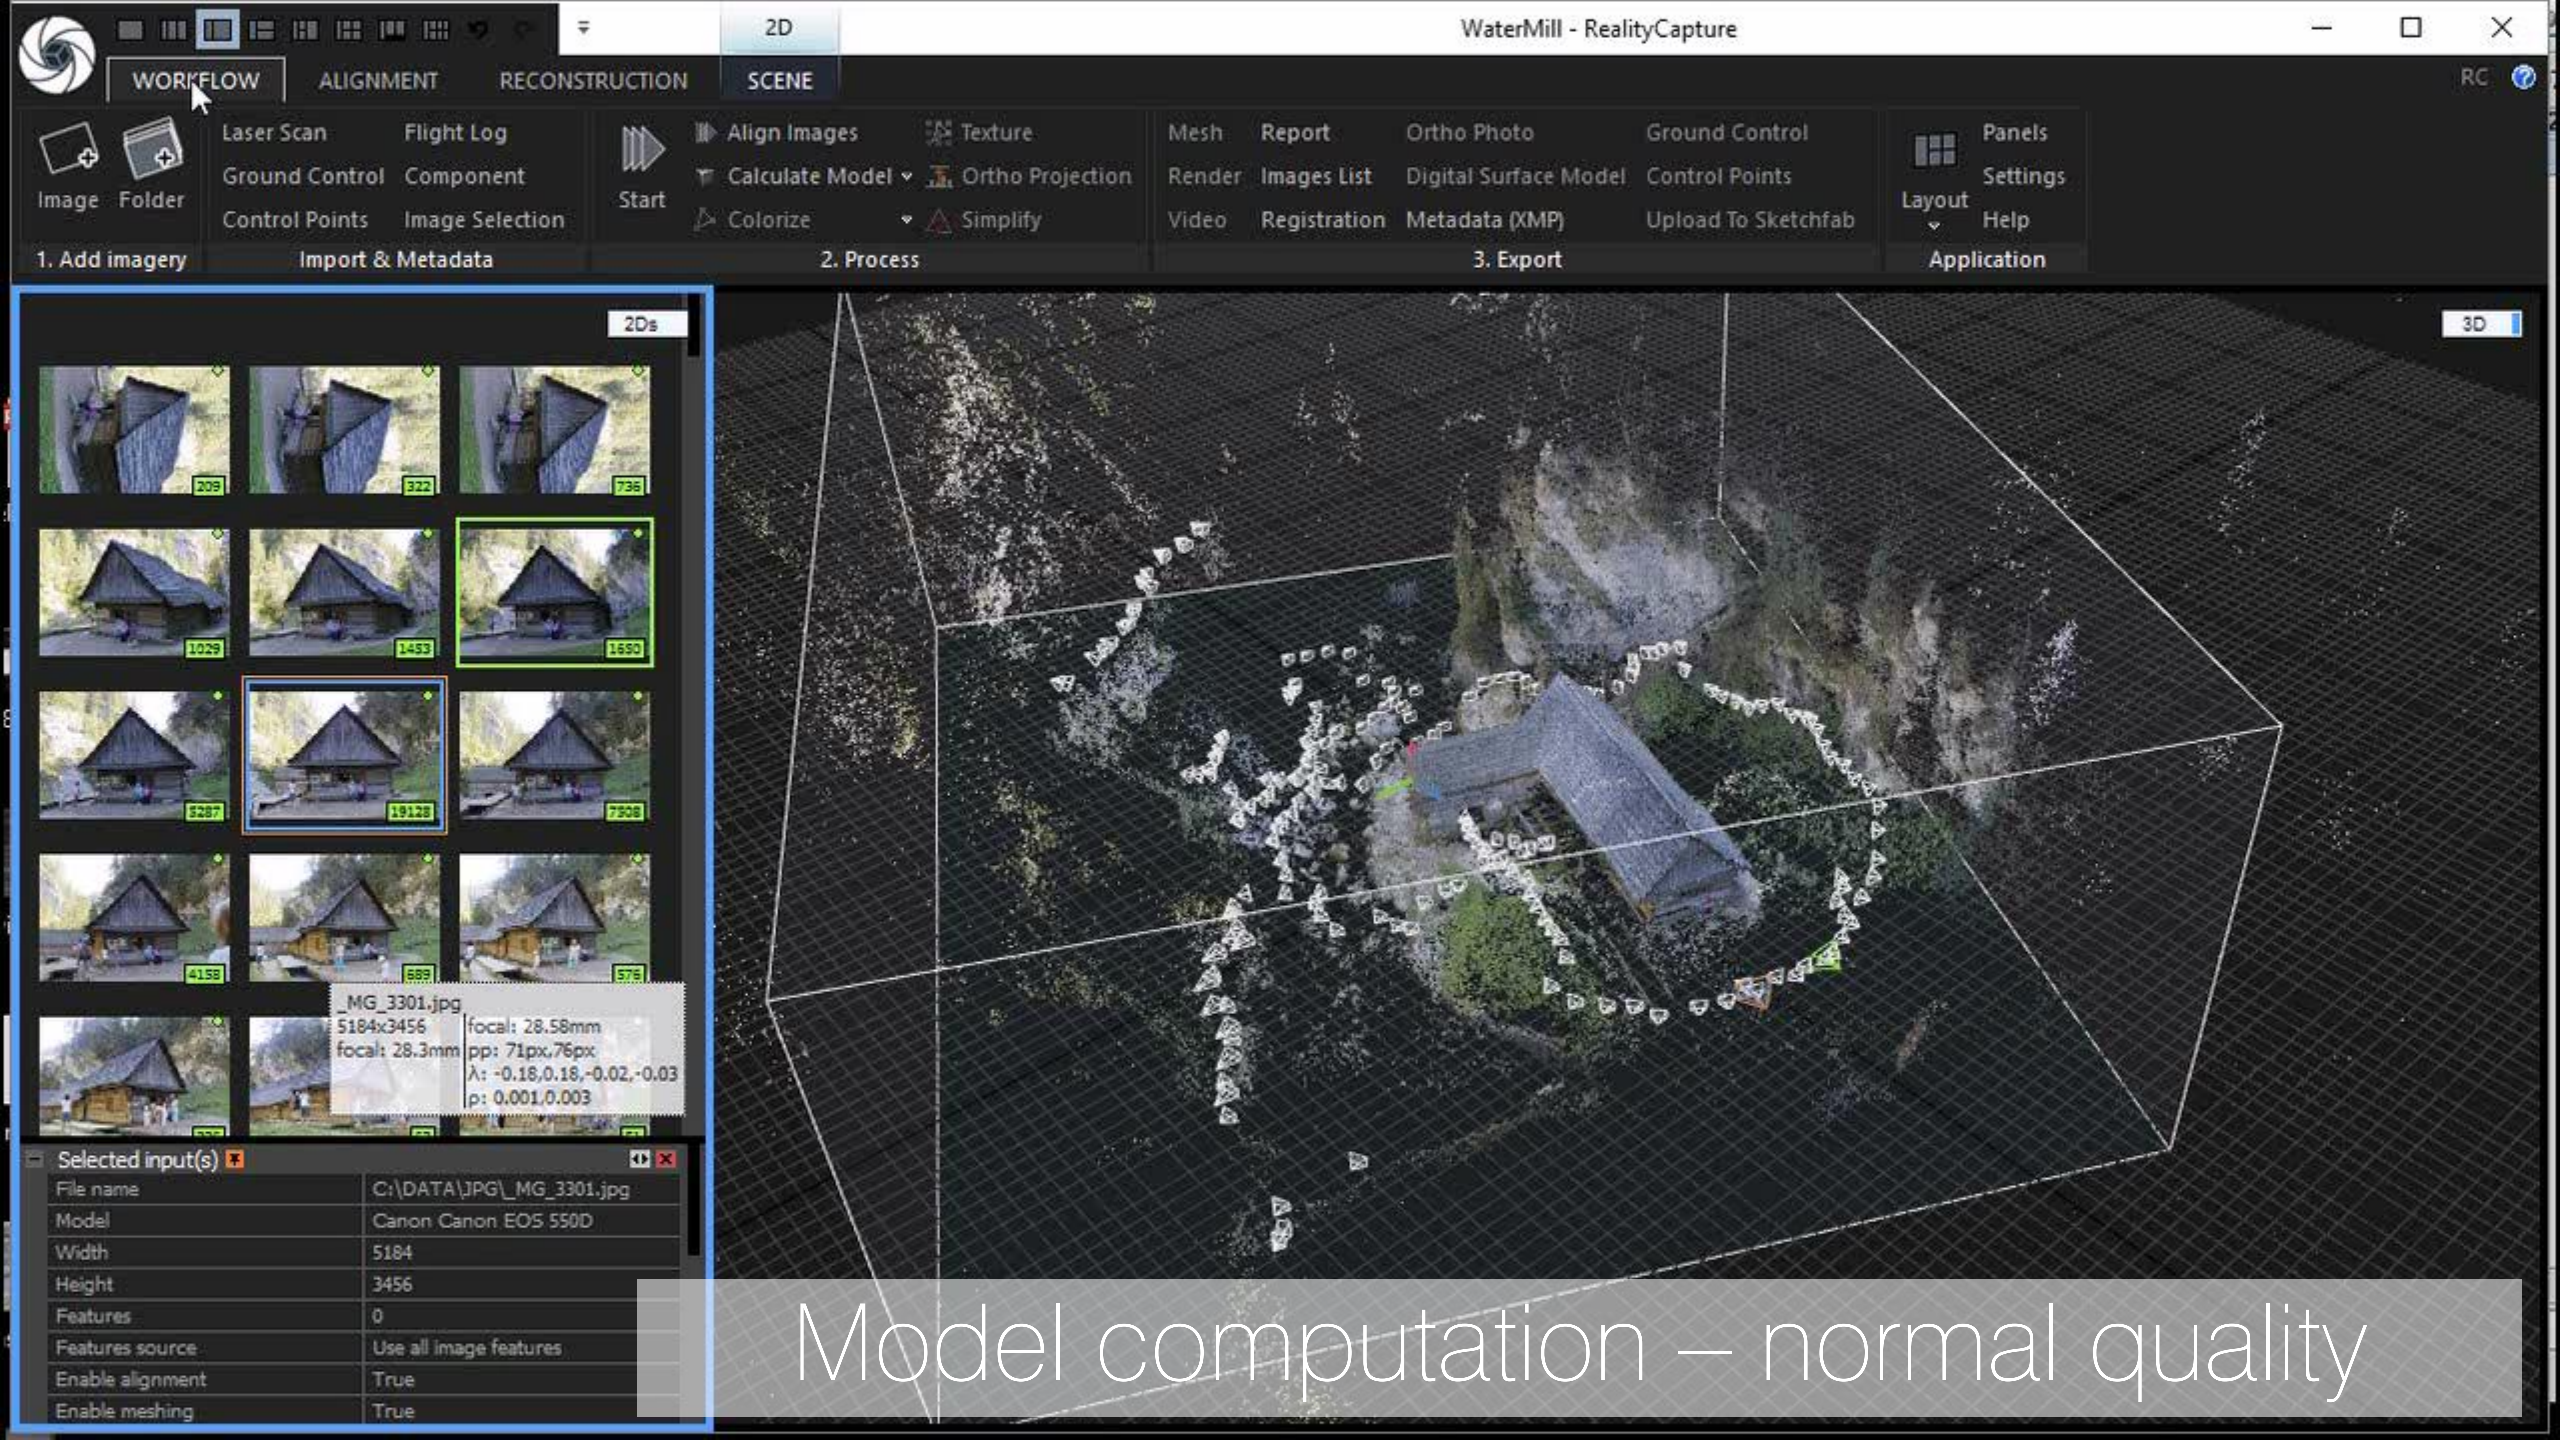
WaterMill - RealityCapture
2D
WORKFLOW
ALIGNMENT
RECONSTRUCTION
SCENE
Image
Folder
Laser Scan
Ground Control
Control Points
Flight Log
Component
Image Selection
Start
Align Images
Calculate Model
Colorize
Texture
Ortho Projection
Simplify
Mesh
Render
Video
Report
Images List
Registration
Ortho Photo
Digital Surface Model
Metadata (XMP)
Ground Control
Control Points
Upload To Sketchfab
Layout
Panels
Settings
Help
1. Add imagery
Import & Metadata
2. Process
3. Export
Application
2Ds
3D
_MG_3301.jpg
5184x3456
focal: 28.3mm
focal: 28.58mm
pp: 71px,76px
λ: -0.18,0.18,-0.02,-0.03
p: 0.001,0.003
Selected input(s)
File name
C:\DATA\JPG_MG_3301.jpg
Model
Canon Canon EOS 550D
Width
5184
Height
3456
Features
0
Features source
Use all image features
Enable alignment
True
Enable meshing
True
Model computation – normal quality

