

Maya 2015

[Interface](#)

[Main](#)

[Hotbox](#)

[Tool Marking Menus](#)

[Object Marking Menus](#)

[Camera Controls](#)

[Projects](#)

[Primitive Objects](#)

[Creation](#)

[Transforming Objects](#)

[Move Tool](#)

[Rotate Tool](#)

[Scale Tool](#)

[Manual Edits](#)

[Transforming Vertices/Faces](#)

[Snapping](#)

[Parenting](#)

[Grouping](#)

[Subdivisions](#)

[Hypergraph](#)

[Hypergraph: Hierarchy](#)

[Hypergraph: Connections](#)

[Outliner](#)

[Node Editor](#)

[Linking/Constraining Attributes](#)

[Image Planes](#)

[Curves](#)

[CV/EP Curve Tool](#)

[Pencil Curve Tool](#)

[Arc Curve Tool](#)

[Polygon Selection](#)

[Selecting Components via Mouse](#)

[Edge Selection](#)

[Vertex Selection](#)

[Face Selection](#)

[Selecting Components via Select Menu](#)

[Edge/Vertex/Face Selection](#)

[Edge Selection](#)

[Vertex Selection](#)

[Face Selection](#)

[Select Edge Loop Tool](#)

[Select Edge Ring Tool](#)

[Select Border Edge Tool](#)

[Select Shortest Edge Path Tool](#)

[Convert Selection](#)

[Select Using Constraints](#)

[Paint Select Tool](#)

[Polygon Extrude Faces](#)

[Polygon Extrude Face Along a Curve](#)

[Polygon Edge Loops](#)

[Polygon Mirror Geometry](#)

[Polygon Normals](#)

[Face Normals](#)

[Vertex Normals](#)

[Manually Editing Normals](#)

[Backfacing](#)

[Polygon Sculpt Geometry](#)

[Push/Pull Tool](#)

[Smooth/Relax Tool](#)

[Pinch Tool](#)

[Slide Tool](#)

[Erase Tool](#)

[Polygon Bridge](#)

[Basic Bridge](#)

[Curved \(Smooth + Curve\) Bridge](#)

[Blended \(Smooth\) Bridge](#)

[Polygon Mirror Cut](#)

[Polygon Bevel](#)

[Polygon Booleans](#)

[Add / Union](#)

[Difference / Subtract](#)

[Intersect](#)

[Polygon Quad Draw \(cleanup geometry / scan data\)](#)

[Extending Polygons](#)

[Relaxing Polygons](#)

[Soft Select](#)

