

## Canoma Frequently Asked Questions

v1.0

August 1, 1999

Robert Seidl

**Q** What format should I use for importing into 3D Studio Max ?

**A** VRML2, 3D Studio reads the geometry, textures and creates a separate camera object for each source image in Canoma.

**Q** I resized some of my source images in a project \*after\* I'd started the project. (The book says you can do this if you don't change the aspect ratio which I didn't) It failed. The images, now 50% their original size, didn't fit anymore, displaced from the upper left corner. Whats going on ?

**A** Dont\* do that ! Canoma doesn't like resizing of **source images**. At all. The comment in the manual refers to **edited textures** that you pass to a pixel editor like Photoshop with the Canoma "brush" tool. For example, you may get a relatively lores texture on some surface, but you want to put some text on there (uhm, bigger than a 3 pt font :-). So you do an Image Resize in Photoshop, making sure you do not change proportions. Then you can add type in -say- 20 point fonts. Save it and the texture comes back, in higher resolution. Source images cannot have their dimensions changed at all once you got started. Very sorry. The reason is that our 3dv files currently save pin & bead coordinates in pixel coordinates, and obviously those are all wrong once you change your source image dimensions. One solution we are considering for the future is to write out floating point 0..1 coordinates. That could help solve the problem. For now, unfortunately you have to stick with your original images.

**Q** Can Canoma antialias the Quicktime animations it creates ?

**A** Yes. When you create an animation, make your window big. Then as you File>Render Animation, type in a smaller number (eg your original window size was 762x524 or whatever, then type in 200 for the width. The height will automatically adjust, to keep your aspect ratio. Now, before each frame is handed off to Quicktime for compression, we actually render it at highres, then do a filtered downsample (i.e. antialias) to your selected output resolution. That way you can get antialiased movie output from Canoma. Of course, taking the Canoma models into any high quality renderer of ours like Ray Dream Studio or Bryce will give you all kinds of additional rendering options...

**Q** What is the memory consumption of Canoma images – I run out of memory even when I use very high JPG compression ?

**A** Whatever the compressed images are on disk (JPG, TIFF, BMP etc.) they are in memory as 24 bit RGBA images. So the math is: height (in pixels) \* width (in pixels) \* 4 bytes = total # bytes for in-memory picture. In addition, there is a fairly substantial amount of working RAM necessary to generate high quality textures. On the Macintosh, you may want to give Canoma more memory by selecting the application icon and changing the value in the Get Info dialog.

**Q** I only have a single picture for my project because it is a historical photograph and it is not even from a very good point of view. Because of that, I cannot really specify depths of objects properly. Any tips ?

**A** You can cheat by creating a second, artificial view from the side in which you can – arbitrarily – set the depths of objects with beads on their back edges. Here is what you do: In Photoshop, create a 200x100 image, fill with medium gray. Then paint the alpha channel so that all the pixels are masked (except a single one, perhaps at the upper left corner - this is to avoid a known problem when all pixels are masked - sorry) Then you can add this image to your project. Since all (well almost all) pixels are masked, Canoma will not use that picture for texturing. But you can "model" in it by adding pins or beads. I would suggest adding only 3-4 pins to fix rotation, and scale of your model and for the rest use beads, for example to determine the depth of objects that was indeterminate from the (single?) other picture. If you use too many pins, you may "overconstrain" the model with wrong information, since you don't really see the scene in that fake picture, you won't know exactly where to place the pins. Also, don't forget to use the trackball for rough orientation before applying the first pins, so the model is already in approximately the right place.

**Q** What is the "Polish Solution" option and when it should be used ?

**A** Ok. Fasten your seat belt. Canoma's main task is to find & maintain reasonable values for so-called free parameters, like position, orientation, and dimensions of various primitives as well as all the camera information (focal length, position etc.) There can be quite a few of those (even hundreds) It has to find these values subject to the constraints that the user has provided (pins, beads, glue, align, stack, concentric etc.) In principle, Canoma tries to find a globally correct solution (i.e. taking all the constraints into account) When this gets too big a problem to solve, Canoma tries to select a subset to solve. This way it can remain interactive while dragging. However, it may not change all the parameters. Polish Solution actually does a full evaluation of the whole problem. It can take longer but will take all constraints and parameters into account. Usually there won't be much difference between before & after polishing.

While you polish, there will be current error value displayed at the bottom of the screen, that will usually become smaller, indicating Polish is actually succeeding in finding a better overall set of parameters. If that number doesn't change any more, you might as well ESC (escape key) out of polishing – the changes will be negligible.

