Using Single Photo Constraints in PhotoModeler

Contents

Introduction

Solving a scene from photographs requires knowledge of the camera internals (e.g., focal length, format size) and externals (e.g., position and angles). With multiple photographs, PhotoModeler can often solve both the internal and external camera parameters without external knowledge. When there is a single photograph, additional information is needed to solve the scene. This additional information is generically called "constraint data." Constraint data tells PhotoModeler something else you know about the scene. Often this information consists of 3D control points from a survey, a laser scan, or another PhotoModeler project.

When you need to solve a scene or create an object model from a single photograph with no 3D control points, PhotoModeler provides the Axes Constraint system. This system works when the object or scene contains several parallel and perpendicular lines. This is based on the principles of vanishing points.

Axes Constraint projects may not be as accurate as multi-photograph or single-photograph-control-based projects, but they can provide useful data when no other method is available.

Vanishing Points

The Axes Constraints system in PhotoModeler is based on the concepts of vanishing points and n-point perspective. The diagram below illustrates the different forms. The number of vanishing points controls which parameters PhotoModeler can solve. Three-point perspective (with 3 vanishing points) allows for the largest number of parameters to be solved.

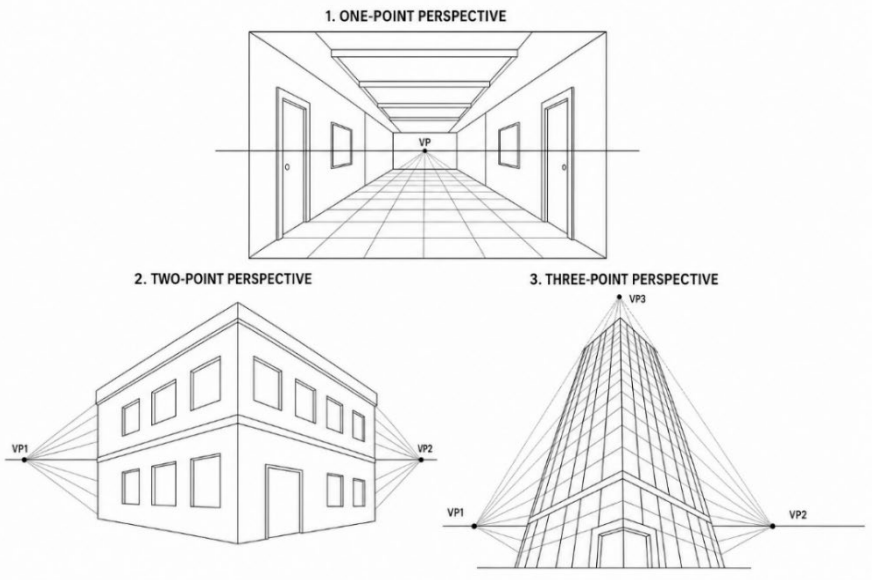

PhotoModeler does not refer directly to the vanishing points (or show the values) but instead the user defines the geometry through axes constraints. That is, the user tells the system what lines drawn on the photograph fall on each axis. For example, in the two-point perspective diagram, the vertical building lines are defined as Z-axis lines, and the horizontal lines are defined as X-axis and Y-axis lines. PhotoModeler uses a right-handed coordinate system:

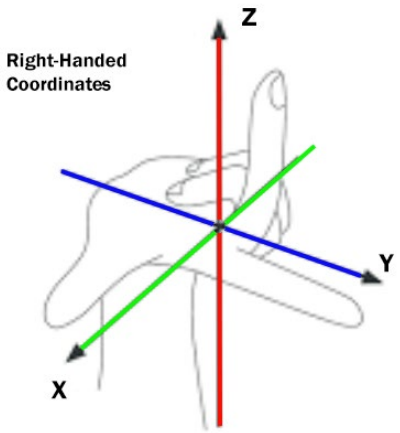

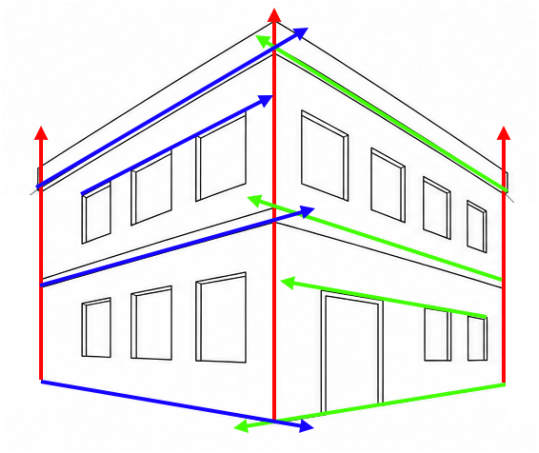

Step By Step Guide

Overview of Steps

The overview of the steps when solving axes constraint projects:

1. Start a Manual Project.
2. Load the Photo and choose the Camera solution.
3. Mark Lines on the Photo.
4. Apply the Axis Constraints.
5. Process and Review.
6. Scale the Project.
7. View, Measure, and Export

Basic Steps

The project used in this section is a photograph of square toy blocks, the perfect example of parallel and perpendicular lines in three-point perspective. More real-world examples are shown later in this document.

If you wish to follow along, you can download the photograph of the toy blocks here.

1. Start a Manual Project

Start PhotoModeler (this process works in Standard and Premium), choose 'Start a new project', and choose a 'Manually Marked Project'.

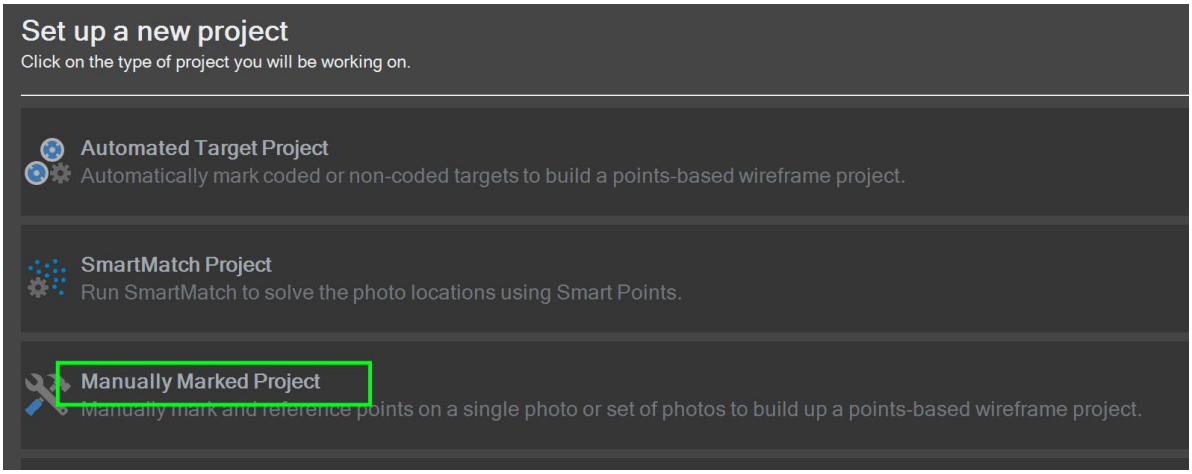

2. Load the photograph and choose the Camera solution

Find and load the photograph.

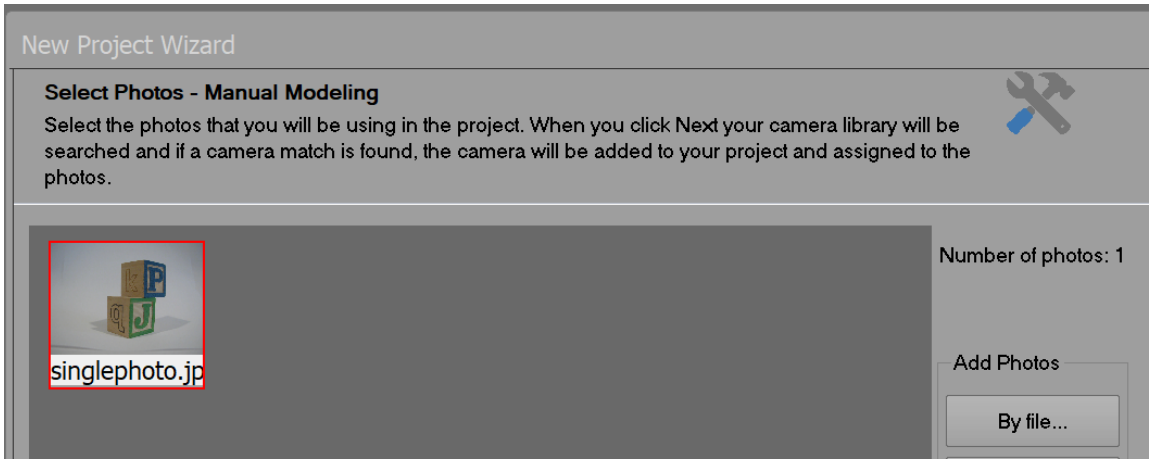

In the next Wizard dialog, choose the 'unknown camera solved by constraints' option:

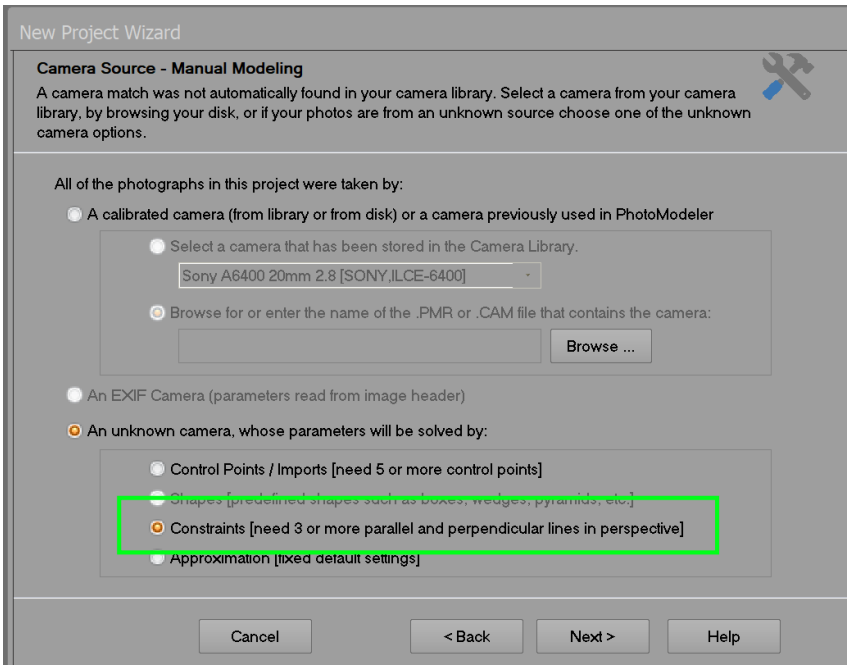

3. Mark Lines on the Photo

Open the photograph and choose Lines mode:

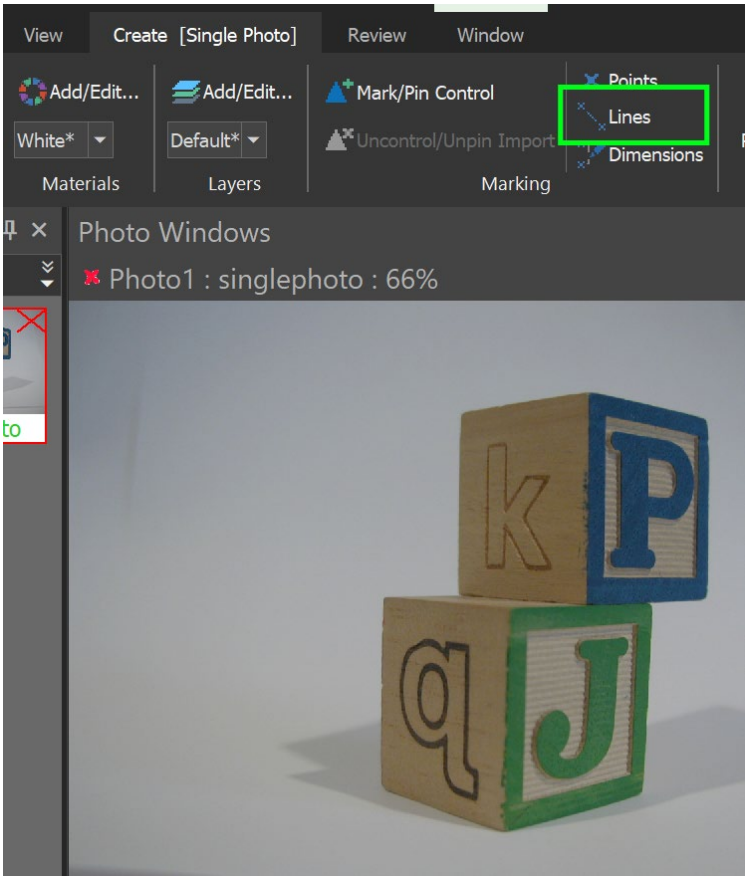

Constraints work with Lines and Edges (Edges are lines that do not have defined end points), but for most purposes with these projects, Lines will suffice.

Draw lines along the distinct edges of your object or scene (for example, the edges of the blocks). Try to connect the Lines at Points, as this helps build the complete model. Here all the visible lines have been drawn (blue dashed lines):

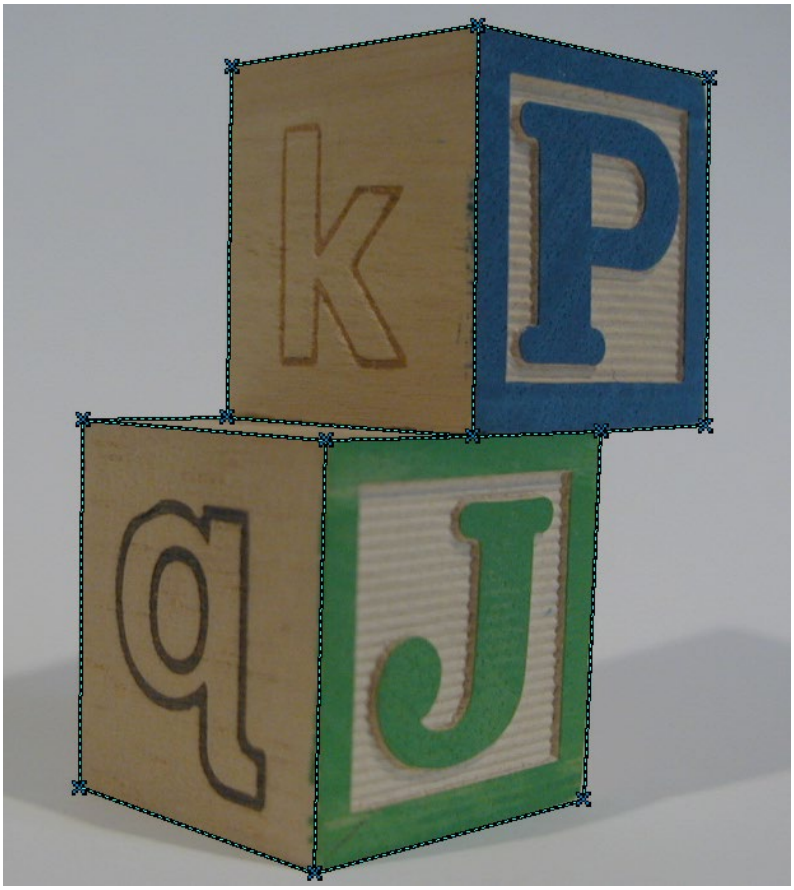

Sometimes it can be more efficient to draw all the lines when the whole photograph can be seen and then zoom in and reposition the points for better alignment using the select cursor.

The image above shows the Lines displayed with the Thick Line Visibility property enabled (see the Tips section below).

In this project, all Lines connect at corners. In some projects you may have lines or edges in the photograph that do not connect to each other but are still parallel. It can be helpful to still draw these. Further information on connection between lines is described below.

4. Apply the Axis Constraints

The next step is to tell PhotoModeler which lines are parallel and perpendicular and which axis they belong to. This is done in Constraints Definition mode. This mode can be accessed by a) using the drop-down menu in the External Geometry section of the "Create [Single Photo]" toolbar, b) typing "Constraint" into the Tool search box at the top right, or c) using the 'C' shortcut key on the keyboard. The C key will be used frequently when working with constraint projects.