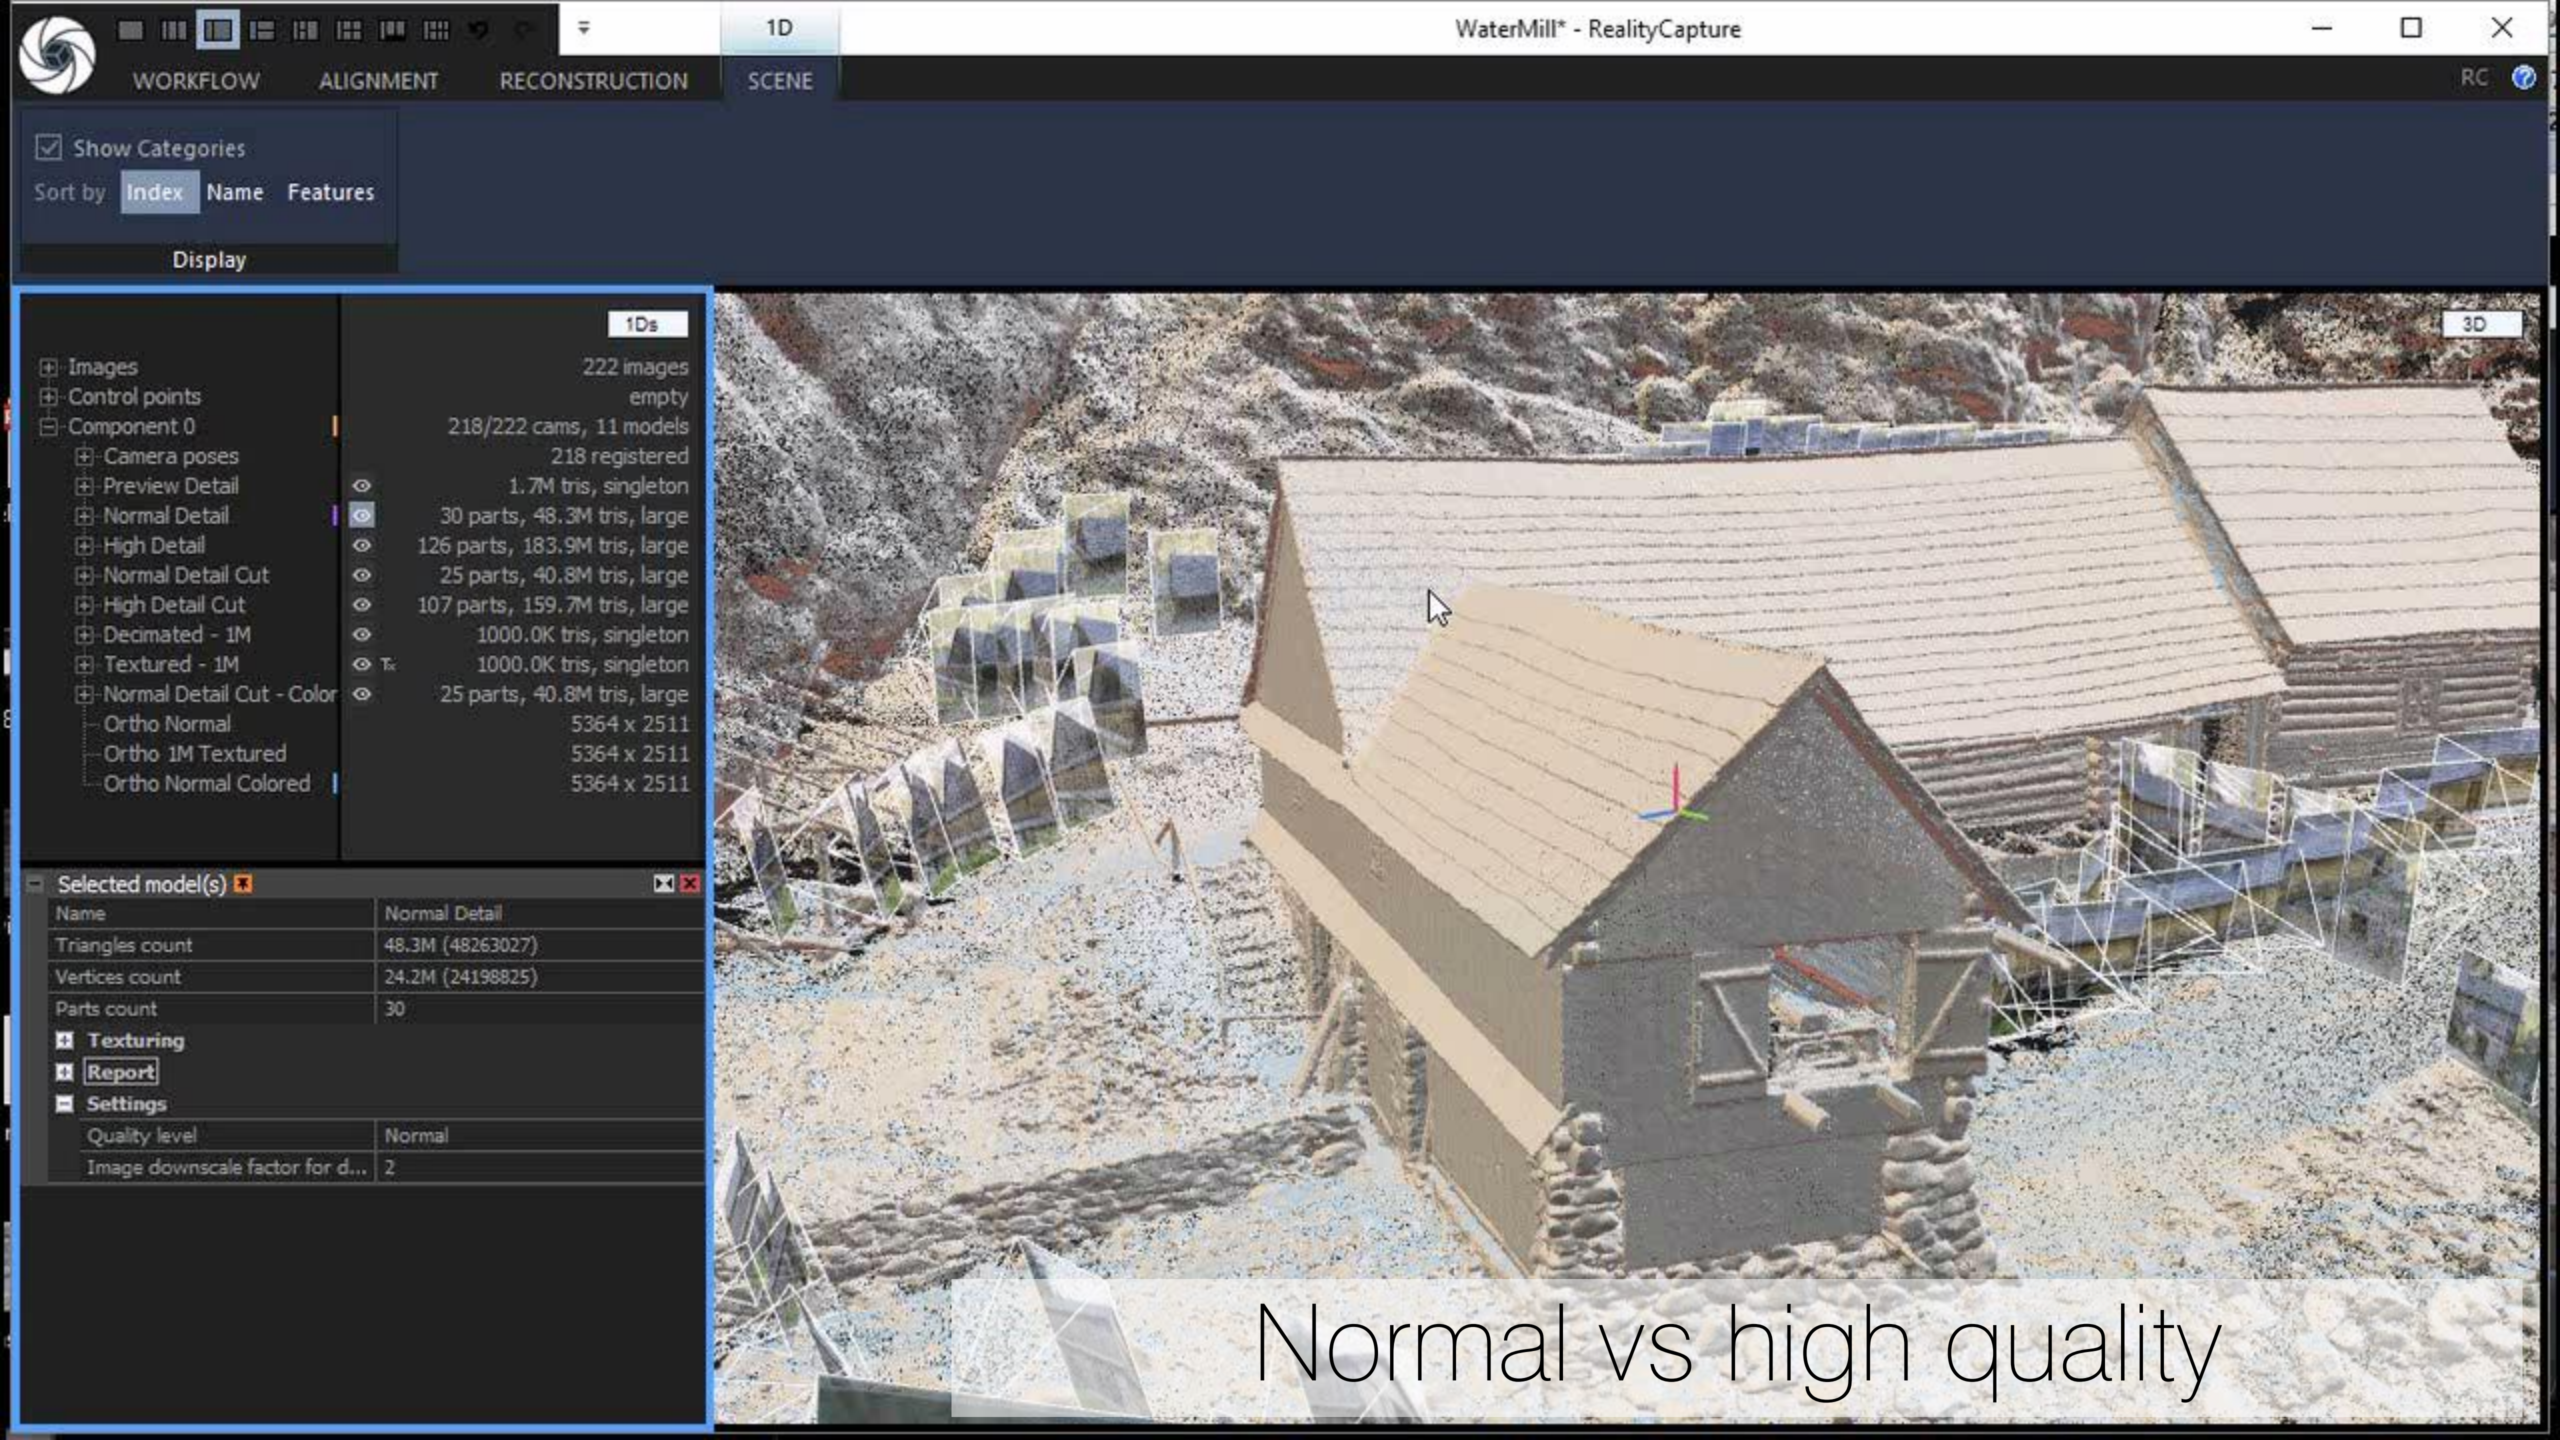
WaterMill* - RealityCapture
1D
WORKFLOW
ALIGNMENT
RECONSTRUCTION
SCENE
RC
Show Categories
Sort by Index Name Features
Display
1Ds
Images 222 images
Control points empty
Component 0 218/222 cams, 11 models
Camera poses 218 registered
Preview Detail 1.7M tris, singleton
Normal Detail 30 parts, 48.3M tris, large
High Detail 126 parts, 183.9M tris, large
Normal Detail Cut 25 parts, 40.8M tris, large
High Detail Cut 107 parts, 159.7M tris, large
Decimated - 1M 1000.0K tris, singleton
Textured - 1M 1000.0K tris, singleton
Normal Detail Cut - Color 25 parts, 40.8M tris, large
Ortho Normal 5364 x 2511
Ortho 1M Textured 5364 x 2511
Ortho Normal Colored 5364 x 2511
Selected model(s)
Name Normal Detail
Triangles count 48.3M (48263027)
Vertices count 24.2M (24198825)
Parts count 30
Texturing
Report
Settings
Quality level Normal
Image downscale factor for d... 2
3D
Normal vs high quality

