insight3d – quick tutorial

What can it do?

insight3d lets you create 3D models from photographs. You give it a series of photos of a real scene (e.g., of a building), it automatically matches them and then calculates positions in space from which each photo has been taken (plus camera's optical parameters) along with a 3D pointcloud of the scene. You can then use insight3d's modeling tools to create textured polygonal model.

This tutorial takes you through the individual steps and describes how to achieve the best possible result (or at least get it to work, it's in beta).

Where do I get the app?

insight3d is a free opensource application. Versions for both Linux and Windows are available on project's website:

`http://insight3d.sourceforge.net/`

The package includes the application, this manual and four photos of a small part of Prague's castle (in the folder `example_photoset`) which can be used to test the application.

What photos should I take?

You should take a series of photos that cover the whole scene you want to reconstruct. The following photos would be ideal for reconstruction of the building on them.

Fun fact: These photos are of "Lesser town square" in Prague. Photos taken there were the first ones to be reconstructed using insight3d. Coincidentally, according to Microsoft's website, the first photos fed into PhotoSynth (which is another application that also implements camera calibration), were of Prague's "Old town square".

The photos can be taken from arbitrary positions and you don't need to perform any measurements when making them. The application figures all of that automatically. However, several important rules should be remembered:

- **The photos should be focused** (not blurred).
- There should be **large overlaps between neighbouring photos**. Notice how most of what can be seen on any of the photos above can also be seen on the next one.
- There should be only **limited angle difference** between subsequent photos. When moving around an object, take a photo every 15-25 degrees. Again note the photos above and how with each photo we look at the building from only slightly different angle.
- Shoot scenes with lots of unique **details and textures**. Old buildings and sculptures are a piece of cake to work with in insight3d. Blank walls and clean cars have low chance to be automatically matched. But if automatic matching fails, insight3d offers tools to enter matches manually.
- **Avoid planar scenes.** When everything on the photo lies on a single plane in space, like on the following photo, insight3d can't properly determine the focal length.

- **Do not crop the images.**
- **Rule of 3.** Every part of the scene you want to reconstruct should be visible on at least 3 photos.

If your dataset meets these criteria and you still can't get insight3d to work with it, please send it to us. The software is in beta, there are lots of different cameras and there are parameters in the calibration algorithm that still have to be tweaked. Information about how to send the dataset is at the end of this document.

What result can I expect?

The following figure shows untextured and textured polygonal model that can be quickly (within 10 minutes) created using insight3d. The pyramids represent the positions and orientations of the cameras. The screenshot on the left also shows automatically reconstructed 3D point cloud of the scene.

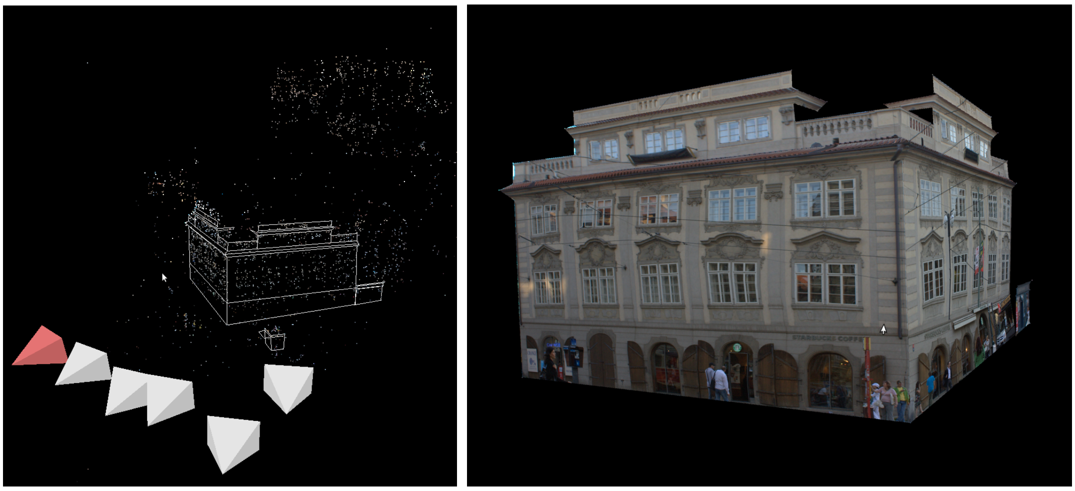

Starting the application and loading photos

Let's create our own 3D model! First start insight3d from start menu. (If you downloaded standalone version of the application, just extract it somewhere on your disk and start `insight3d.exe`.) You should see something like this:

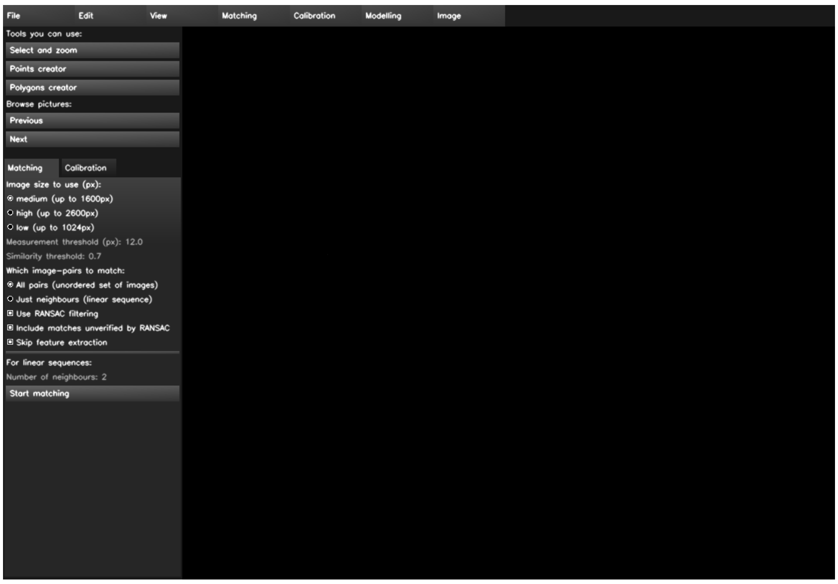

At the top, there's main menu with items like File, Edit, View, etc. On the left are several sidepanels, which we'll describe during the tutorial. Note that a console window also appears. This is important because there you can often see error listings or info about the progress in computations. On Linux, you might want to start the application from terminal (but it's not crucial).

insight3d starts with an empty project. Clicking File > New in the main menu gets us back to this stage. Standard installation comes with four testing photos located in the example_photoset folder in the installation directory (on Windows, usually C:/Program Files/insight3d/example_photoset). We'll use these photos during the tutorial:

As you can see, the building on the photos is very textured and there's only small angle difference between the photos – thanks to that, it should be very easy to automatically match them.

First we have to add these photos to our project.

You can add them one by one using File > Add image in the menu. This opens a dialog that lets you select a jpg file. After you confirm your selection, the photo is added to the project. You can then proceed to add other photos of the scene in the same manner.

There's a faster way to add several photos at once: You can create a text file with the names of the files and then load this file in insight3d. To do this, go to the folder with the images, create a new text file and put the names of the photos (like 100_4741.JPG) each on a separate line. Save this text file (files like this usually have the extension .ifl, but you can use .txt if you want) and open it in insight3d using File > Add list of images in the main menu. This file is already prepared for the example photos, so just use the file list_of_files.ifl in the example_photoset folder.

After adding the photos, one of them should appear in the main working area of the application. Use the Previous and Next in the sidepanel on the left to browse through them.

Automatic matching and camera calibration

Here is the magic. insight3d can automatically determine all parameters of the cameras used to acquire the photos (like their orientations in space or focal lengths) without knowing anything about them.

This is done in two steps. First the photos are *matched* – the application looks at each image, finds important points in it and then tries to track these points on other photos.

The photos are then *calibrated*, which means that the positions of the camera (and its other parameters) when taking each picture are calculated from these tracked points.

If everything goes well, all this can be done with two clicks. To match the photos, click `Matching > Start matching` in main menu. The matching process will take some time. For the example photos, it should take about a minute. During this the application might become unresponsive, please check the progressbar or the console window to see if the computation is still going on. Once it's finished, you should see detected points marked on the current image:

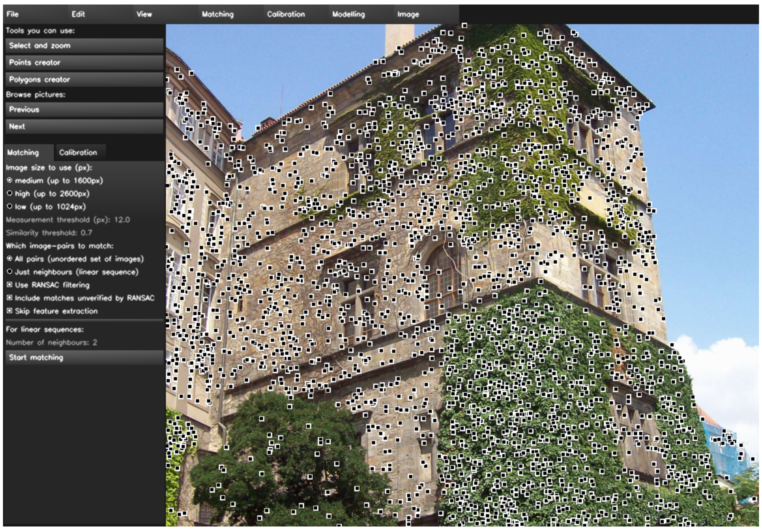

As you can see, there's a lot of square shaped markers, each representing a point that has been sucessfully tracked on other photos. If you move your cursor over any point (and keep it still for a second), insight3d will show thumbnails of all the corresponding points. If these thumbnails don't appear, it's probably because currently selected tool is not the "Select and zoom" tool. Click on the "Select and zoom" button in the sidepanel on the left to activate it. The thumbnails look like this:

We can see that this point has been tracked correctly – all the points correspond to the same physical point (vertex) in space. Sometimes automatic matching produces mismatches, but thankfully the camera calibration algorithm is robust and such outliers shouldn't affect the result.

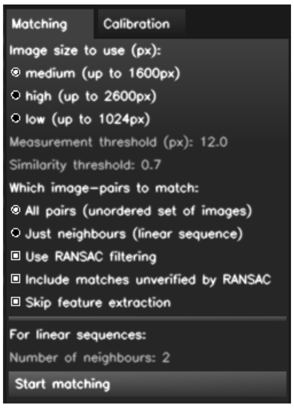

It might happen that only small number of tracks (potentially not covering the whole photo) is detected and it's impossible to correctly calibrate the cameras (in the next step). In that case, you might want to check if your photos satisfy the recommendations listed in the beginning of this tutorial. It might also help to adjust the matching settings. This is done in the matching tab in the sidepanel (see the image on the left).

You might want to change the image resolution used to extract the features (bigger resolution means more points to track). If you have a lot of photos (9 or more), it makes sense to match them *as a sequence*. By default, `insight3d` matches all photos against each other. If you select the "Just neighbours" radio button, it will match only the pairs of photos that are near each other (in the order in which they have been loaded). This option is used very often. The other options are not as usefull and we'll ignore them for now.

Now comes the second click. To calibrate the cameras, use `Calibration > Automatic calibration` in the menu. Once again, the application should perform some computations for a minute or so. When the main window becomes responsive once again, the cameras should be calibrated.

Immediately, you won't see any difference. To check out the result, switch into "Overview mode", where you can see the scene in 3D. To do this, go once again to main menu and use the `Edit > Mode... > Overview mode` item. This shows a rotating 3D overview of the current scene, like on the next image:

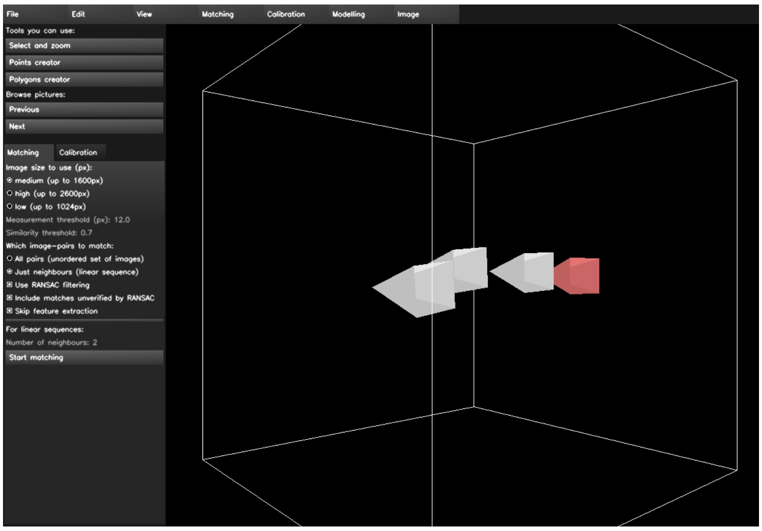

There should be four pyramids, each representing one camera in our scene. The camera belonging to currently selected photo is highlighted.