The Constraints toolbar (green box below) will open:

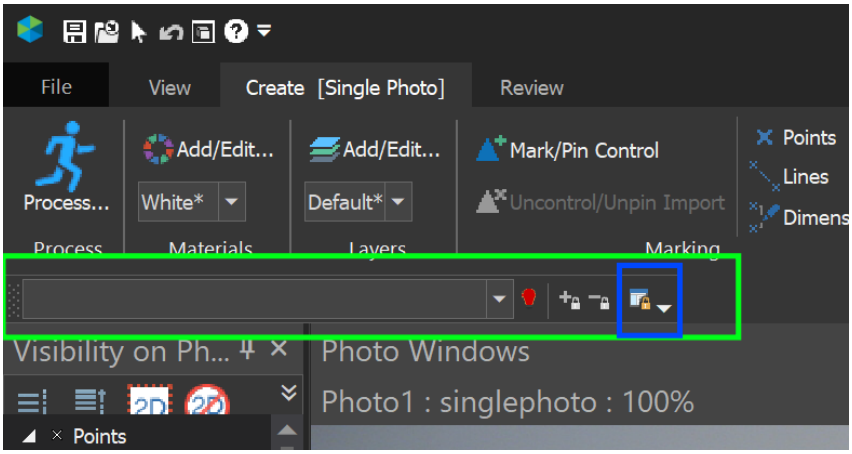

The drop-down box on the left lists all Constraints in the project, and the buttons on the right are used to work with those Constraints. The first button to use is the Constraints Dialog (blue box above).

Press the Add button on the Constraints dialog to open the Constraints Type dialog:

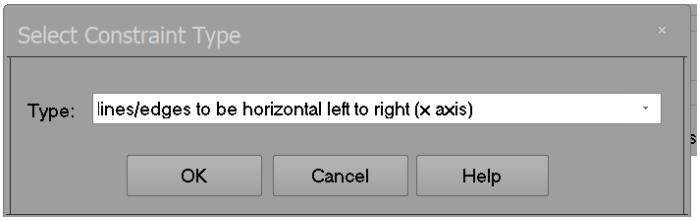

For this project we will add the three axes' constraints: x axis, y axis, and z axis. All Axes Constraint projects will have at least two of these. The left side of the Constraints dialog will look like this:

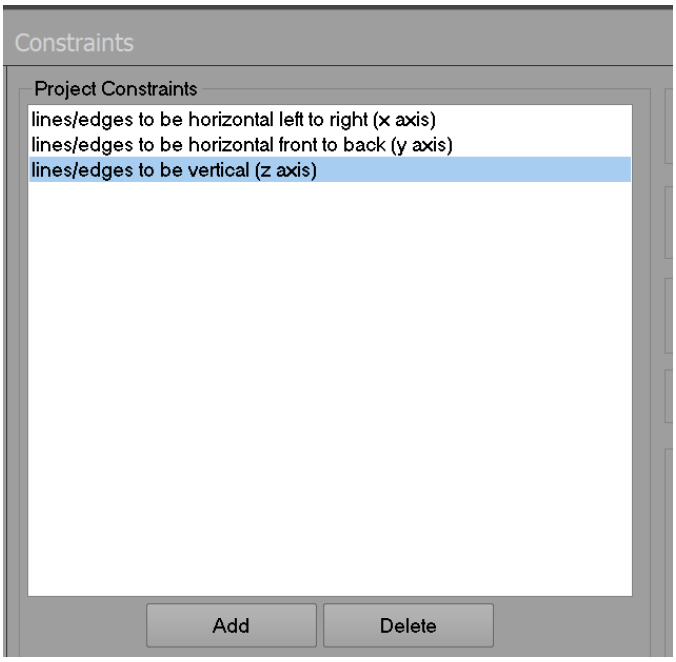

In the Constraints toolbar, you can now choose which Constraint to add Lines to (green box below). Below you start with the vertical Lines, so you choose the Z axis Constraint. Then you can select Lines in the photograph. You can do this one Line at a time or all Lines at once. Here all vertical Lines have been selected (they are all red).

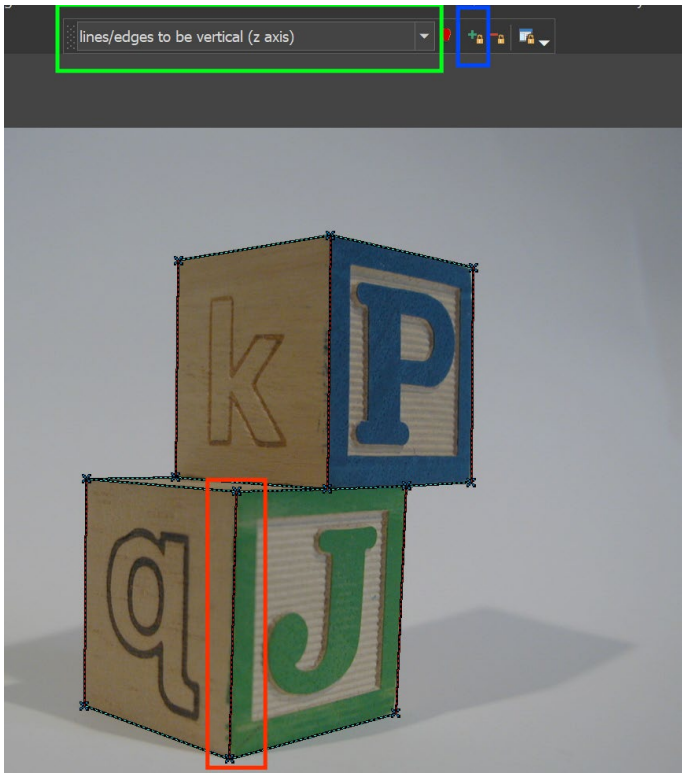

Then press the ‘Add object(s) to constraint’ button (blue box above). All the Lines that are part of the Z axis Constraint will now have the green highlight on the photograph. If the Line is not green, it is not part of the Constraint (a good time to check and review):

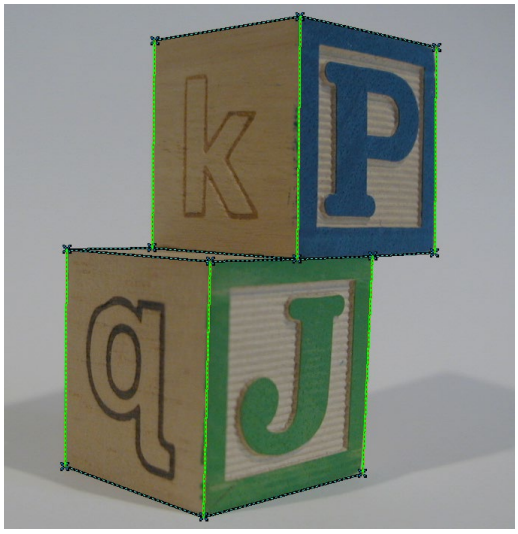

Repeat this for the X-axis (remember the right-hand rule; see the diagram above). Here the X axis is for the Lines left to right.

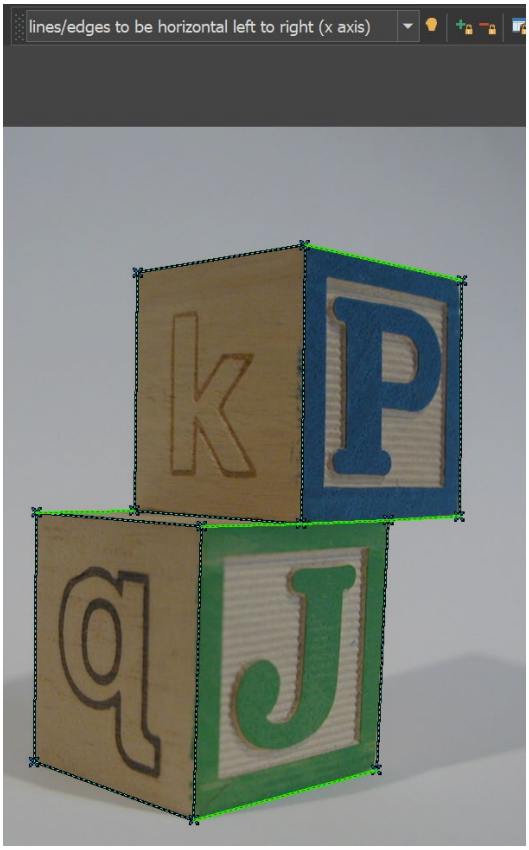

Lastly repeat this for the Y-axis. Here the Y axis is for the Lines front to back on the blocks.