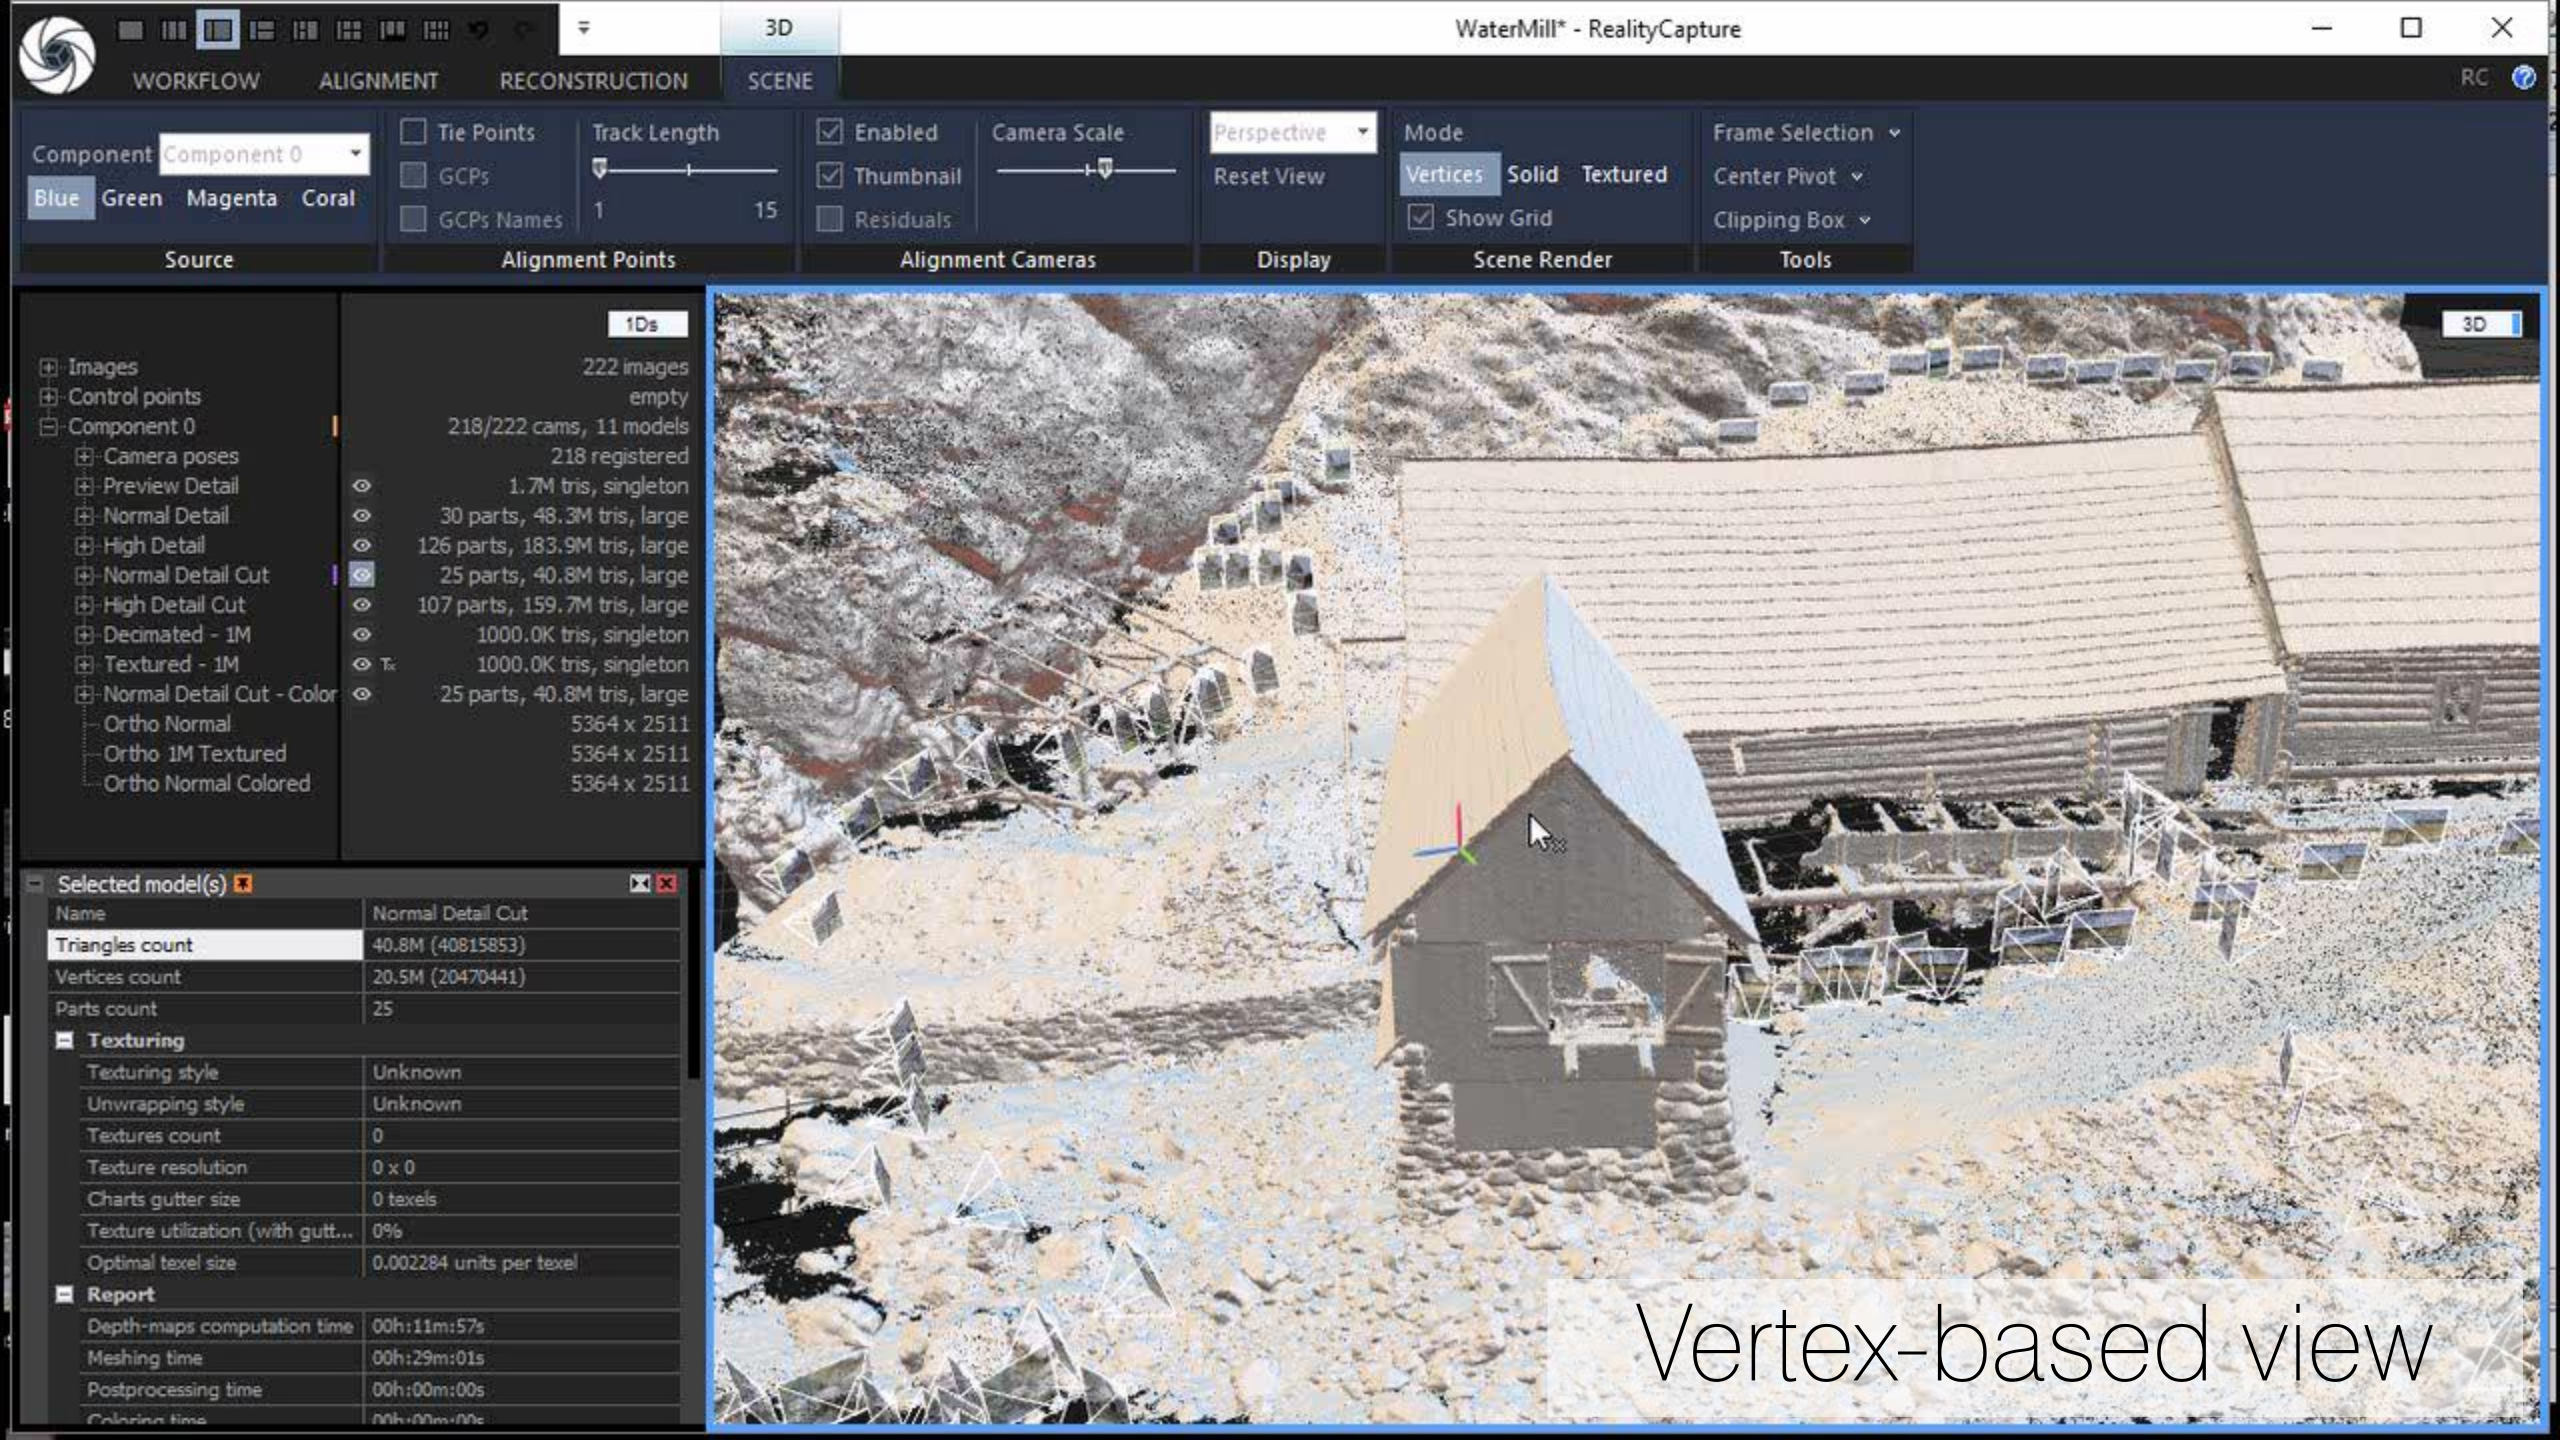
WaterMill* - RealityCapture
3D
WORKFLOW
ALIGNMENT
RECONSTRUCTION
SCENE
RC
Component
Component 0
Blue
Green
Magenta
Coral
Source
Tie Points
GCPs
GCPs Names
Track Length
1
15
Alignment Points
Enabled
Thumbnail
Residuals
Camera Scale
Alignment Cameras
Perspective
Reset View
Display
Mode
Vertices
Solid
Textured
Show Grid
Scene Render
Frame Selection
Center Pivot
Clipping Box
Tools
1Ds
Images 222 images
Control points empty
Component 0 218/222 cams, 11 models
Camera poses 218 registered
Preview Detail 1.7M tris, singleton
Normal Detail 30 parts, 48.3M tris, large
High Detail 126 parts, 183.9M tris, large
Normal Detail Cut 25 parts, 40.8M tris, large
High Detail Cut 107 parts, 159.7M tris, large
Decimated - 1M 1000.0K tris, singleton
Textured - 1M 1000.0K tris, singleton
Normal Detail Cut - Color 25 parts, 40.8M tris, large
Ortho Normal 5364 x 2511
Ortho 1M Textured 5364 x 2511
Ortho Normal Colored 5364 x 2511
Selected model(s)
Name Normal Detail Cut
Triangles count 40.8M (40815853)
Vertices count 20.5M (20470441)
Parts count 25
Texturing
Texturing style Unknown
Unwrapping style Unknown
Textures count 0
Texture resolution 0 x 0
Charts gutter size 0 texels
Texture utilization (with gutt... 0%
Optimal texel size 0.002284 units per texel
Report
Depth-maps computation time 00h:11m:57s
Meshing time 00h:29m:01s
Postprocessing time 00h:00m:00s
3D
Vertex-based view