Interface

Main

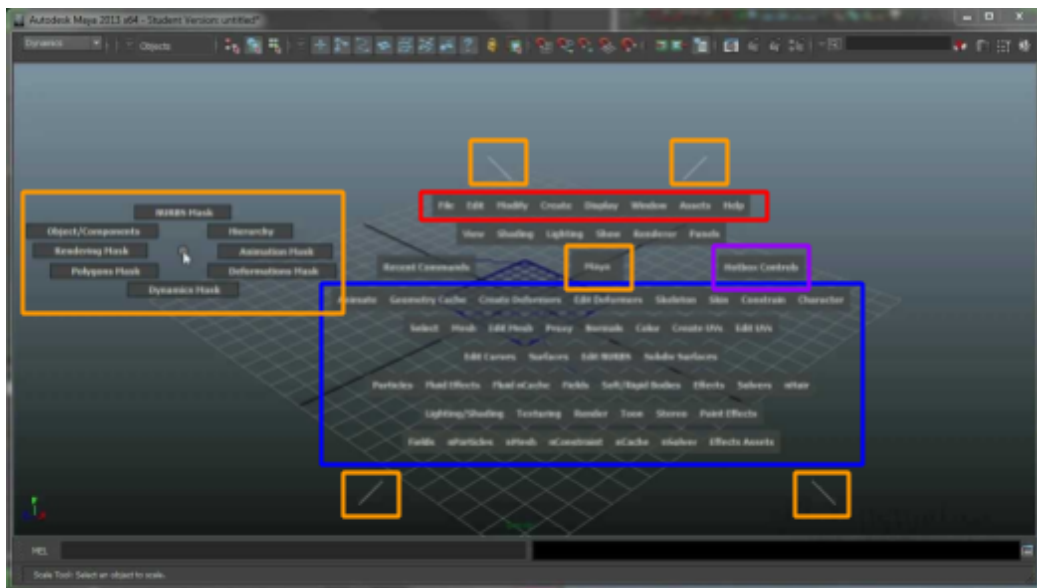


- Menu** Provides access to all of Maya's functionality: everything from object creation to rendering a final image. The first 7 menus and the last menu never change (common functionality). The menus in between change based on the option selected in the menu set.
- Menu set** Controls which menus show up.
- Toolbar** Typical toolbar: saving/loading, selection, magnets, etc...
- Shelves** Commonly used tools: creating objects, deforming objects, etc. You can create custom shelves, but Maya comes with several shells already setup for you.
- Channel box** Streamlined access to object attributes for quick edits.

- Layer** Organize objects, renders, and animations in layers. Layers are key to managing all of your assets inside of Maya.
- Toolbox** Selection + transform tools (move/rotate/scale).
- Quick layout** Shortcuts to commonly used viewport configurations.
- Timeline** Timeline, playback controls, and range slider. Together, these let you set the length of your animation and playback.
- MEL** Maya expression language (programming interface). Gray bar on the left is the input, while black bar on the right is the output.

Hotbox

The hotbox is an action menu that contains every “action” from Maya’s interface. You can access the hotbox by pressing and holding the spacebar.



Across the top of the hotbox, you get each of the **common menus**. On the bottom, you get each of the individual **menu sets**.

NOTE: Remember that common menus are the the first 7 + last main menus, while the menu sets are the menus that show up when you choose a specific item in the menu set drop down. Check out the interface section if you can’t remember what common menus / menu sets are.

The 4 white lines represent the **designated marking menus**. Left clicking in the ...

- center region gives you options for switching between different cameras
- top region gives you options for predefined viewport layouts
- left region gives you selection masks

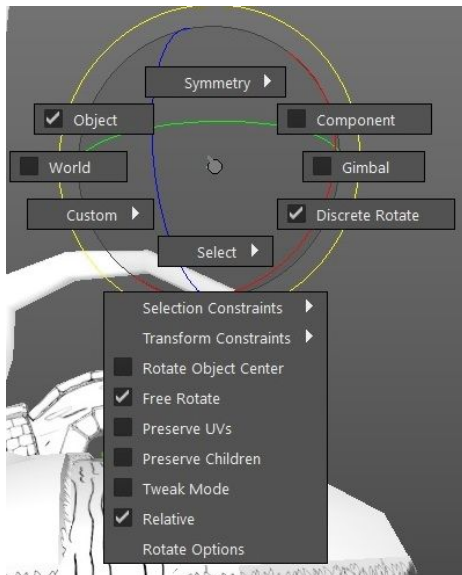
- right region gives you options for turning on/off UI elements (e.g. get rid of timeslider)
- bottom region gives you options to switch to different windows/editors (e.g. graph editor)

You can also customize the hotbox by using [hotbox controls](#).

Tool Marking Menus

Depending on the tool, you can get menus to pop up (similar to the designated marking menus talked about in the previous section) that let you access the configuration options for that tool.


For example, the W key in the hotkey for the move tool. If you press and hold the W key, you can left click anywhere in the viewport to show a pop-up menu similar to a hotbox designated marking menu. This menu will give you access to all of the tool's functionality, similar to if you double click the move tool in the toolbox and have the options show up in the tool settings pane.



This works with all of the transformation tools (E for rotate and R for scale).