3D point cloud

The cameras are now calibrated. That is: their positions, orientations and internal parameters are known. But are they correct? It's quite easy to check this. With another click of a button, insight3d can calculate the positions of all points tracked on our photos. The resulting point cloud should resemble our scene. Use the `Modelling > Triangulate all vertices` menu item to do this (it's the second one in the `Modelling` submenu, not the first one). After a few seconds, there should be 3D point cloud of the scene rotating on your screen (note that you have to be in Overview mode):

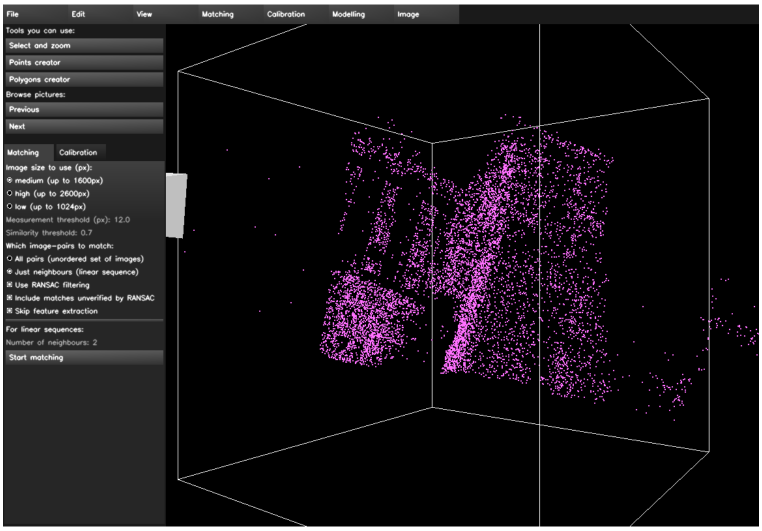

Sometimes, the point cloud seems deformed or just looks like random noise. That might be because there's a lot of mismatched tracks. The cameras still might be correctly calibrated, though. Use the `Modelling > Triangulate, only trusted` command to triangulate only the vertices that have a very high probability of being correctly tracked.

If you want, you can use the `Image > Colorize vertices` menu item and insight3d will produce a colorized pointcloud. Thus, if all goes well, you can have quite precise 3D pointcloud of your scene with only minimal effort.

Modeling

Polygonal model is created by defining vertices and then joining them into polygons. To create new vertex, it has to be marked on at least two photos, so start by going back to Shot mode by clicking `Edit > Mode... > Shot mode` in menu. We don't want to be distracted by the vast number of automatically detected points, so we'll hide them using menu command `View > Show/hide automatic points`.

To mark the position of a scene vertex, we'll need to use the "Points creator" tool. Activate it by clicking the "Points creator" button in the sidepanel. Now, mark the corners of the building on the first photo by clicking on them using left mouse button. This will create a point for each of them. When moving the cursor over the photo with the "Points creator" tool activated, you should see a detailed view of the spot underneath the cursor. This should help you mark the point more precisely. Use dragging by middle mouse button to scroll the photo, zoom in and out using the mouse wheel (the detailed view dissapears if you zoom in so close that it loses its purpose). The middle button/mouse wheel works the same no matter which tool is currently used.

The first photo with points marking the positions of the building's corners could look like this:

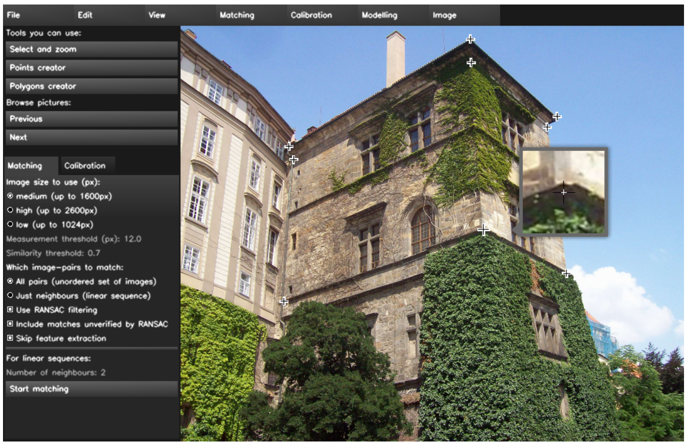

We have to mark these vertices on at least one other photo to be able to triangulate them. Move to a different photo using the "Next" and "Previous" buttons on the left. With the "Points creator" tool still activated, move the cursor over the new photo. You'll notice something different. Besides a detailed preview of the spot under the cursor, there's also a thumbnail of one of the vertices we've just marked on the previous image. This is the vertex we should now mark on current image. So, click on the place in the new photo where this corner is visible:

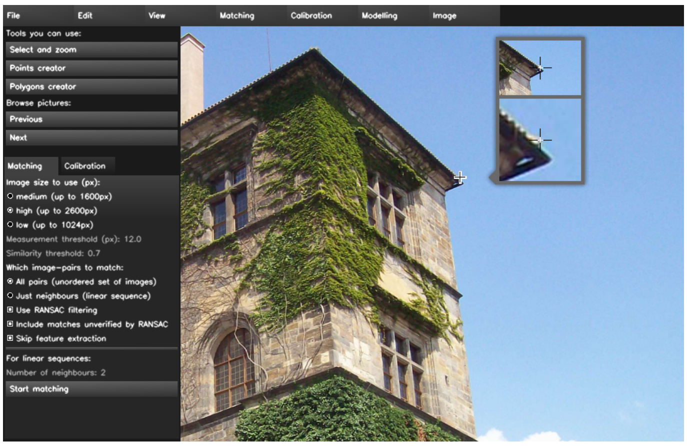

After you do this, insight3d will automatically move to the next vertex that is still unmarked on current photo. In this way, mark all the building corners for the second time. Once you've marked all the vertices, insight3d lets you create new ones like we did on the first photo. It often happens that not all vertices (such as all building corners) are visible on all photos due to occlusions. In such cases, you can skip to the next unmarked vertex by pressing PageDown key on your keyboard. The PageUp key gets you to the previous vertex, the Home key to the first unmarked vertex and the End key lets you to mark new vertices (e.g., corners that were not yet marked on any photo). You can edit the position of an existing point by moving the cursor over it and dragging it to the new position.

When the vertices are marked on at least two calibrated photos, we can triangulate their positions. Use the menu command `Modelling > Triangulate user vertices` to do this (you can use also the command `Triangulate all vertices`, but that computes also the automatically tracked vertices – since there usually is a lot of them, it might take some time and so it's quicker to triangulate only the user-defined ones). Correctly triangulated vertices should have a green dot appear close to them, like on the following image:

This green dot represents the *reprojection* of the vertex – that is the position on which the triangulated vertex would be visible using the calibrated camera. Ideally, vertices would be visible precisely on the position of their marked points, but inaccuracies prevent such perfect fit. To achieve maximum precision, reprojection error should be as small as possible. That is, the green dot should be really close to the center of the cross representing the point. (Reprojection is displayed even for automatically generated vertices and thus can be also used to verify the correctness of camera calibration.)

If a point is marked incorrectly, insight3d will probably discard it. For example, if a vertex is marked on four photos (using four point markers) and on one of them was positioned incorrectly, insight3d will detect this and use only the correct ones. Note that the vertex in bottom right in the figure above was not triangulated (it doesn't have a green dot connected to it). This happens when all or most of the points belonging to a vertex are marked impreciselly. In these cases, insight3d will leave the position of the vertex undetermined and it's upon the user to correct the position of some of the points.

To achieve good results, it's usually enough to precisely mark the position on three photos. If you want, you can now go ahead and mark the building's corners on another photo. Now, it's even easier to mark the vertices, since insight3d now knows their approximate position and once you press the left mouse button down, it'll immediately show the reprojection of the vertex being marked. Thus, it's harder to make mistakes like marking the wrong corner of the building. Use the command `Modelling > Triangulate user vertices` once more to recompute the vertices to reflect their newly marked points.

Sometimes you may need to delete a point that was created mistakenly (perhaps it belongs to a vertex that is not even visible on current image). To do this just select the vertex by clicking on it while the "Select and zoom" tool is active and then use `Edit > Erase selected points`. You can select more than one point using rectangle select by pressing down the right mouse button and dragging to define rectangular region. Rectangle select works with all tools, not only with "Select and zoom". (Also, it can be combined with the Shift and Ctrl keys.) If you want to delete the whole vertex, you should delete all its points. This can be quickly done by selecting one of them and then using the `Edit > Select... > Add corresponding points` menu item to add all corresponding points to the selection. A single delete of all selected points then erases all points of the vertex (or vertices).