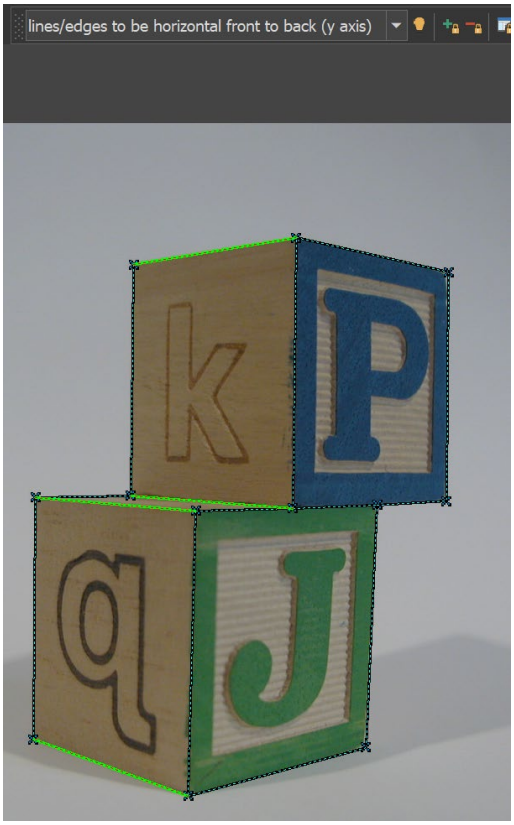

This completes the definition of the Lines and their Axis Constraints.

5. Process and Review

Press the Process button on the toolbar to open the Processing dialog:

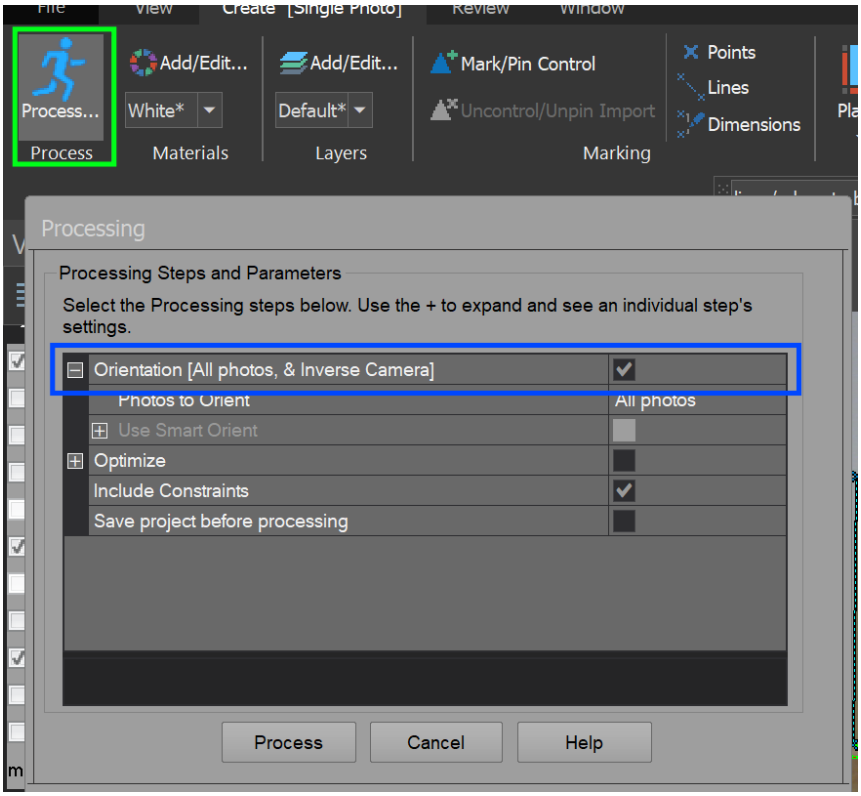

No changes are needed here by default. The 'Orientation & Inverse Camera' is all that is needed. When processing is completed, the Processing Finished dialog will appear:

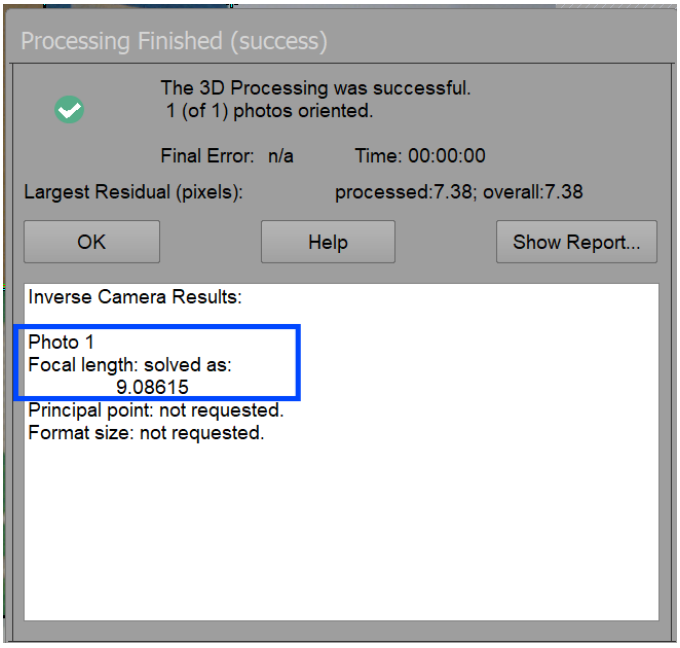

For a single-photograph, constraint-based project, the Inverse Camera results are shown (e.g. which camera parameters were solved).

It is important to review the result to see if the solution looks correct. Incorrect Line markings, incorrect Constraint assignments, or a camera that is too far from a pinhole camera model may result in poor solutions. Before you depend on the solution, it is important to review it carefully.

First check the maximum residual (as shown on the bottom right of the program window):

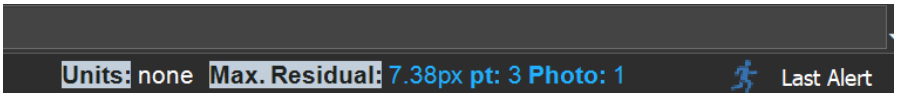

These projects are not as accurate as other PhotoModeler projects so you cannot expect a low residual here. However, if it climbs over 20 pixels, you should investigate. You can click on the Max Residual value on the toolbar to select the point with the highest residual error to see if it can be adjusted and improved. You can also open the Points Table as well to review all the residuals.

Then open the 3D Viewer, turn on the Camera Stations, and review the Camera position and the solved 3D model.

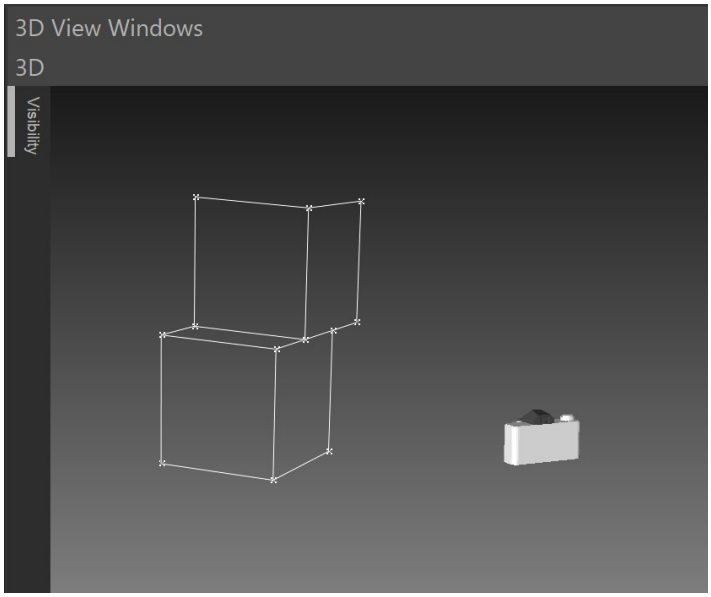

Lastly, turn on 3d projections on the Photo to see how the solved 3d data aligns with the photograph.