Object Marking Menus

Similar to the tool marking menus discussed above, you can right-click on an object to get a marking menu specific to that object.



Camera Controls

Maya has 4 cameras by default...

- perspective
- top
- left
- side

You can toggle back and forth between them by using the space bar. To manipulate the camera, you can use the Alt in combination with any of the mouse buttons...

- Alt + LMB = tumble/rotate (only in perspective view)
- Alt + MMB = pan
- Alt + RMB = zoom (mouse wheel also works for this)

You can also access these tools in the menus just above the viewport's toolbar...


View -> Camera Tools -> (tool)

The history for the camera's position is retained, you can go back and forward using the bracket keys ([and]).

Sometimes you want to frame up objects -- have it take up the entire viewport. To do this...

- Press F to frame up the currently selected objects.
- Press A if you want to frame up all objects.
- Hold Ctrl+Alt and select area (top-left first) in viewport to frame up an area.
- Hold Ctrl+Alt and select area (bottom-right first) in viewport to frame up out of an area.

There are also camera controls in the viewport's toolbar.



1. First icon selects the camera as if it were an object.
2. Second icon selects the camera and open the attribute editor. Use this to set camera attributes like FOV.
3. Third icon bookmarks the current view.
To go to a bookmark, just above the viewport's toolbar select View -> Bookmarks -> (bookmark here)

Projects

When working inside of Maya, you always want to start by establishing a project directory. Once you have that project directory established, you only need to set it from there on out. Maya will remember that project directory and every time you start Maya it'll directly open up into that project.

NOTE: This is handled by setshot at SPI?


What is the purpose of having a project? It's so that Maya's directory paths are set to your current 'scene'. This makes sure that all of the external assets (e.g. textures) are loaded properly.

NOTE: This is a fancy way of saying that paths are relative, referenced against the project directory.

If an asset can't be found, Maya will use a placeholder. For example, if a texture can't be found you'll get an all black placeholder.

To change the current project... File -> Set Project.

To create a project... File -> Project Window. Press New, filling out the name/location fields, updating the location fields (if you want to), and hit Accept.



- Primary Project Locations -- this is the individual folders Maya will create for us
- Secondary Project Locations -- similar to primary but for lesser used things like fur
- Translate/Custom Data Locations -- other types of directories we want Maya to create

NOTE: At SPI, it looks like you can't do this. The system already sets it up for you when setshot is used and Maya starts. If you try to make a new project folder in that shot, you won't have the permissions to do so?

NOTE: A project's information is kept in the workspace.mel file. You'll find that file in the root project folder. Without that file, Maya won't be able to recognize the project structure.

NOTE: The two really important locations are scenes and source images. These are the ones that you usually access most often.

- Scenes -- scenes
- Source images -- external textures / images

Primitive Objects

There are 3 types of primitive objects in Maya...

- NURBs -- algorithmic curves, you did your gfx project on these algorithms
- Polygons -- basic flat polygon surfaces (triangle, triangle fan, etc..)
- Volumes -- implicit surface (used for fog/lighting effects)

NOTE: What is an implicit surface? It means you can't change any of the components/vertices that make this surface up. You can scale/rotate/move, but you can't modify.

Creation

To create a primitive object...

1. Create -> (type) Primitives -> (object type)
2. Click and drag on viewport

NOTE: You'll notice an checkable menu item called Interactive Creation near the bottom. If you turn this off, objects will be created at origin instead of click-dragging.

Primitive objects all come with a creation node (accessible under INPUTS in channel box). Changing the properties under this creation node modifies the components that make up the object.

For example, a NURBs sphere will have the properties sections and spans. These properties add/remove segments for the NURBs surface.

HINT: If you select an integer item in the channel box like sections/spans, you can hold the middle mouse button down in the viewport and move the mouse left/right to decrease/increase.