You can now take a quick look at the scene in the Overview mode (as you might remember, it's triggered by `Edit > Mode... > Overview` in the main menu). There should be our vertices (the corners of the building) rotating inside a 3D cube.

Let's connect them into polygons. Switch back to Shot mode (in main menu, use `Edit > Mode... > Shot mode`) and activate the "Polygons creator" tool by clicking on its button in the left sidepanel. We want to create polygons for both of the building's walls seen on the photos. To do this, click with the left mouse button on the four points forming one of the walls. insight3d will start to draw violet polygon connecting the points as you're adding them to the new polygon:

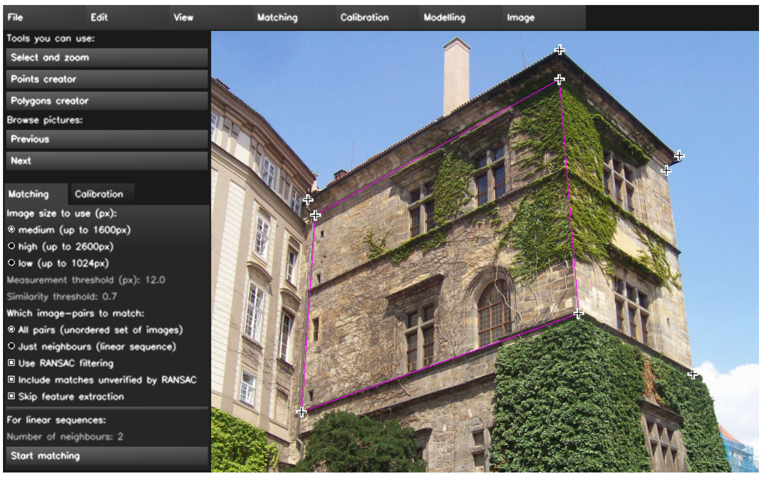

After you've added all four points, press Enter to finish the polygon. It will turn its color from violet to white and you can start creating the next one.

The polygon highlighted by the violet color is the one that is currently being edited. You can select it using the PageDown and PageUp keys on keyboard. In this way you can add a point in an already finished polygon. To remove a point from a polygon, click on the point while pressing Shift on keyboard. To delete a polygon (but not its vertices), use `Edit > Erase current polygon`.

The photo with all polygons defined could look something like this:

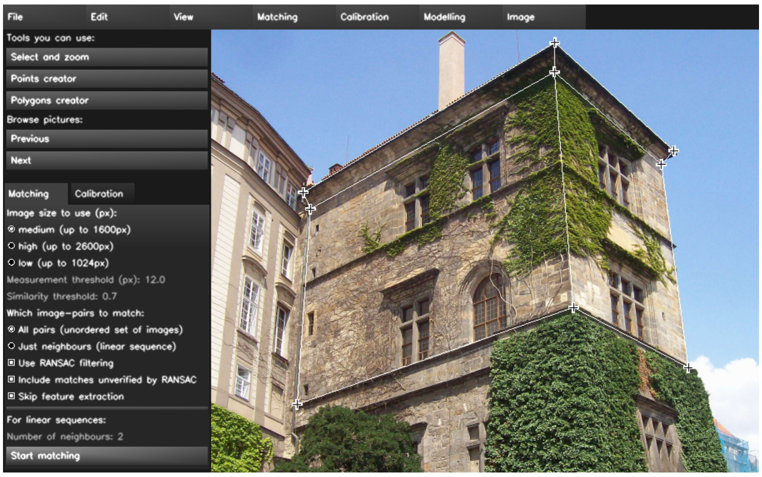

You don't have to mark the polygon on every photo where its vertices are visible. It suffices to do it once. Note that the vertices of a polygon are connected by lines even when some of them are not visible on the current photo. This might be a little confusing sometimes, if you have an idea how to better visualize this (you'll probably run into such situation at some point), write it on our wiki (it's on the project's website).