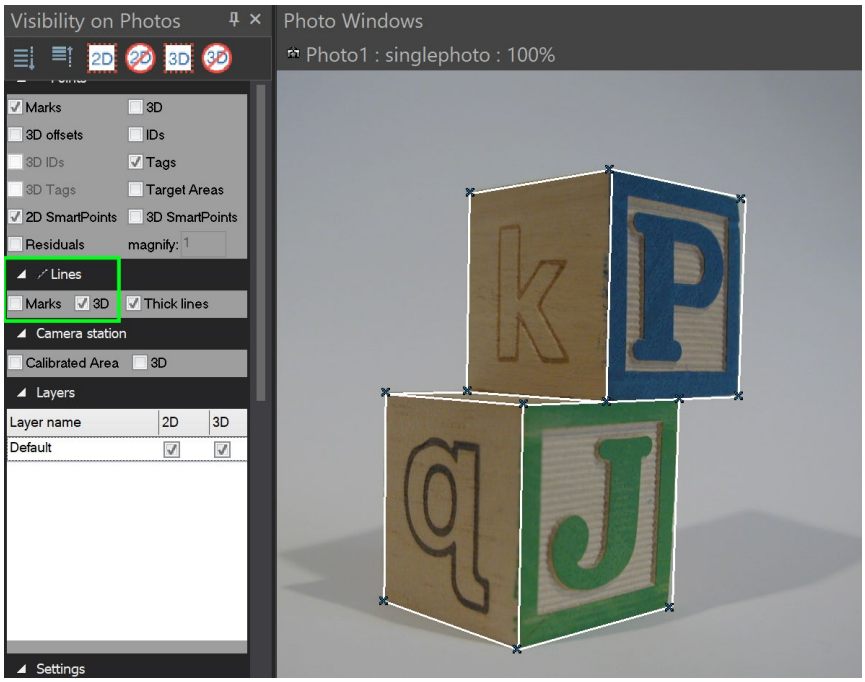

In the figure above, in the Visibility on Photos pane, the Line Marks have been turned off and the Lines 3D projection has been turned on. This makes it easier to see the solved 3D model on top of the photograph to check alignment. See the Tips below for ways to make the projection easier to see.

If you have other measurements in the scene or object, it is useful to use them as verification. Get the measurement from PhotoModeler and compare it to the known value. This is especially useful for proving accuracy in forensic cases.

6. Scale the Project

By this step you may have what you need. But if you wish to make measurements or export the data for external measurement and diagramming, you will need to add a scale.

Setting up scale is documented elsewhere (e.g. this video), so we will not describe all the steps here. One line at the bottom of a block was selected and a scale of 2cm was added.

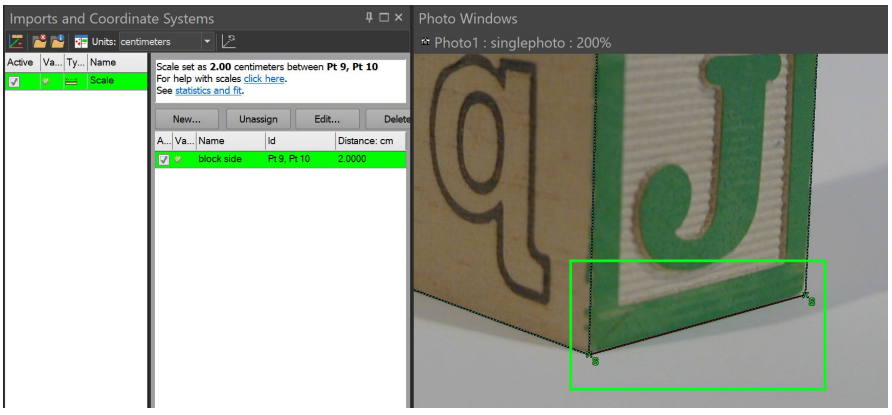

7. View, Measure, and Export

You can view the data in the 3D Viewer and in Tables, and you can export it. You can also use the Measure or Dimensions tools in PhotoModeler. A section below describes how you can add Surfaces and Surface Draw/Outlines to the models for additional detail. Here we have added two Dimensions and opened the Measurements pane.

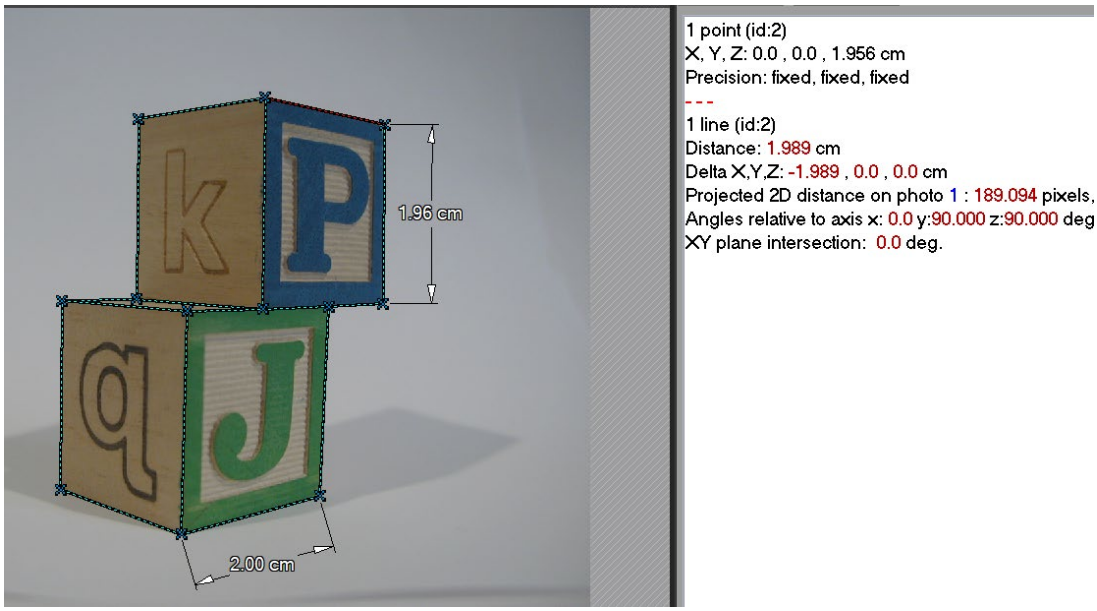

What Becomes 3D and the Surface Constraint

An important concept with these single-photograph constraint solutions is how objects become 3D (i.e. have a solved XYZ position). With the Axis Constraint types, as described above, Lines must be connected to each other to be 3D.

If there is a Line that is part of a Constraint but not connected to other Lines, it will contribute to the camera solution but will not be 3D itself (it will not appear in the 3D View and you cannot measure or dimension it, etc.). Think of this as a propagation of information. PhotoModeler chooses the largest collection of connected Lines and selects an origin point. It then propagates information along the connected parallel and perpendicular Lines to generate the XYZ positions of other Points.

Here is an example of an interior scene (a simulated crime scene captured with a surveillance camera). Several parallel lines have been drawn on partitions and furniture. The cabinet has the most connected Lines, so it is solved (blue box). Notice that the other Lines (green boxes) are not 3D and do not appear in the model. They are included in a Constraint definition but are not connected to the cabinet Lines, so they do not solve (though they contribute to the camera internal and external parameters).

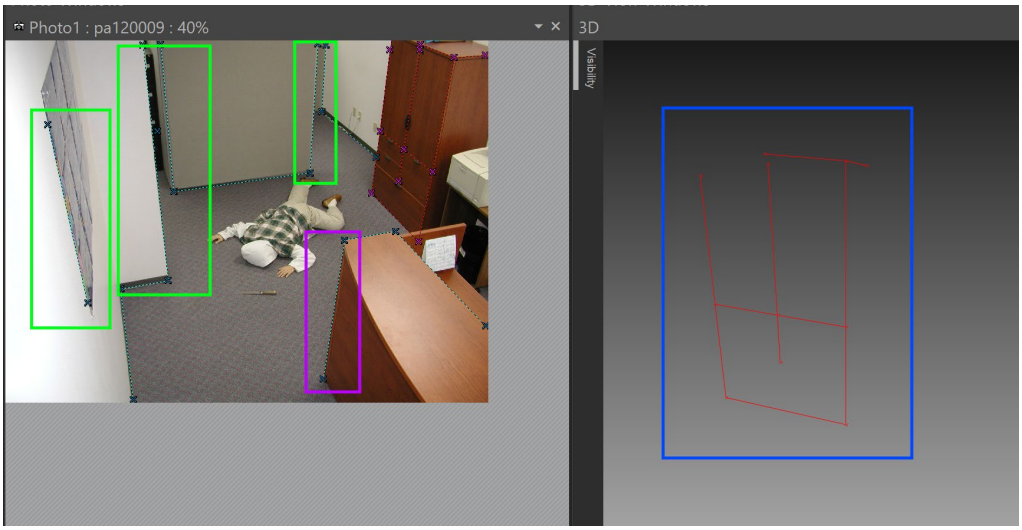

There are a couple of items we wish to solve in this scene: the height of the desk, and the position of the knife relative to the body. We use an additional Constraint to pull most of the remaining lines into the solution (and give us a Surface to draw on). First, a Path Surface is created using all the points on the floor (esp. including the two points on the cabinet that were solved before). The Surface is added to a 'surfaces to be horizontal (xy plane)' Constraint. When the scene is processed again, the Surface is computed and all Lines connected to the Surface can now be computed.