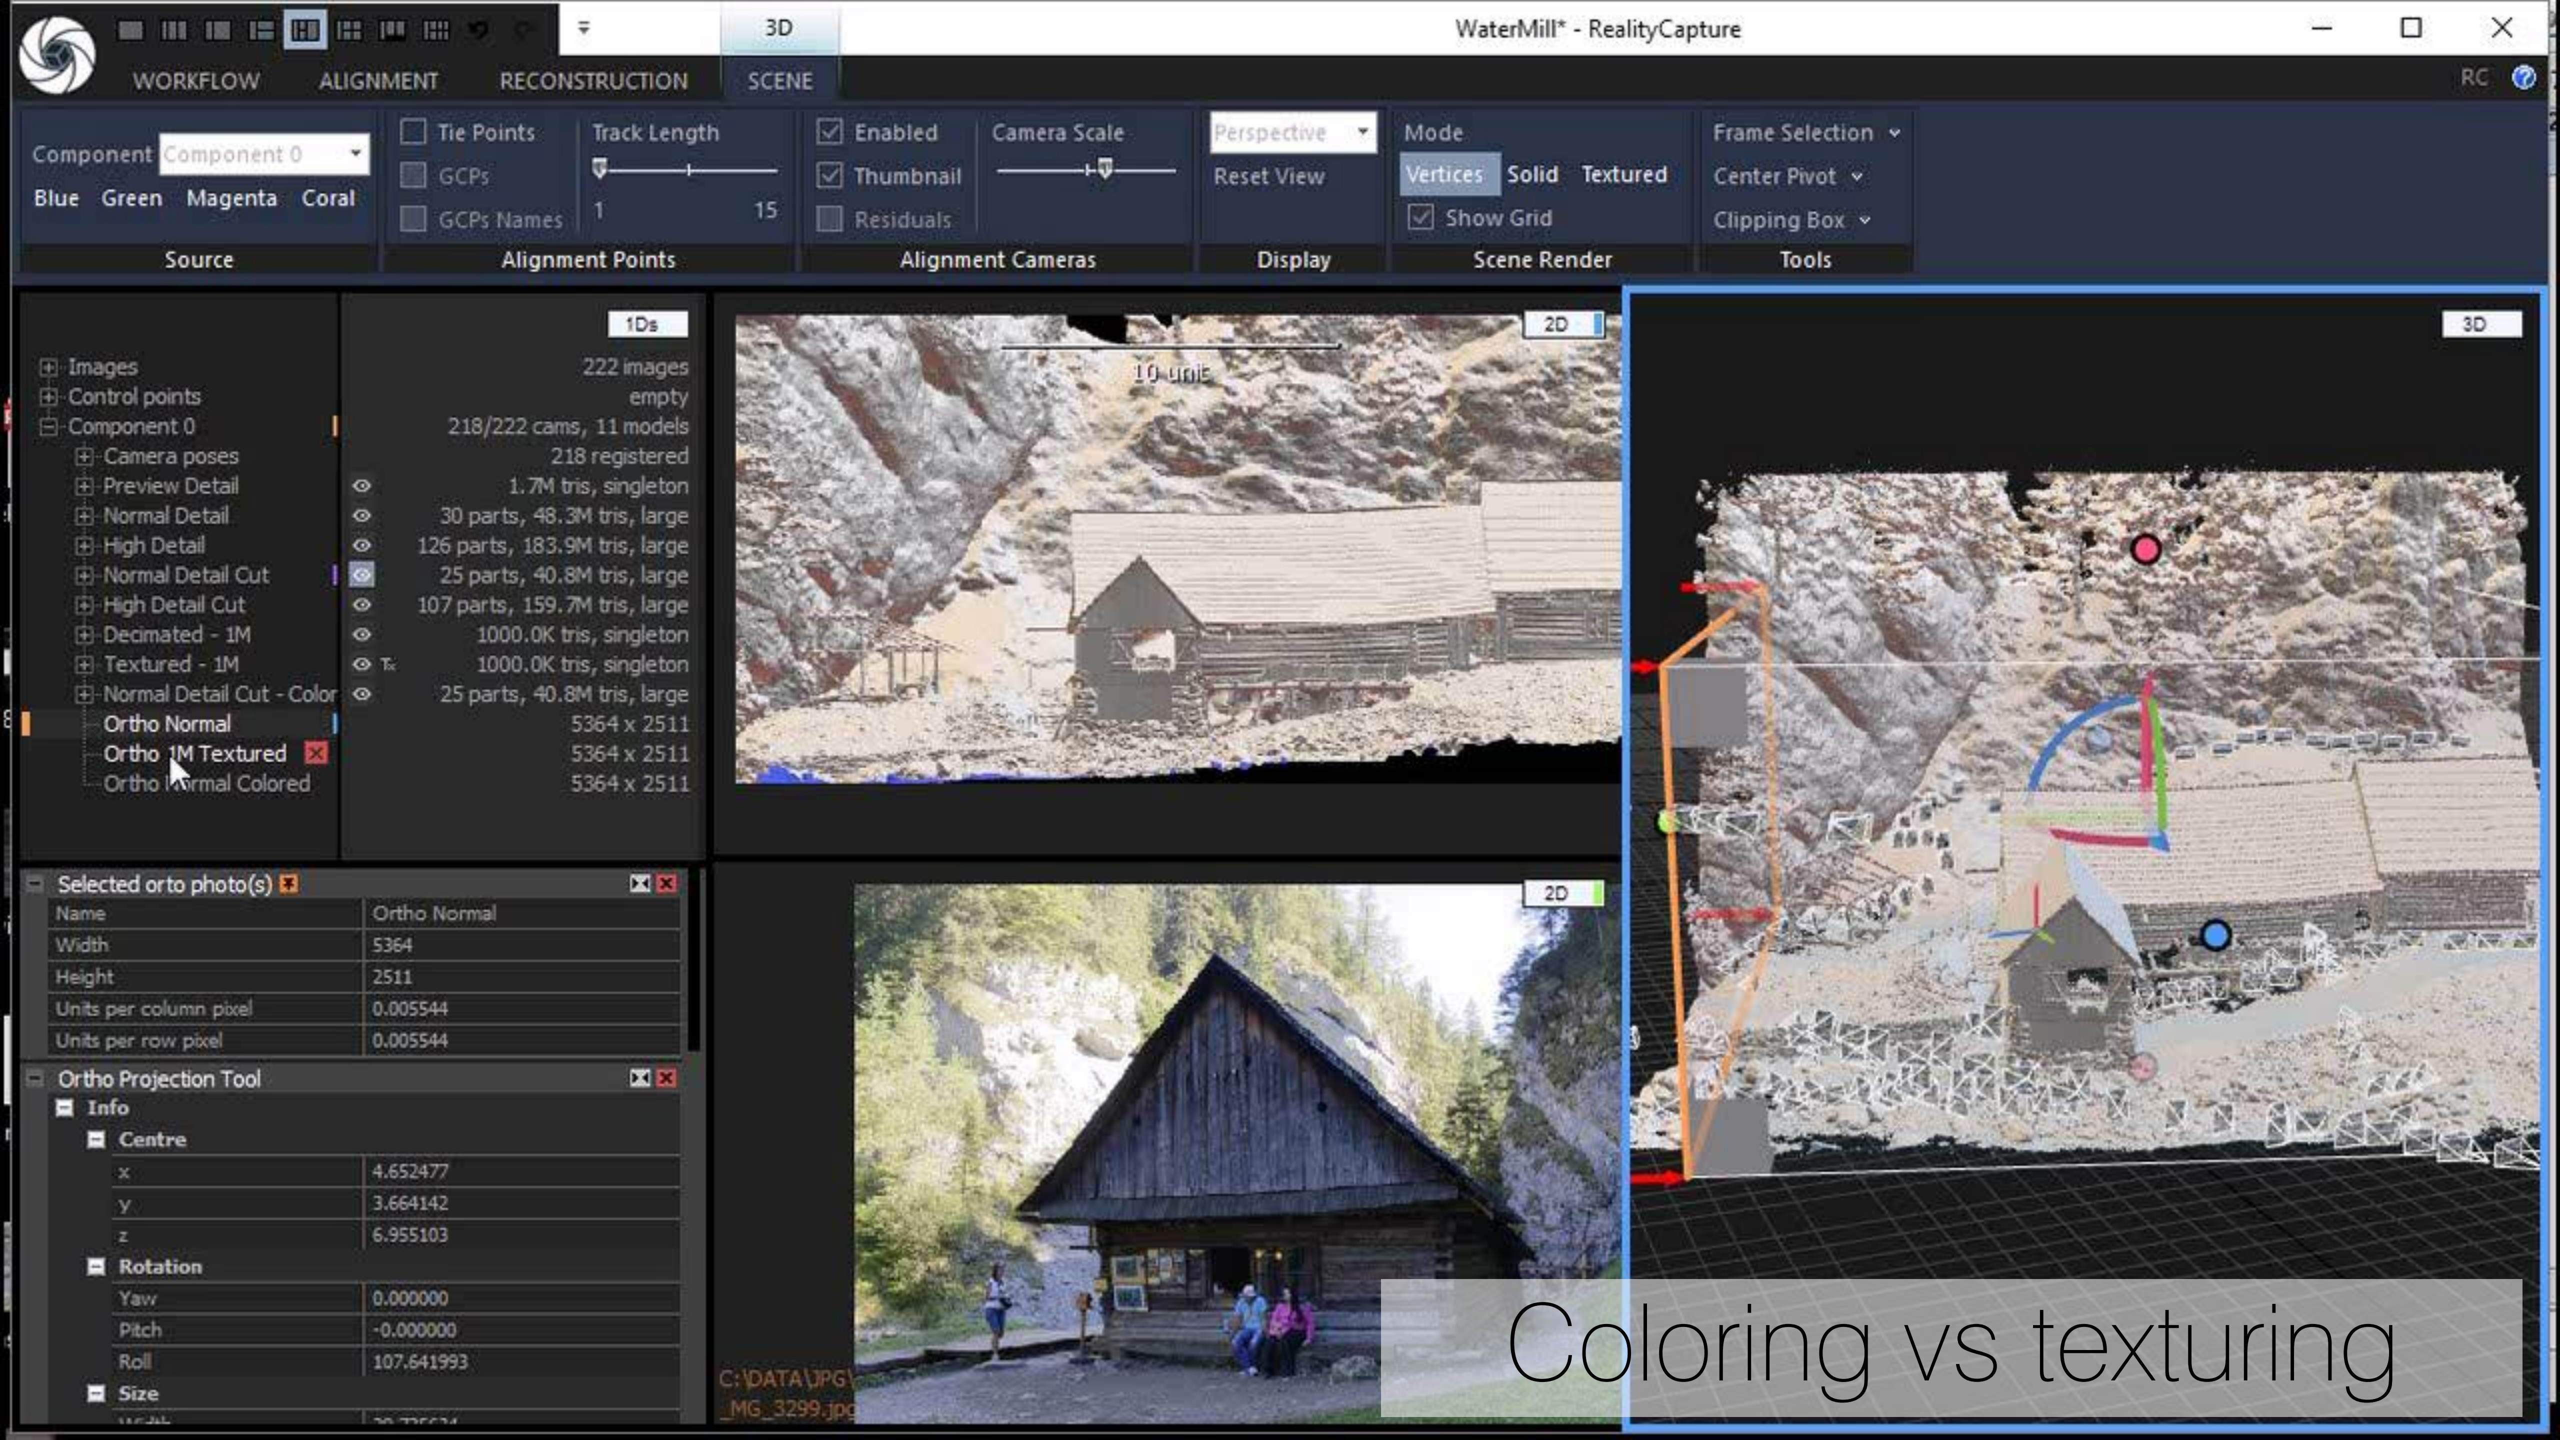
WaterMill* - RealityCapture
3D
WORKFLOW
ALIGNMENT
RECONSTRUCTION
SCENE
Component
Component 0
Blue
Green
Magenta
Coral
Source
Tie Points
GCPs
GCPs Names
Track Length
1
15
Alignment Points
Enabled
Thumbnail
Residuals
Camera Scale
Alignment Cameras
Perspective
Reset View
Display
Mode
Vertices
Solid
Textured
Show Grid
Scene Render
Frame Selection
Center Pivot
Clipping Box
Tools
1Ds
Images
222 images
Control points
empty
Component 0
218/222 cams, 11 models
Camera poses
218 registered
Preview Detail
1.7M tris, singleton
Normal Detail
30 parts, 48.3M tris, large
High Detail
126 parts, 183.9M tris, large
Normal Detail Cut
25 parts, 40.8M tris, large
High Detail Cut
107 parts, 159.7M tris, large
Decimated - 1M
1000.0K tris, singleton
Textured - 1M
1000.0K tris, singleton
Normal Detail Cut - Color
25 parts, 40.8M tris, large
Ortho Normal
5364 x 2511
Ortho 1M Textured
5364 x 2511
Ortho Normal Colored
5364 x 2511
Selected orto photo(s)
Name
Ortho Normal
Width
5364
Height
2511
Units per column pixel
0.005544
Units per row pixel
0.005544
Ortho Projection Tool
Info
Centre
x
4.652477
y
3.664142
z
6.955103
Rotation
Yaw
0.000000
Pitch
-0.000000
Roll
107.641993
Size
2D
10 unit
2D
3D
Coloring vs texturing