Another hint: turn on the "stress display", which is the bottom of the three gray/purple buttons on the top right hand side of the editing window. Canoma will then display faint gray wireframes where IT thinks primitives are. If you modeled everything correctly, you wont see a discrepancy between your yellow/red wireframes and these gray ones. But if you "stress out" canoma by giving wrong input, you will see the gray frames "tearing" away from the usual yellow/red wireframes. You can try this even with a single box: pin 4 corners then noodle around with the 5th to see it stress.

**Q** What is a translation sweep ?

**A** A translation sweep or extrusion is a planar contour (in Canomas case a polygonal one) that is extruded straight back in a perpendicular direction from the contour plane. Its handy for modeling logos, irregular roof shapes, people "cutouts", cars etc. The translation sweep primitive is to the left of the polygon primitives in the Canoma primis. They start life as a triangular prism, but with the pen tool (or accelerator SHIFT key) you can subdivide its contour edges to create more contour points (even approximate a curve). You can also delete existing contour points by clicking on them. (Polygons and Curtain work the same way with SHIFT/pen tool)

**Q** I'd like to see a second use for Alpha channels: could they be used to make certain areas of an image transparent such as lattice work, small railings, phone wires, etc. I have a scene where a building has a sign on it's roof that is just letters held up with a frame of wrought iron struts. Too complicated to model all the pieces of course, but a single card plane mapped with the image and an alpha to cut out the spaces around the letters would give the illusion of a complicated piece of detail with little effort except making the Alpha channel. Or is this best done in another program after the model has been exported? Maybe so, because I know Bryce could do this in a snap.

**A** This is an excellent idea. Yes, it should be done as a postprocess on the final textures that Canoma emits to OBJ or VRML file formats. It is a very effective trick to make foliage seem semi-transparent (just massage a noisy alpha channel a little and you can look through portions of a tree canopy) The export format for Canoma textures is JPG, which is nicely compressed on disk, but does not allow for an alpha channel. So a little bit of extra work is required: first you edit the appropriate texture (you will probably want an image browsing application to find it) to add the alpha channel and save it in a format capable of storing the alpha channel. For VRML use, the recommended format is PNG and for use in Bryce, TIFF is recommended. Then you have to change the reference to the texture from the old jpg file to your new PNG/TIFF file. Since both VRML and OBJ are simple text files this is easy. Find all occurrences of the old filename and replace them with the new filename (A texture might be used more than once, for example when texture stealing is on or when it is used for both the ambient and diffuse channels in an OBJ file)

**Q** When should I use pins, glue or beads ?

**A** Use pins to fix corners of the wireframe, beads to pull edges to the right place and glue to attach two objects together in some fashion. Always start modeling with pins, then perhaps beads and last glue if needed.

**Q** I gather that it's advisable to have a minimum of pins, that is don't over pin things. Am I right? Does each pin give Canoma another (set of) number(s) it has to crunch? Sometimes I can move an existing pin just a bit and it moves an unpinned corner to the right place. Is this advisable? i.e. Don't put pins on corners that are already lining up okay?

**A** Exactly right, fewer pins mean less work. In the very beginning of a project however you really need to provide some pins to tell Canoma about the perspective in a scene. But in general you only need to pin things that are "off" their correct positions on the picture.

**Q** Is using Glue a good idea whenever possible? Does Glue help Canoma in some way do less work?

**A** Glue allows you to specify a spatial relationship between two objects without having to use pins. For example, you could glue the corners of two objects together with a single glue instead of two separate pins. If you know two surfaces are aligned, glueing their edges together or at least a corner to an edge is a good way of telling Canoma about this geometric relationship.

Glue can be very handy to tell Canoma relative position, like "this buildings front is in line with the other buildings front". If you do not see bottoms of buildings, glue can be useful to fix their relative position (since the y-coordinate of the bottom corners will determine depth away from viewer, getting y wrong a few pixels in a shallow perspective could slide them hundreds of meters forward or backwards. Glue allows you to say "the bottom front of this building is the same distance as the front or back of this other building" without actually specifying a pin which would likely be imprecise.

The other common use for glue is fixing stacked objects to their base in the right place: roofs, columns etc.

One word of caution: dont use too much glue too early in a project. Canoma will shrink the whole scene to make the glue smaller (its aim) but that will make the model less controllable. Not a hard and fast rule, but you should probably have 2 or 3 well pinned objects before adding glue.

**Q** Is holding down the Shift key (or checking the Align icon) when creating an object a good idea? Does the constraints it places on that object assist it in making less work for Canoma?

**A** Absolutely. You should use stack/align/concentric whenever you can. Why? Because it reduces the number of free parameters Canoma has to figure out values for. For example, when you stack one primitive on top of another, Canoma doesn't have to figure out the height at which to place the second primitive - it just uses the top of the first one.