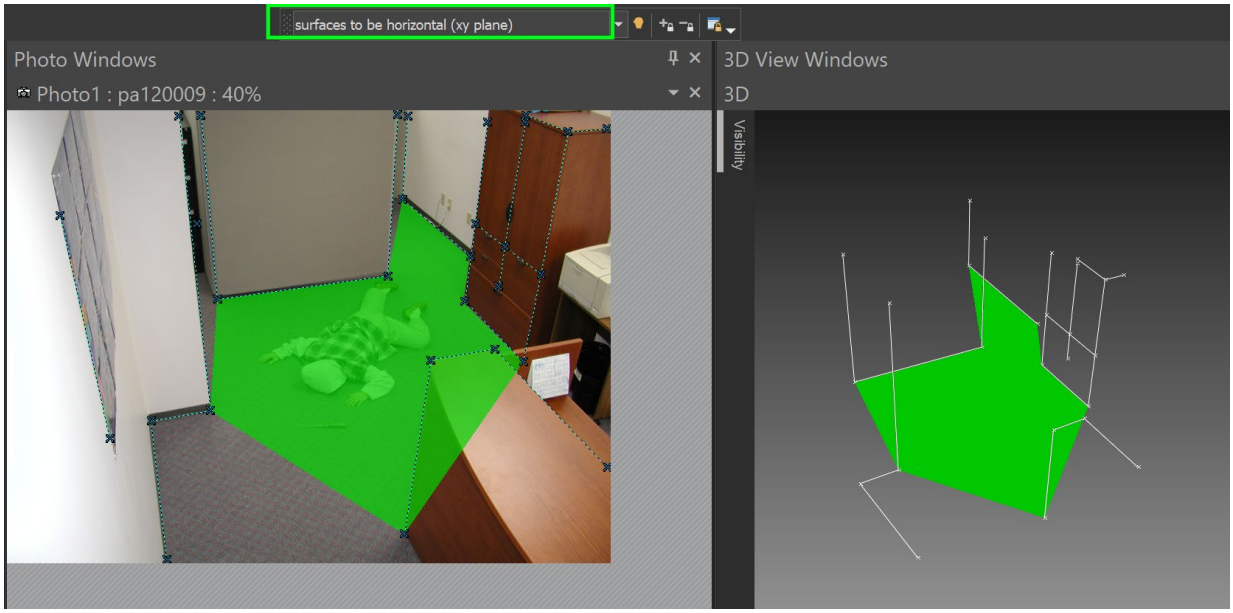

With the other Lines solved and a Surface known, we can use Surface Draw for the body and knife as well as get the height of the desk:

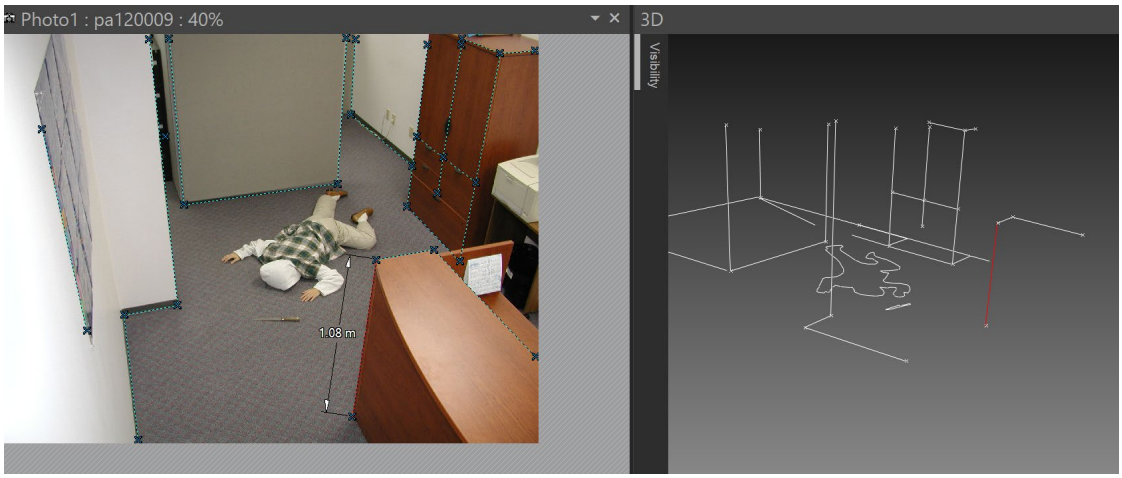

Much of the work in single-photograph, constraint-based projects is determining whether you can obtain the measurements you need from what you can define. For example, in the project above if we needed the height of the printer (right middle in the photograph) above the floor, that would be difficult to get because the point under the printer on the floor is obscured by the desk. We could get the height of the poster on the left wall or the height of the wall plates on the right wall by drawing lines to see where they intersect the ground plane and working from there. Sometimes you may be willing to make assumptions to obtain any data, but you must always be aware of those assumptions and the resulting lack of accuracy.

Tips

Here are a couple of tips for setting up the user interface to make these projects easier to work with.

Thick Photo Lines: To make the Lines easier to see on top of the photos, use the "Thick lines" option on “Visibility on Photos”.

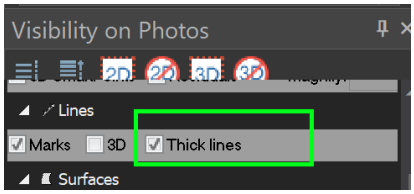

The thickness of those lines is set in Preferences under User Interface / Image / Photo display:

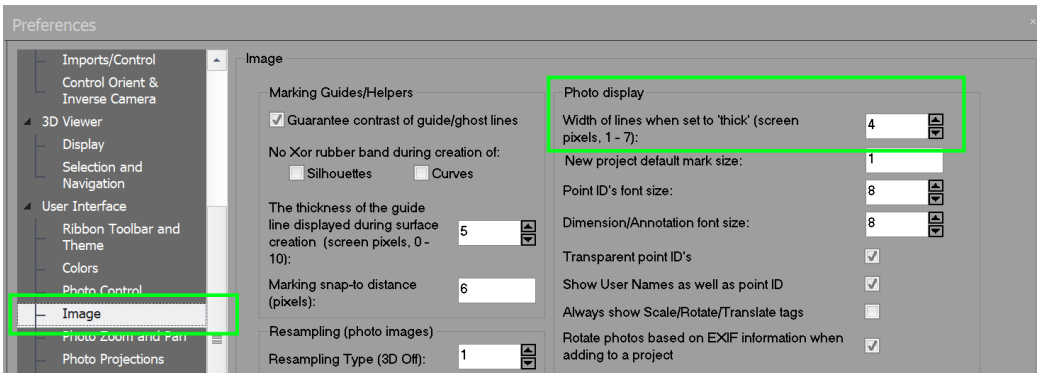

Thick and Visible 3D Projections: when the 3D projections of Lines are displayed on the Photo, it can be helpful to reduce the transparency and increase the 3D line thickness. Both settings are controlled on the 'Visibility on Photos' pane:

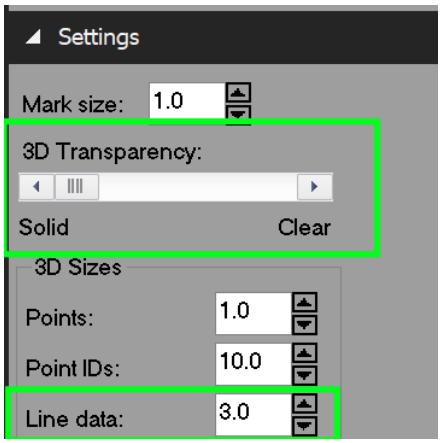

Adding to the Project

After the basic project structure is solved, you may wish to add to the model beyond what has been solved by the constrained Lines and Surfaces. Other videos and documents describe these features in detail, so this section is not step-by-step but instead provides some examples.

Here we have added Surfaces (Path Surfaces and a Best Fit Plane) to the block points.