The resulting textured 3D model

Let's take a look at what we've created. Switch into Overview mode (`Edit > Mode... > Overview`) and you should see a rotating cube with the polygonal model inside:

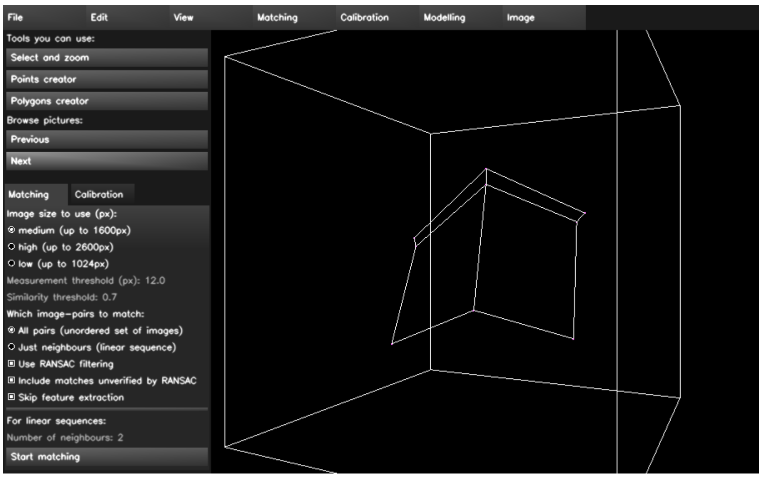

If some of the polygons are missing, it's probably because one or more vertices couldn't be correctly triangulated because they are marked inaccurately. Fix this by going back to Shot mode and correcting the points (since their position couldn't be trianguated, they do not have a green dot near them).

Textures for the polygons can be easily generated by clicking `Image > Generate textures` in main menu:

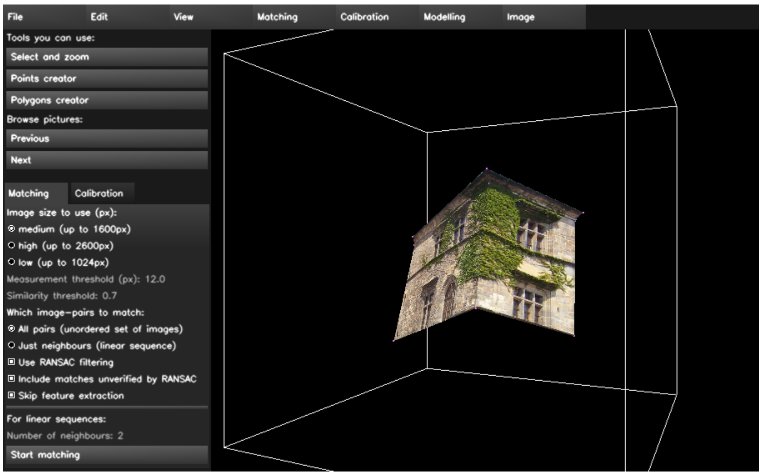

You might want to see how your model looks along with the 3D point cloud from automatically tracked points that we've hidden when starting modeling. Clicking again on `View > Show/hide automatic points` will bring it back. You might also want to see the model in Inspection mode, which is triggered by clicking `Edit > Mode... > Inspection mode` in menu. In this mode, you can walk around the scene using the keyboard's cursor keys.

Saving and exporting

Save your project often. insight3d never asks for confirmation of user's actions (this will be fixed). Saving the project can be done using the `File > Save project` menu item. This will open a dialog that lets you select the file name for the project. insight3d's files usually have the extension `.i3d`. Saved project can be opened using `File > Open project`.

The final 3D model can be exported into several file formats. The most usefull is probably VRML, which you can use to import the constructed 3D model into other applications – for example, into conventional 3D modeling software like Blender. If the pointcloud is visible when exporting into VRML, the pointcloud is exported along with the polygonal model. If it's hidden, it's not exported.

Final thoughts

insight3d has a lot of features that haven't been described in this tutorial (a lot of these undocumented features are still under development, though). If you have suggestions on how to improve insight3d or this tutorial, feel free to send an email or write in our wiki on project's website: `http://insight3d.sourceforge.net/`

If you happen to use insight3d for something nice, please send me an email. It'll motivate further developement of this application – and there is a ton of stuff that could make insight3d better and much more fun to use.

Since the application is in beta, it would be great if you could send in photos taken by your digital camera intended for reconstruction in insight3d. There are many, many different cameras with different internal parameters and the calibration algorithm still has to be tweaked a little. You should always send the whole dataset (all the photos you tried to use in insight3d). The photos should satisfy the recommendations from the beginning of this tutorial. They shouldn't be altered in any way (for example resized in Photoshop). Both the photos that were successfully reconstructed in insight3d and the ones with which it had difficulties are valuable. You can send them by email (if they aren't that large, gmail has 25MB limit) or using `rapidshare.com` or similar uploading service.

Thanks for using insight3d.