NOTE: Between NURBs and Polygon objects, polygons are more versatile and have the greatest amount of tools. NURBs aren't used that often anymore because they're harder to work with. With polygons, you can take any surface and subdivide it to do whatever you have to.

Transforming Objects


Remember that the basic controls for transforming an object are...

- Move tool (hotkey W)
- Rotate tool (hotkey E)
- Scale tool (hotkey R)

NOTE: You can access all these tools through the toolbox as well as through the hotkeys given.

When one of these tools is selected, you have control markers show up. For example, the move tool will give you arrows and planes that map to the world coordinates.

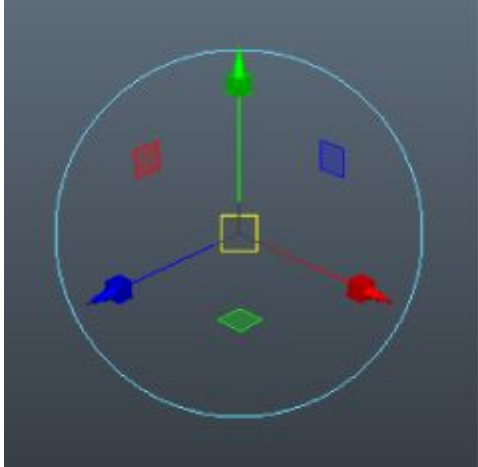
- Pressing + makes the markers bigger
- Pressing - makes the markers smaller
- Pressing INS allows you to move the pivot point in the center
 - You can also hold D and move the pivot point
 - If you want to go back you can go to the main menu and Modify -> Center Pivot



Move Tool

The move tool allows you to move an object. Depending on the marker you choose, you move in a certain axis/plane.

- Selecting arrows makes you move in that arrow's direction
- Selecting planes makes you move in that plane
- Selecting the pivot makes you move in all 3 dimensions (relative to viewport?)




One special thing to note is that you can change the axis to something other than object/local (e.g. local vs world). In the Tool Settings pane (double click the scale tool), you can find the Scale Axis option under the Scale Settings rollout (at the top).

Rotate Tool

The rotate tool is similar to the move tool. You can use the different markers to rotate in the desired.

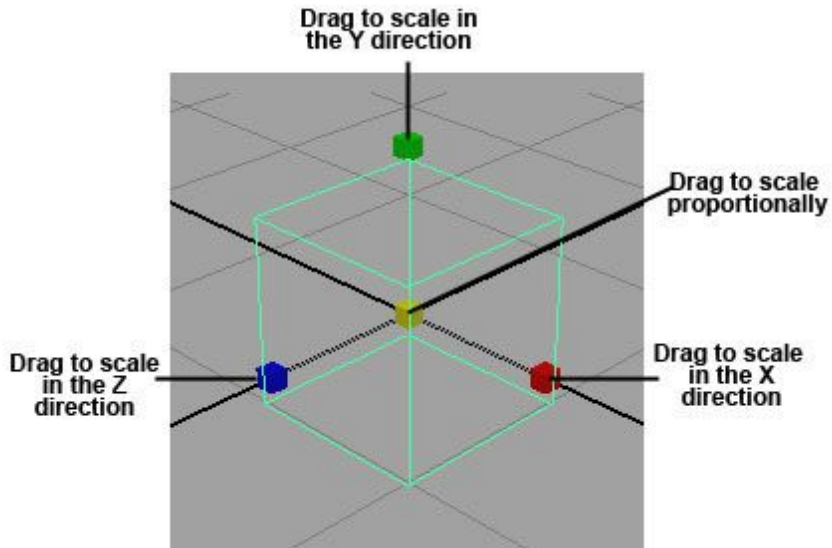
BE AWARE that this rotates your object around the pivot point, so if you moved the pivot point the markers will be off and you'll be rotating around the pivot point instead of around your object's center.



Scale Tool

The scale tool is similar to the move tool. You can use the different markers to scale in the desired direction.