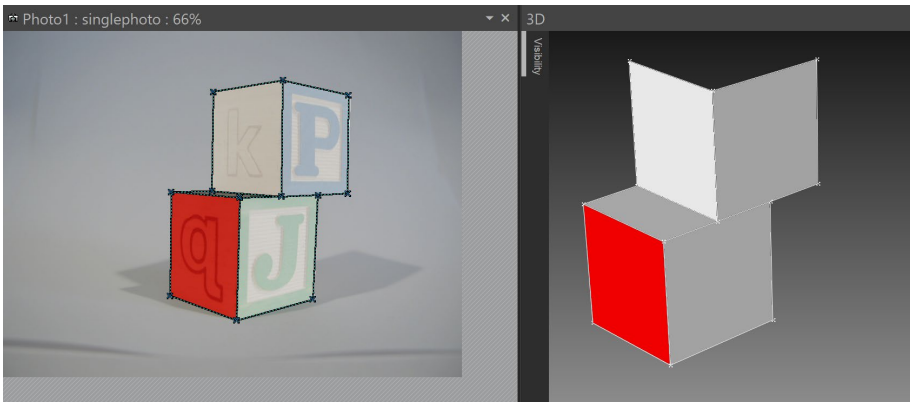

With surfaces you can do several things, such as create photograph textures and ortho-photos, and draw on the surfaces with Surface Draw and Outlines. Several tutorial videos are available on texturing, Outlines, etc. Here are two 3D views, one with extracted textures and one with Outlines.

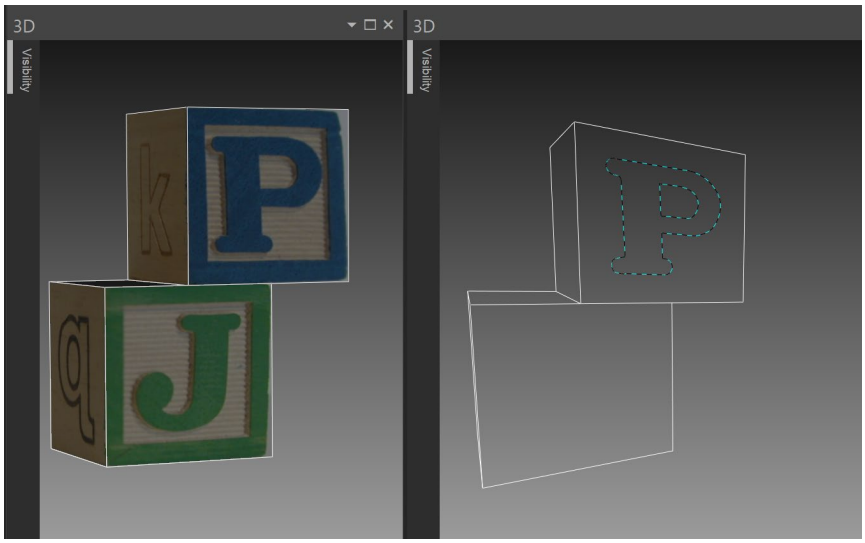

You can also offset a Surface to extract data from another area, provided you know the offset amount. And Surface Offset Points can be useful for extracting heights. See the examples below. Here is a video on Offset Marks.

Solving Camera Parameters with Constraints

Inverse Camera is the process in PhotoModeler that solves a Camera's internal parameters from photographic data. Inverse Camera works with data from one photograph and Constraints or Control Points on that photograph.

When an Inverse Camera process is complete you will see the results of what camera parameters have been solved. For example, for all constraint-based projects at least the focal length will be solved, and the Processing Complete dialog will show something like:

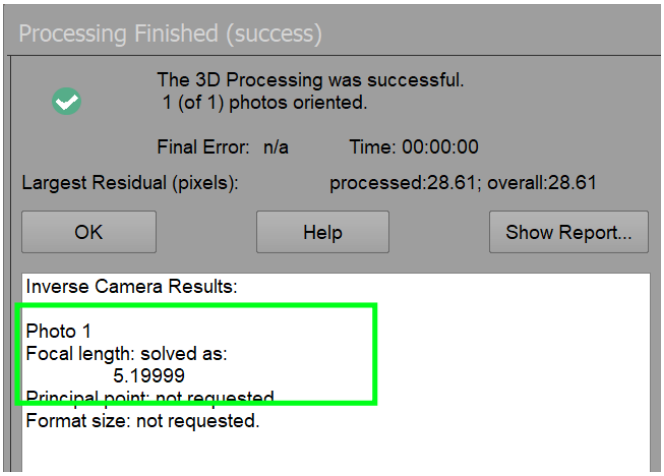

You may also see the statement "requested but not solved". This has two meanings: a) the first is that there was not enough data or geometry to solve the parameter, and b) the parameter did not change from the last processing – this is common in Constraint projects if no editing has been done. In this case it is not an error; it is simply indicating that the parameter did not change in the last processing.

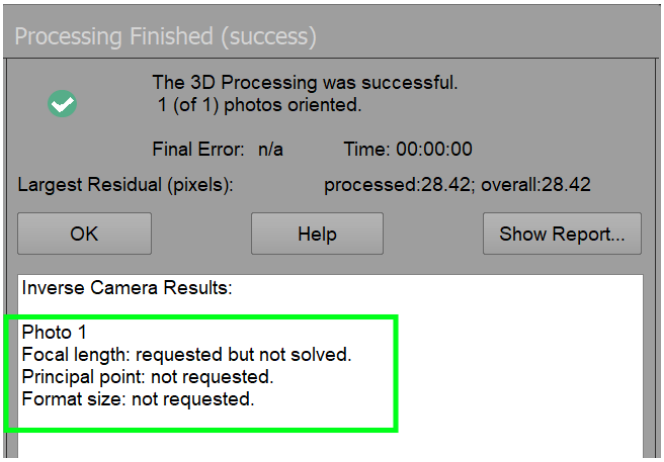

The statement "not requested" means that this parameter was not requested for solution. The number of parameters that can be solved depends on the vanishing point model (see above). A photograph with two-point perspective can be used to solve focal length. With a strong three-point perspective photograph, you may be able to solve for the principal point and format size. Solving extra parameters may be necessary at times to reduce the error and improve the results. You cannot solve any parameters with a one-point perspective photograph, but the project will solve and can give useful results otherwise. With a one-

point perspective photograph solution you cannot depend on the position of the solved camera.

The other parameters that can be solved in a Constraints project are:

- “Principal point” – the location of the optical axis of the lens where it hits the imaging surface (normally near the center but in some cases the photograph may be a print that has been cropped (i.e. offset center)). If it is a non-cropped digital photograph, normally the default centered principal point is sufficient.
- “Format size” – the size of the imaging surface in real world units. This can also be thought of as solving the pixel aspect ratio. An image with square pixels will have a format size in the same ratio as the number of image pixels. But if the camera had non-square pixels or the image is the result of an uneven scan, the image format width and height would not be the same ratio as the width and height of the image in pixels.

To request these extra parameters to be solved, open the photograph's properties (right-click menu on the photograph window or the image chip) and expand the Inverse Camera settings:

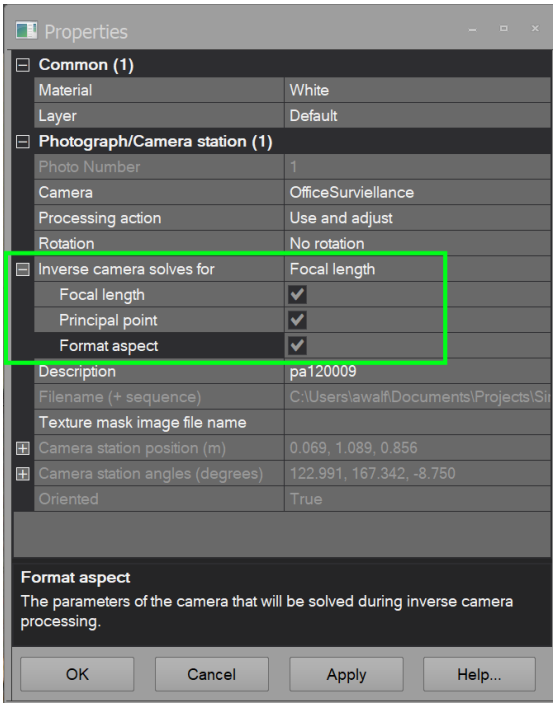

After re-processing, in the Processing Finished dialog we can see the results of requesting these additional parameters:

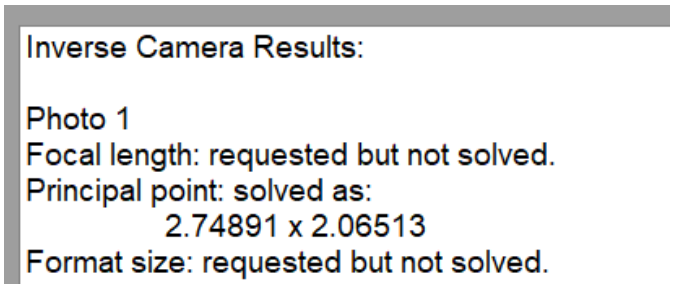

Here, the focal length did not change, the principal point was solved, and there was insufficient data to improve on the default format size.