So how far we are from "the ultimate goal"?

Why it is so difficult?

I will show it on a simple problem:

P3P - pose of a calibrated camera from 3x 2D-3D point correspondences

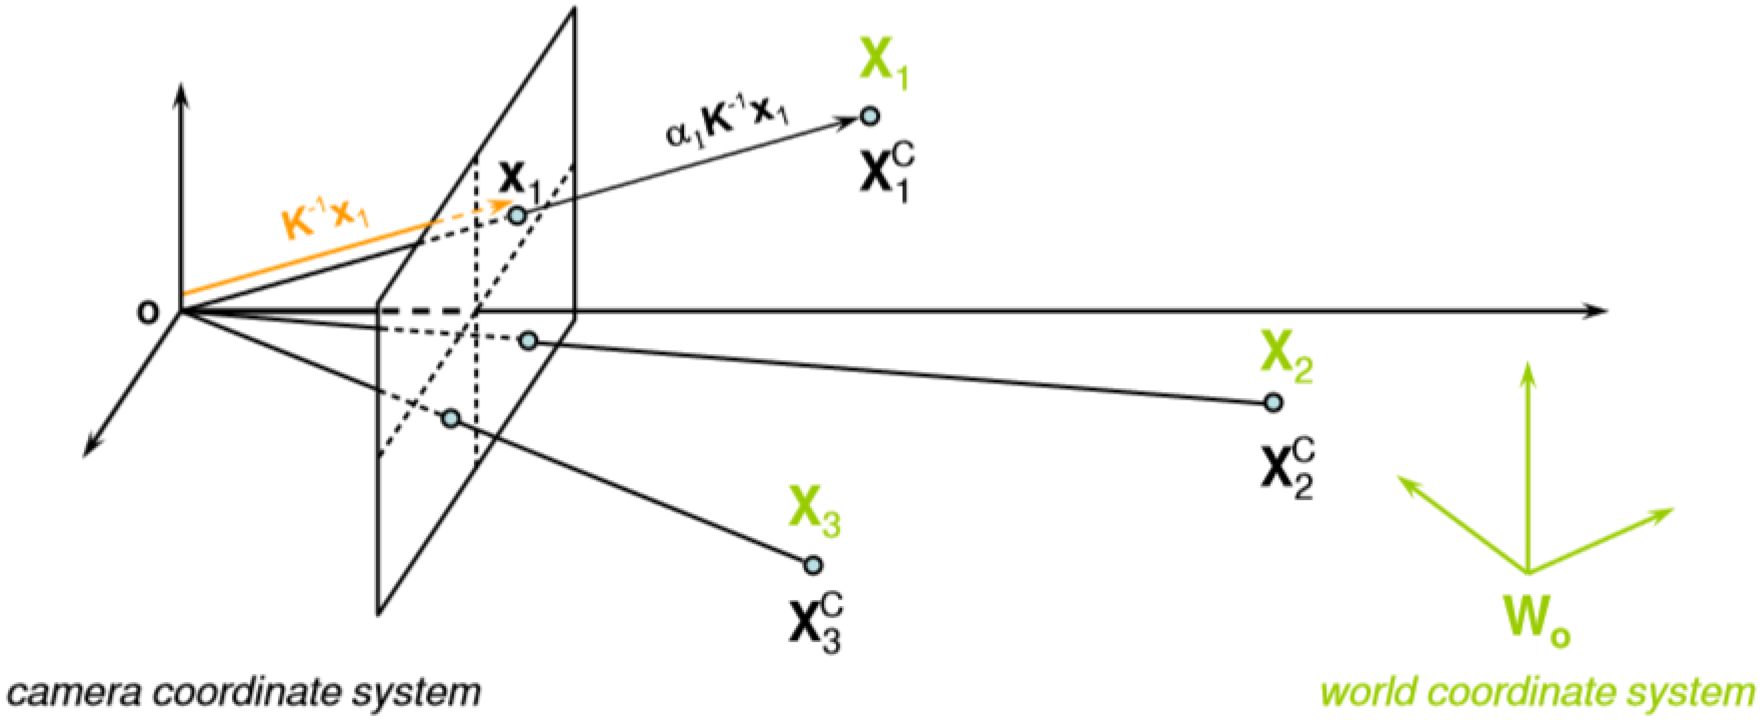

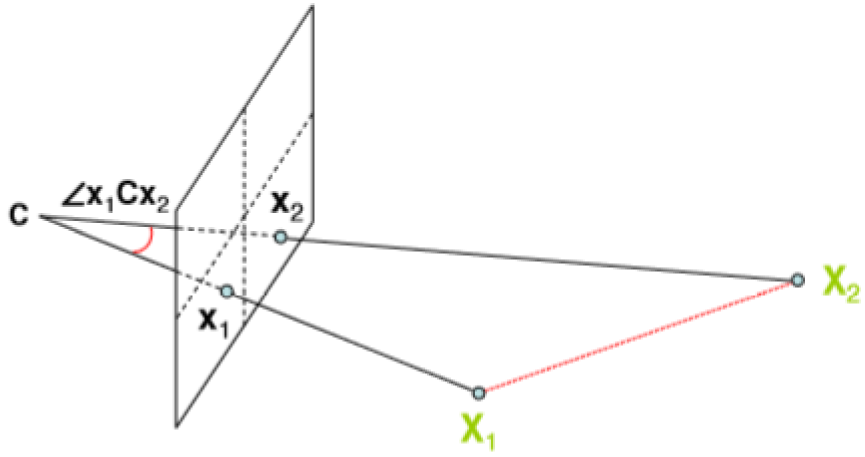

$$\cos\phi_{ij} = \frac{(\mathsf{K}^{-1}x_i)^\top(\mathsf{K}^{-1}x_j)}{\|\mathsf{K}^{-1}x_i\|\|\mathsf{K}^{-1}x_j\|}.$$

$$d_{ij} = d_{ij} = l_i^2 + l_j^2 - 2\,l_i\,l_j\cos\phi_{ij},$$

10-page solution in Fischler & Bolles RANSAC paper (1981)

Open my PhD thesis at page 34 (http://cmp.felk.cvut.cz/~bujnam1/)

You will find a **10-line** solution

10-page solution in
Fischler & Bolles
RANSAC paper (1981)

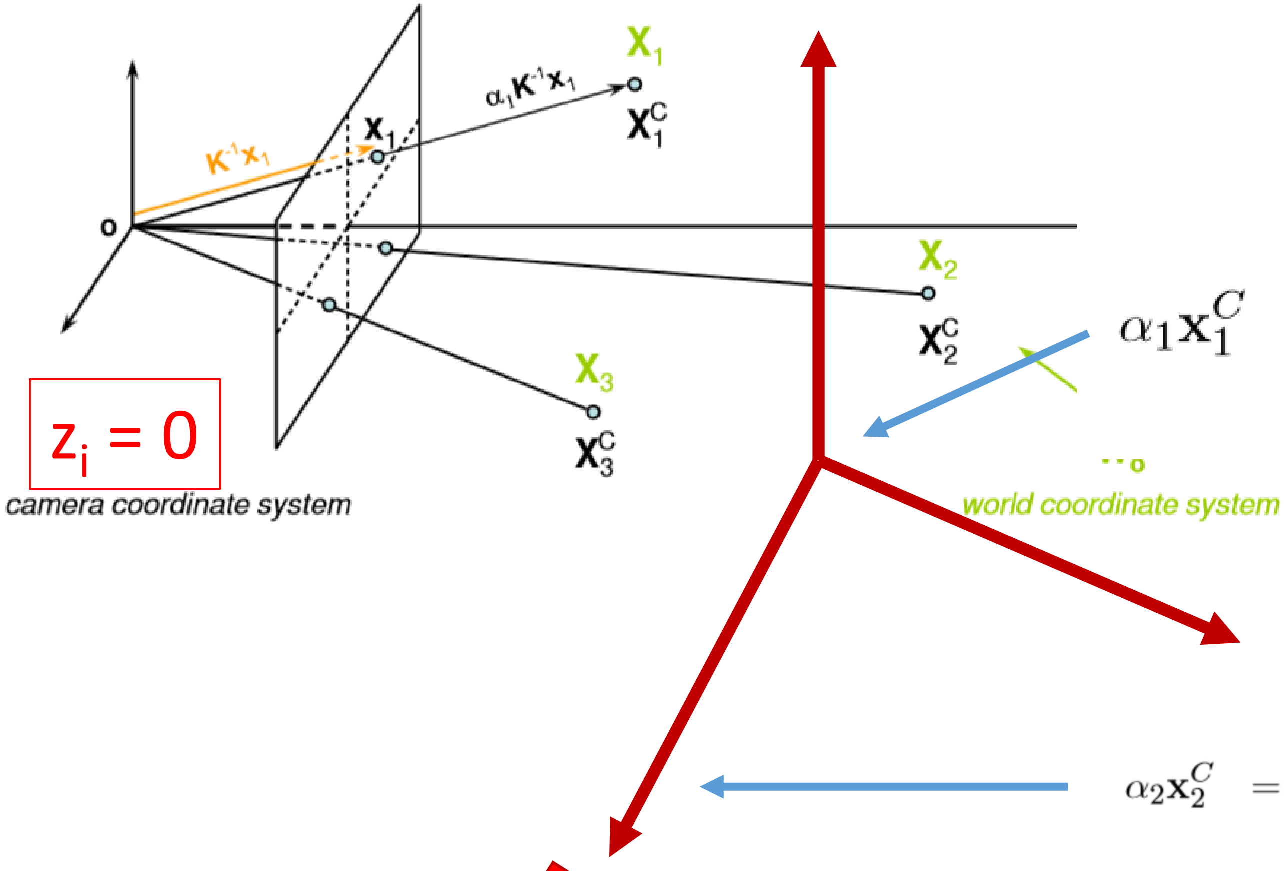

$$\alpha_i \mathbf{x_i} = \mathtt{K}\,[\mathtt{R}|\mathbf{t}] \begin{bmatrix} x \\ y \\ 0 \\ 1 \end{bmatrix}$$