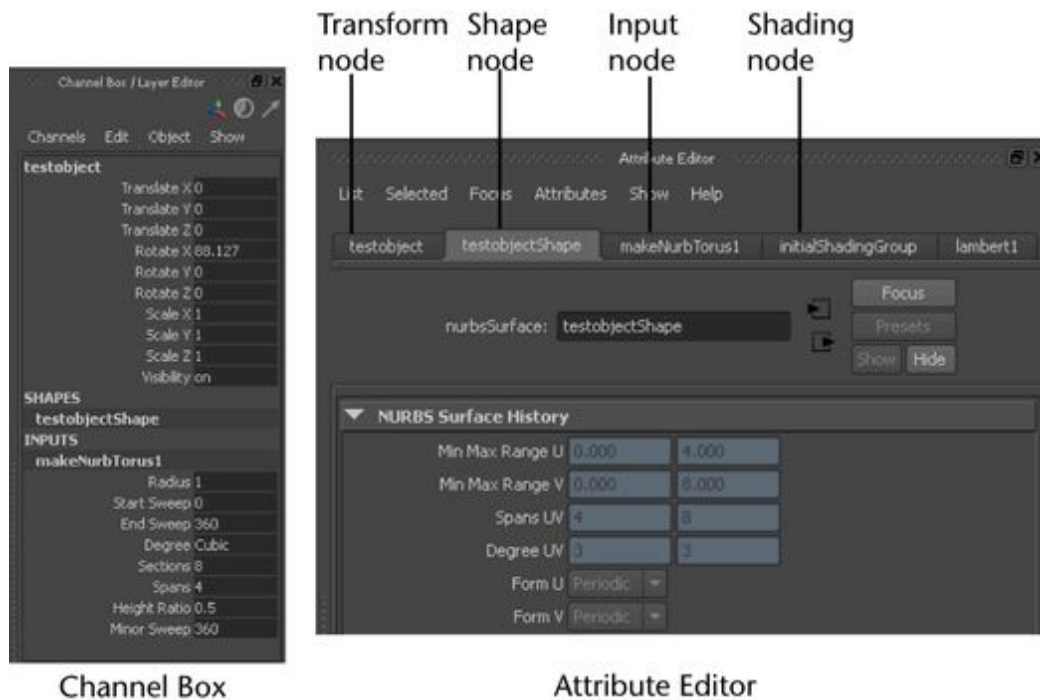
BE AWARE that this scales your object around the pivot point, so if you moved the pivot point the markers will be off and you'll be scaling based on some arbitrary point rather than the local object.



One special thing to note is that you can change the axis to something other than object/local (e.g. world). In the Tool Settings pane (double click the scale tool), you can find the Scale Axis option under the Scale Settings rollout (at the top).

Manual Edits

Manual edits to the translate/rotate/scale can be done via the channel box / attribute editor.



For example, if you want to move the object back to (0,0,0), you can change each translate coordinate back to 0.

HINT: You can select all 3 inputs at once and type 0 (channel box only)

HINT: You can select the **field label** (not the input field) and hold+move the middle mouse button in the viewport to increase/decrease numeric values (channel box only)

Transforming Vertices/Faces

To move a vertex:

1. Go into vertex/component selection mode
2. Click vertex
3. Press W for the move tool
4. Move vertex

To move a face:

1. Go into vertex/component selection mode
2. Click vertices that make up face
3. Press W for the move tool
4. Move vertices

You can't really rotate or scale a face?

Snapping

You can snap to different things by using the snapping menu / snapping hotkeys.