Aligned was perhaps an unfortunate choice of word: we really mean "having the same orientation or rotation" rather than "colinear". Perhaps a linguist out there can help out with a better term for v2?

Especially architectural scenes and cityscapes tend to be extremely parallel, so use align (==shift key creation or click on the align icon before creating a primitive) a lot. Not only will your model look better - because primitives will have exactly the same orientation, you will get there faster because you need fewer pins to make the wireframes match and Canoma will compute more quickly.

Concentric is usually used together with stack, for example when making a tower or column with multiple stacked primitives that all share the same central axis.

**Q** I just used a vertical polygon to create an object -- In View Mode I found it BEHIND the curtain that defines the background of the scene. Why?

**A** When Canoma creates new primitives, it \*try\* to do it in a reasonable location, somewhere in the middle of the screen. If you are zoomed in, Canoma will try to create a new primitive somewhere in your current view. When you are creating a new primitive stacked on top of an existing one, it will try to place it right over the base primitive.

Canoma does not try to place new primitives "in front" of any others. So what may have happened is it just chose a default location midscreen, which may be behind an already placed primitive. Pulling down on the bottom two corners should bring the polygon closer to you.

There is one extra tip: the horizon line will usually be somewhere in your edit view (you won't see it - perhaps we should show it - but you can imagine it there) unless you have a pretty top down view. Now as you drag an object up with a bottom pin, it will be forced to move backwards, away from you (since the pin must be on the ground plane, and that is receding into the horizon). In fact, it will get smaller quite fast as you drag it closer to the horizon. Once you are dragging it higher than the horizon line, you are asking the program to make the object lift off the ground plane or (if that is still an option) make the camera look further downwards so it forces the horizon line further upwards.

Anyway, this is getting a little too intricate, but what I mean to say is: try not to push bottom pins up beyond the imaginary horizon line (there should not be any real world objects there anyway - unless they are flying) And don't forget in case you or Canoma get really confused, there is always the undo key (10 steps of Undo)

**Q** What formats should I use to export my Canoma models to other 3D applications ?

**A** For Bryce and Ray Dream Studio, use Wavefront OBJ file format (you may want to change the amount of ambient and diffuse lighting the texture maps on the objects produce - see Release notes dealing with " Better OBJ and SCN output control ").

For Poser3, use Poser 3 Prop format and follow the instructions in the manual. For Poser 4, things have become much simpler: export as Wavefront OBJ and just import directly into Poser. Canoma is a great room/architecture, furniture and prop generator for Poser users !

For many other 3D programs, OBJ import works fine. Canoma outputs a .mtl (material) file and all necessary textures as JPG files. Be aware that some older applications ignore the material file and hence do not show the objects textured.

For 3D Studio Max, use the VRML2 output format - it loads just nicely into Max and retains geometry, texturing and a camera object corresponding to each source image. See also Canoma 1.0.1 Release Notes on "Better VRML output"

Some architectural modeling programs only allow importing DXF files. This is not as good as the OBJ format as it does not support textures, but sometimes the only way (and sufficient for massing studies) Since some programs rescale DXF files on import, you may have to adjust the size of your imported Canoma model.

**Q** Where are the curved shapes in Canoma ?

**A** We know about the need for curved shapes. The math is an order of magnitude harder and that is why they were not included in v1 - but we are working on it. Meanwhile, you can try this easter egg: Stairs N=98 and 99. Often, you can also approximate rounded shapes with translation sweeps or polygons or (by building some "scaffolding" with rects & line primitives) arbitrary polyhedra. For example, the guitar on the Canoma CD looks very round, it was modeled using a translation sweep for its body.

**Q** Why are all the textures lost when I load a project even though I previously generated them ?

**A** Canoma forgets the textures (on purpose). They are easy to regenerate however, just hit the turtle or rabbit, generally this is acceptably fast. The reason Canoma wants you to regenerate rather than save textures is that as soon as you edit your scene, any applied textures will be wrong (stretched etc.) Try it: do tutorial 1, texture the box & ground plane. Looks fine. Go back into edit mode, and make the ground plane bigger in one dimension (or smaller) Go back into view mode. Squashissimo ! Now retexture. All is well again.

So keeping the textures really only makes sense if you promise not to modify the model anymore (for example you only use it to demos)  
So you promise ?

Ok, here is a trick: hold down the ctrl key (cmd on the mac) while hitting the turtle. In addition to creating the textures, Canoma will write them into a Cache folder on disk (same dir as 3dv project). Now quit Canoma. Restart, load the 3dv file and hit ctrl/cmd-turtle again. Canoma checks for existence of the cache. If its there, it sucks all the textures in from there immediately, no computations. This is also a nice feature if you want to batch process all textures with a Photoshop filter, like a dry brush or one of Kais hairy things. Just make an action and batch apply it to the cache folder. One word of caution: if you change your scene after all (you \*did\* promise...) then you have to manually delete the cache folder.