$$\alpha_1 \mathbf{x}_1^C = [\mathbf{r}_1 \mathbf{r}_2 \mathbf{t}] \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = \mathbf{t},$$

$$\alpha_2 \mathbf{x}_2^C = [\mathbf{r}_1 \mathbf{r}_2 \mathbf{t}] \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} = \mathbf{r}_1 + \mathbf{t} = \mathbf{r}_1 + \alpha_1 \mathbf{x}_1^C.$$

$$\begin{aligned} \mathbf{r}_1 \perp \mathbf{r}_2 &= \left(\alpha_2 \mathbf{x}_2^C - \mathbf{x}_1^C\right) \perp \left(\alpha_2 \mathbf{c_a} + \alpha_3 \mathbf{c_b} + \mathbf{c_c}\right), \\ \|\mathbf{r}_1\| = \|\mathbf{r}_2\| &= \|\alpha_2 \mathbf{x}_2^C - \mathbf{x}_1^C\| = \|\alpha_2 \mathbf{c_a} + \alpha_3 \mathbf{c_b} + \mathbf{c_c}\|. \end{aligned}$$

$$\downarrow$$

$$\mathtt{M}'(\alpha_2)\mathbf{v} = \begin{bmatrix} 0 & c_{11}\alpha_2 + c_{01} & c_{20}\alpha_2^2 + c_{10}\alpha_2 + c_{00} \\ d_{02} & d_{11}\alpha_2 + d_{01} & d_{20}\alpha_2^2 + d_{10}\alpha_2 + d_{00} \\ c_{11}\alpha_2 + c_{01} & c_{20}\alpha_2^2 + c_{10}\alpha_2 + c_{00} & 0 \end{bmatrix} \begin{bmatrix} \alpha_3^2 \\ \alpha_3 \\ 1 \end{bmatrix} = \mathbf{0}. \tag{5.30}$$

Now, a non-trivial solution exists if and only if the determinant of matrix $\mathtt{M}'(\alpha_2)$ vanishes, i.e.

$$det\left(\mathtt{M}'(\alpha_2)\right) = 0. \tag{5.31}$$

>> 4^{th} degree polynomial - can be solved in closed form

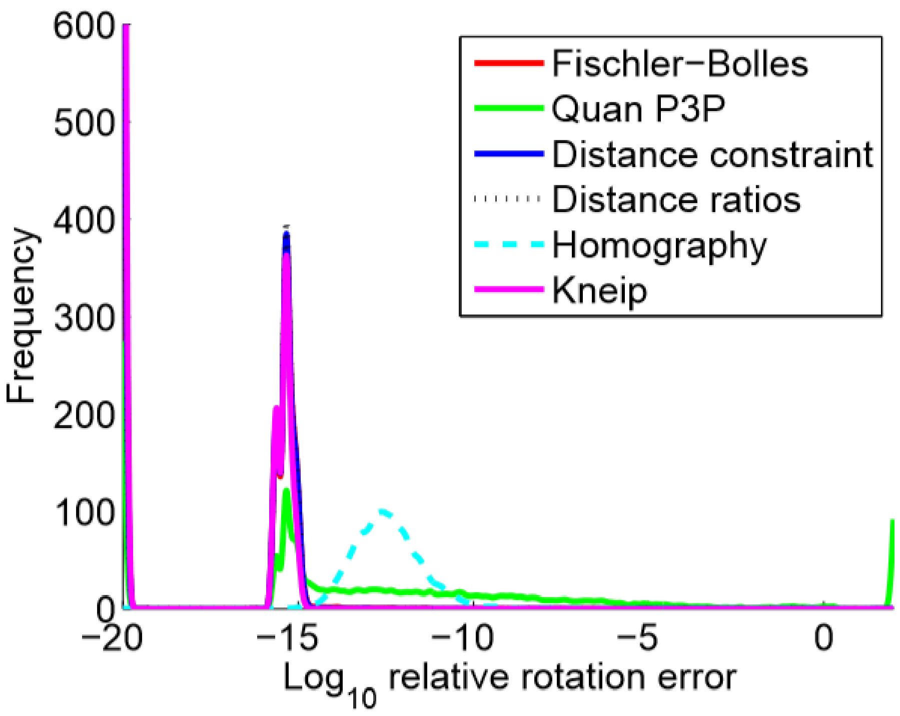

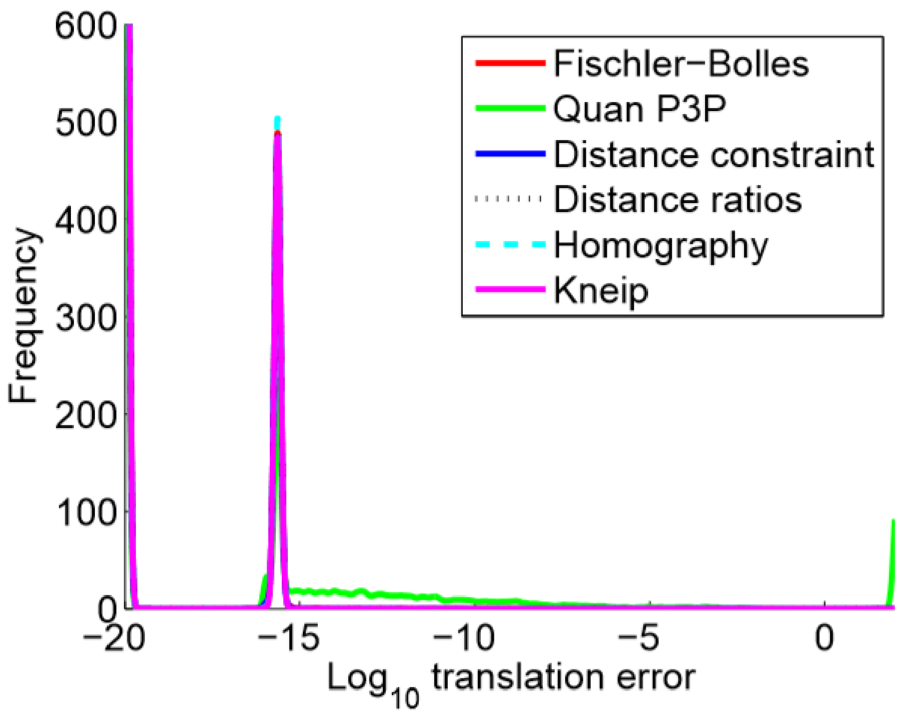

If properly implemented!

Actually –

closed form 4^{th} degree polynomial uses reduction to 3^{rd} degree polynomial which is then solved using inverse to a goniometric function

- numerically not precise (especially around 0)
- magnitude differences between translation vector vs. rotation part
- results in non-orthogonal rotation matrix

Blind use of unknown algorithms can kill your pipeline!

Many glitches:

You need to understand the method

You need to understand its weakness

You need to understand how to implement it

You need to implement it damn good!

But why to re-invent the wheel?

Keep focused on the area where you are the best

Reality Capture SDK - RealityCapture Engine

- Registration
- Meshing
- Tools

C++
SDK

- develop for price of a lemonade
- deploy for free

- Windows, Linux
- C++

This winter!

RealityCapture SDK in hands of nvidia for 1 month

Google Nvidia + RealityCapture

I can imagine you already have full head of ideas how to use it

Opened topics?

- Automatic image masking – trees in the wind, growing grass, …
- Moving object masking,
- Water, glass, mirrors detection,
- Clever simplification (context sensitive)
- High level image clustering for super huge reconstructions
- Real time experience
- …

Thank you!

Looking for a job? jobs@capturingreality.com