Snap to grids - (hold X) will snap to intersection points on the ground grid

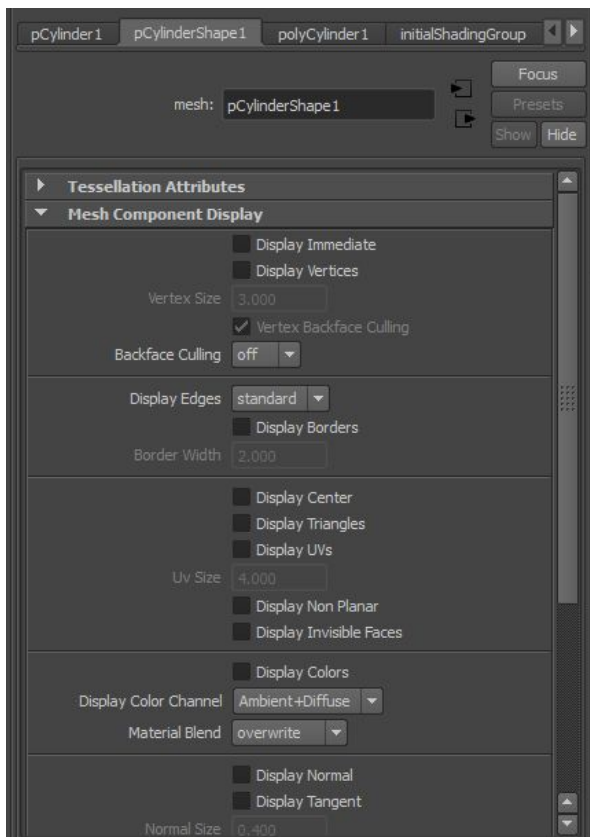
Snap to curves - (hold C) will snap to curves -- you can slide along the curve

Snap to points - (hold V) will snap to points -- this is useful, see note below

Snap to projected center - (no hotkey) unsure what this snaps to

Snap to view plane - (no hotkey) unsure what this snaps to

For snap to points, you can snap to various points on an object. You need to first select the object, go into the attribute editor and select the shape node (e.g. CylinderShape1), expand Mesh Component Display, and select what you want to display.



You can snap other objects to whatever you choose to display. For example, if you choose to display center, you can now snap to the center of each face on that object by using **snap to points**.

Parenting

Parenting allows you to make objects children of some other object. For example, object C is a child of object P, now whenever object P is selected object C is selected along with it (the opposite doesn't apply).

Moves/rotates/scales applied to the parent will be applied to the child as well. When this happens, the pivot of the parent object is the pivot of the child objects (e.g. children's rotation will be pivoted around parent's pivot).

You can still move/rotate/scale children independently by selecting them individually.

To attach child objects to a parent...

1. Hold Shift
2. Select the child objects
3. Select the parent object (MUST be selected last)
4. Hit P

To detach child objects from a parent...

1. Select the child objects (BUT NOT the parent object)
2. Hit Shift+P

HINT: As you've probably guessed, you can have a hierarchy chain of parents. You can use the arrow keys to traverse the hierarchy. Pressing up will take you to the parent.

Grouping

Grouping is similar to parenting, except that the parent will be a newly created null object.

NOTE: A null object is an object without any geometry (nothing renderable). It's sole purpose in this case is to anchor the other objects.

To group objects together...

1. Hold Shift
2. Select the objects
3. Hit Ctrl+G

To ungroup objects...

1. Hold Shift
2. Select the objects
3. Hit Shift+P

Subdivisions

Subdivisions are like a hybrid between NURBs features and polygon features. For example...

1. Create a polygon cube
2. In the attribute editor, go to Smooth Mesh rollout
3. Enable the Smooth Mesh Preview checkbox (hotkey is 3 to enable / 1 to disable)

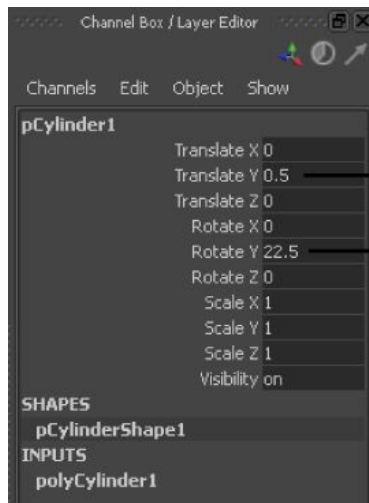
You should now see that the cube is smoothed into a sphere. But, if you try to move the vertices of the sphere, you'll notice that the vertices show up where they would be on the original cube.