Single-photo, constraint-based projects do not solve for lens distortion, so these methods will not work on photos with significant distortion. Note that there is the potential of using the single-photo-line-based distortion estimation method of PhotoModeler to correct distortion first on the image and then use the Axes Constraints on the corrected image (contact PhotoModeler).

While having a calibrated camera is unlikely in these reconstruction scenarios, if you do have a calibration for the camera (or like-camera), you can load it, assign it to the photo, and the processing will take it into account. You may wish to turn off all the Inverse Camera parameters solution settings in this case.

Multiple Photo Projects

What is the best method when you have more than one photograph of the object or scene with no camera calibration information, and you wish to use Constraints to solve for the Cameras?

If the photographs have enough overlap for a manual multiple-photograph project but not enough camera information (the cameras might even be different for each photograph), a good approach is to solve for each camera's internal parameters (its calibration) with separate constraint-based projects for each photograph, then load those cameras into a standard multiple-photograph project and use manual marking and referencing. The Manual Project tutorial shows the process with a single camera, but multiple camera definitions can be loaded, even one for each photograph.

You can use Axis Constraints in a multiple photograph project. Each photograph (and any Inverse Camera) is solved first. This is similar to having separate projects for each photograph. Then any Optimization is done on the common marks/references between photographs (with the Axis Constraints being ignored).

Project Examples

Earlier, a pair of toy blocks and an indoor crime scene were used to demonstrate the basic constraint setup steps. A couple of real-world examples are shown below, focusing on results without step-by-step details.

Ambulance Example

Vehicles that have a box (ambulances, trucks, etc.) can sometimes be modeled or partially modeled with Constraints. And with some assumptions, other parts not on the box can be modeled as well. Below we added the X, Y, and Z Axis Constraints (the Z Axis Constraint is visible in green), and then added two Surfaces so Outlines could be used for the rear doors and windows:

The left 3D View shows the two Surfaces and some Outlines.
The right 3D View shows the photo-texture of the side with no perspective distortion.

As noted, sometimes you may be willing to make assumptions to add to the model. In this example, the two wheels have been outlined with circles on the plane of the side of the truck box. The wheels may not be in the plane of the box side, so this is an approximation.

Building Example

Buildings tend to have strong lines that work well for vanishing point analysis and PhotoModeler Axis Constraints. However, you may be limited to what information you can obtain from the building. If the base of the building is visible and straight (that is, not on slanted-earth or sidewalk), you may be able to derive information from the bottom plane, as in this example.

Here the goal was to find the height of the traffic light. We can obtain an approximation of the height of the traffic light above the plane of the bottom of the building. If the sidewalk is not too sloped, then the measurement will be close for height above the sidewalk.

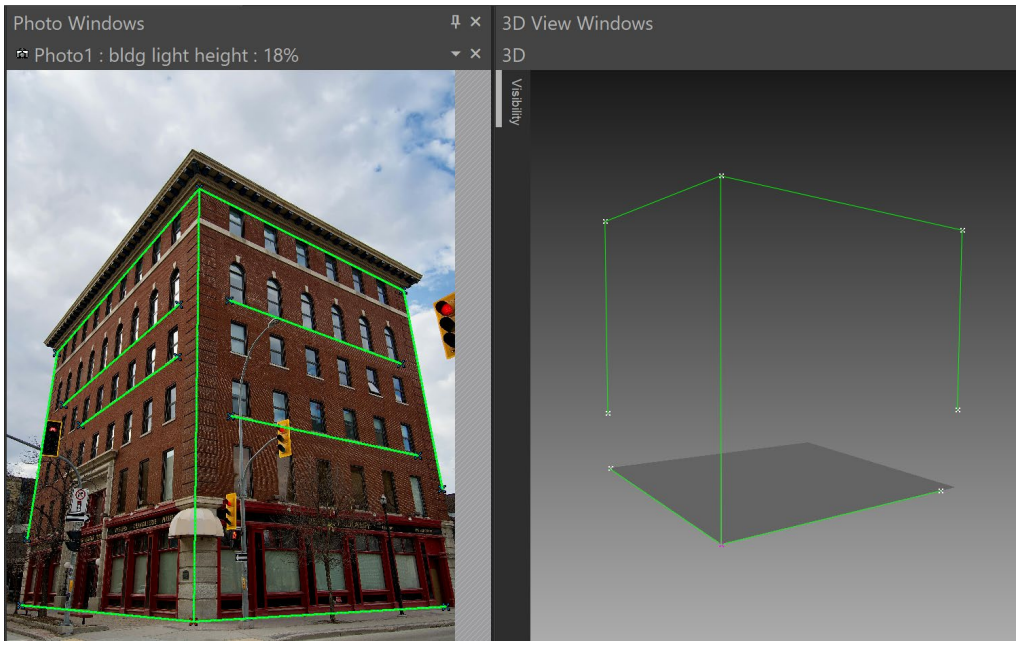

All the constrained lines are highlighted in green and the XY ground plane was created.

We know the width of the building, which was used to set the scale, and it is solved on the XY plane. Then a Surface Draw point was placed at the base of the light standard (with assumption that the bottom is roughly in the plane of the bottom of the building). An offset in the Z direction (height) was then added and dragged to the height of the clamp for the traffic light. Giving a height of approximately 12 feet. The offset point method used here can be seen in [this video on Offset Marks].

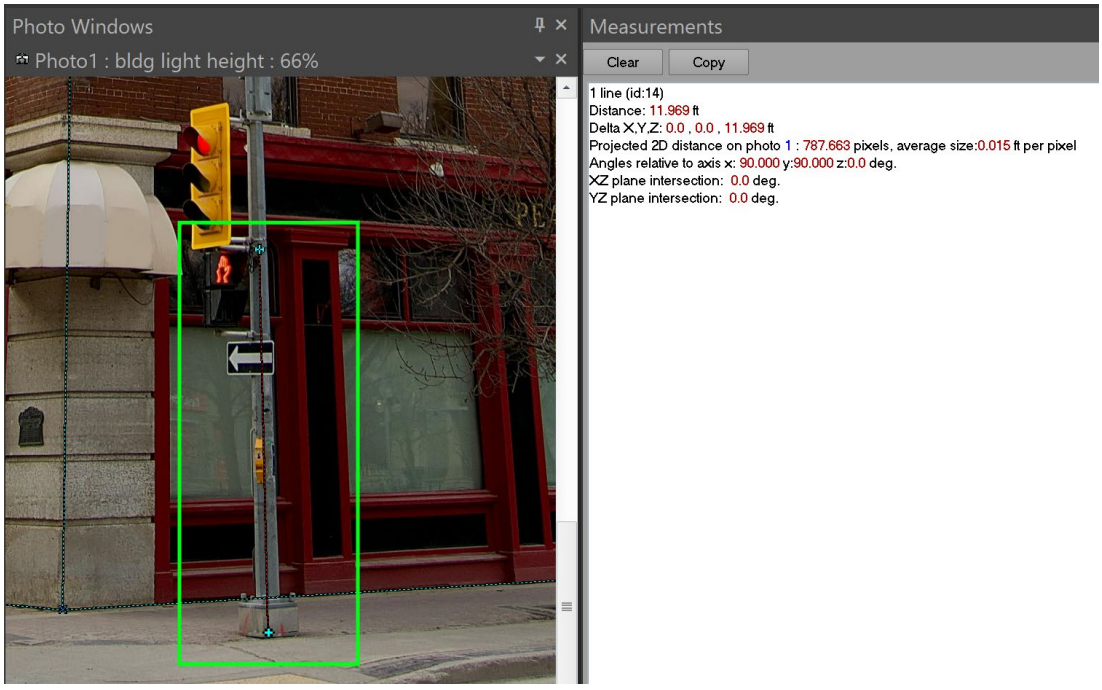

The green box outlines the Offset Mark. The blue point is a Surface Draw point, drawn on the plane at the bottom of the building. The Measurements pane on the right shows the height of that line.

Summary

- Axes Constraints allow single-photo projects to solve useful 3D data when parallel and perpendicular lines are present.
- Correctly marked and connected Lines are essential; only the largest connected network becomes 3D and measurable.
- Residuals should be inspected, and projects with high residuals should be reviewed for marking or Constraint issues.
- Additional Constraints, Surfaces, and Surface Draw can extend what can be modeled and measured.
- Inverse Camera with Constraints can solve key internal camera parameters, which can then be reused in multi-photo projects.