The other time textures are memorized is when you edit them with the Canoma brush tool & Photoshop or Painter etc. Edited textures appear immediately on load.

**Q** Is there a connection between Canoma and all the Renaissance Art on the box and samples ?

**A** Yes. Canoma uses Renaissance imagery from Italy (and the Vermeer example from Holland) because that is when the concept of perspective was first introduced to painting and science (by Alberti & Bruneschelli, and later da Vinci and Albrecht Durer (who went to visit da Vinci). The cover art for the Canoma box is a composite made from a Canoma model of the Campo (tower) in Siena which is at the center of the Tuscany region in Italy and which flourished during the renaissance together with its bigger rival, Florence. The textures of the Campo model were somewhat massaged in Photoshop then brought into 3D Studio Max where the composit with the room was made. The view through the window shows the town square of Siena.

Over 300 years later, Canoma reverses the process of bringing physical objects onto a canvas using perspective. With Canoma you start with an image and let the program & its math reverse the process, extracting 3D shapes back out of the flat image.

For example, we have used "The Music Lesson", Johannes Vermeer ca. 1662-65

He painted in his studio in Delft, Holland. Many of his paintings were done in the same room, so you could find soome similarities between them. It is also believed he used some form of Camera Obscura (pinhole camera) for his precise perspective work.

Check out <http://www.grand-illusions.com/vermeer1.htm> for some background on Vermeer, the Camera Obscura and a reconstruction of his studio (pre Canoma, a physical reconstruction)

or Illusions of Renaissance Art: <http://tgd.advanced.org/3257/illusion.html>

**Q** Is there a Canoma demo available ?

**A** Yes. Its free. Tell all your friends ! It is available from:

<http://www.metacreations.com/products/canoma> Mac and Win version are available. This demo version does not time out. It comes with several tutorials (sample files + step-by-step documentation in pdf) which guide you through the process of image assisted modeling that Canoma uses. What makes it a demo version is that it only loads the tutorial files we supply with it. All features of the full version are available, including export to other file formats, texture editing with Photoshop or Painter, Animations etc.

**Q** Canoma models in Poser - lighting suggestions.

**A** Canoma models are "prelit", that is the photographs from which they derive already show the scene lit, so quite often just ambient light will do. We suggest using white lights only for Canoma scenes and turning shadow casting for the scene itself off (Select the Canoma prop, double click to bring up the properties dialog for the prop, and turn off shadow casting).

**Q** Canoma models in Bryce - importing tip.

**A** When you import Canoma models into Bryce as OBJ files, everything is fine the first time. But say you changed your model in Canoma, export it as OBJ again, switch back to Bryce, delete the old Canoma object in Bryce, and import again, your textures may be wrong. This is because Bryce has a texture caching feature which in this case gets it wrong: it thinks it already has loaded some textures (typical Canoma name is "name0001.jpg") and will not load it again - which is fine if the texture has not changed. But if it has, then you will see the old, wrong texture in Bryce.  
Solution: you have to close the Bryce model after you deleted the old Canoma object, then open again and import the new one. Closing will delete the texture cache.

For the same reason, do not name two Canoma objects you want to import the same since the texture names will overlap and Bryce will not load the second imported objects textures correctly.

**Q** Canoma models in Poser - texture limitation.

**A** Poser 4.0 can only handle up to 255 textures which is fine for the usual figures since each has one texture or "skin", but Canoma models come with a texture per polygon, and that quickly adds up. If you import a more complex model into Poser and notice a "missing texture" that is probably the cause. There is no workaround right now other than to simplify your Canoma model so it has fewer polygons and thus textures.

**Q** What other Canoma resources can I find on the net ?

**A** The FAQ was summarized from ongoing discussions on the Canoma listserver at

[http://intech.stockton.edu/scripts/lyris.pl?enter=canoma&text\\_mode=0](http://intech.stockton.edu/scripts/lyris.pl?enter=canoma&text_mode=0)

and the Canoma online user forum at

<http://www.poserforum.com/messages.ez?ForumID=9139>

These are great places to find other users, answers to your questions and to share your ideas and work with others.

Some Canoma user sites:

<http://2k.si.edu/canoma/>

<http://www.tilenut.com/canoma>

<http://loki.stockton.edu/~ken/canoma/>

<http://hometime.net/canoma.htm>

<http://iravictor.com/pages/virtual.html>

<http://home.sol.no/~bjoernk/canoma.htm>

<http://www.halcyon.com/jeanluc/house/house.html>