Directly below the Smooth Mesh rollout is the Subdivision Levels rollout. Enable the Display Subdivisions checkbox to see the subdivisions being created.

Hypergraph

Maya is a node-based editor -- this means that everything we create is a separate node. The hypergraph gives us a schematic view of the nodes in our scene file.

Take a basic polygon sphere as an example. If you create a cylinder and look in the channel box, you'll see multiple headings: pCylinder, pCylinderShape, and polyCylinder. Each one of these is a separate node.



We can look at these specific nodes inside the hypergraph.

HINT: Always keep the hypergraph open and keep an eye on things. It helps identify problems / helps you understand what's going on in your scene / helps identify things that you thought you got rid of but are still there (e.g. null objects?).

You can consider the hypergraph as the "hub of Maya." It's super important that you get use to using it.

Hypergraph: Hierarchy

Open the hypergraph by going to the main menu: Window -> Hypergraph: Hierarchy

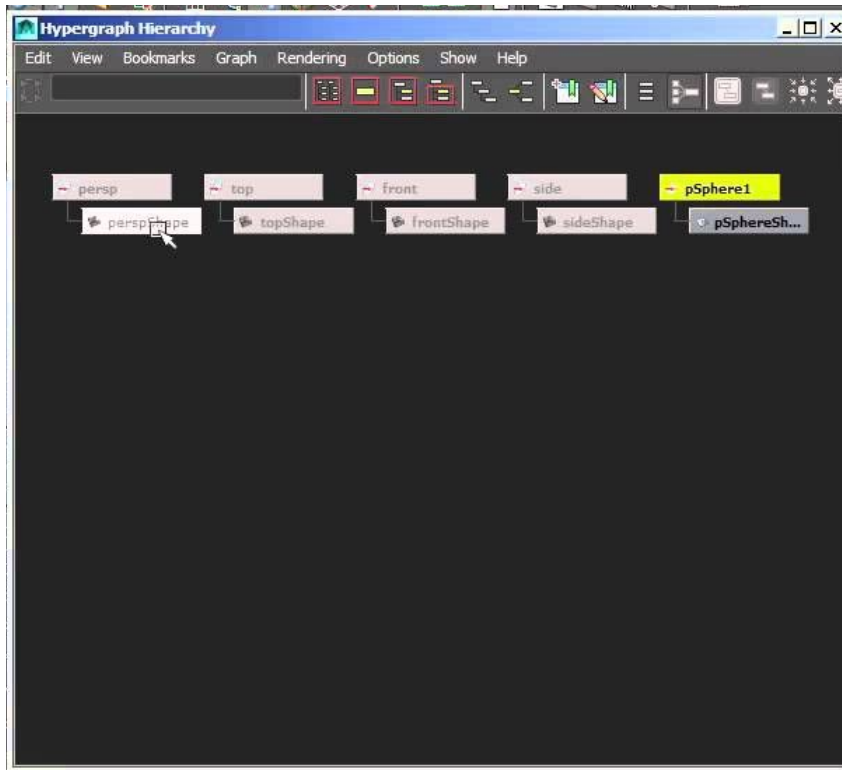
NOTE: Notice that there are 2 different hypergraph options: Hierarchy and Connection. Ultimately, they both lead to the same thing. You can switch between the two via the toolbar in the hypergraph window.

The window that pops-up is the top level hypergraph view. You'll see the objects you created here.

HINT: To maneuver around the hypergraph, you can use the typical camera controls to pan/rotate/zoom (ALT + LMB/MMB/RMB). Hit A to frame all. Hit F to frame the currently selected nodes.

If you want to see all nodes, check the following boxes in the main menu...

1. Options -> Display -> Hidden Nodes (brings up stuff like camera nodes)
2. Options -> Display -> Shape Nodes (brings up the nodes we saw for our cylinder -- pCylinderShape)



Top level nodes (e.g. pSphere or pCylinder) are transform nodes. They contain our transform values like translation and rotation. The node underneath it (e.g. pSphereShape or pCylinderShape) is the node that actually holds the geometry.

HINT: Click and hold the middle mouse button to a transform node and move it under another transform node to automatically create a parent/child relationship. To undo, click and hold middle mouse button and move to empty space.

HINT: Double-click node to hide/expand children. Right-click for further options (e.g. expand all)

The hypergraph shows us everything that makes up our scene. There can sometimes be thousands of nodes in here, so it's best to not show hidden nodes and use the collapse/expand feature. You can also turn off the automatic layout and switch to free-form layout, which lets you move around items and organize them yourself (available in the toolbar of the hypergraph window).

Hypergraph: Connections


Open the hypergraph by going to the main menu: Window -> Hypergraph: Connections (MAKE SURE YOU HAVE AN OBJECT SELECTED WHEN YOU DO THIS)

NOTE: Notice that there are 2 different hypergraph options: Hierarchy and Connection. Ultimately, they both lead to the same thing. You can switch between the two via the toolbar in the hypergraph window.

The window that pops-up is the connections hypergraph view. You'll see the objects you created here + the connections that make them up.

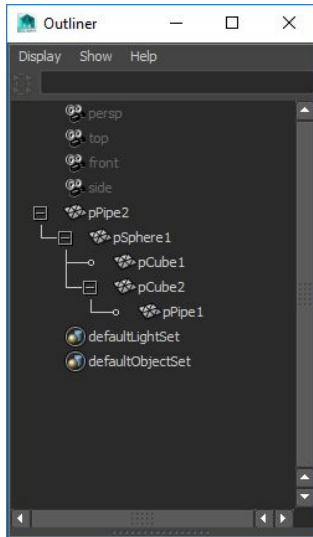
NOTE: A lot of the nodes that are being fed in are construction history nodes? These contain the history of steps that were taken to get the object to its current state??? If you go to Edit -> Delete By Type -> History, it seems to get rid of them but the objects all remain.

HINT: To maneuver around the hypergraph, you can use the typical camera controls to pan/rotate/zoom (ALT + LMB/MMB/RMB). Hit A to frame all. Hit F to frame the currently selected nodes.



Outliner

The outliner is very similar to the hypergraph. It gives you a tree of nodes in your scene.



Just like with the hypergraph, the outliner allows you to see parent child relationships between your nodes. Unlike the hypergraph, the outliner won't allow you to see connections between the nodes.

A lot of what you can do with the hypergraph interface. Here are the important things you can do in the outliner...

1. Click and hold MMB on a node and move it onto another node to make it a child
2. Click and hold MMB on a node and move it in between other nodes to re-order
3. Selecting a node in the outliner selects it in the perspective view (and vice versa)
4. Right-click and choose "Reveal Selected" to have the outliner go to the specific node

If you want to clean up the outliner a bit, you can hide nodes by

1. Select nodes
2. Right-click -> Hide in Outliner -> Hide

You can unhide nodes by...

1. Display -> Ignore Hidden
2. Select all nodes
3. Right-click -> Hide in Outliner -> Unhide


OR

1. Select node in perspective view
2. Go back to outliner and Right-click empty space -> Hide in Outliner -> Unhide


You can limit the type of objects shown in the outliner by ...

1. Show -> Objects -> (object type)

For example, if I choose Cameras, everything that isn't one of my camera nodes will not show in the outliner. Note that when you do this, the icon to the left of the textbox turns green to indicate that some constraint is active. You can click it to clear the constraints.



NOTE: If you look closely, you'll notice a divider at the very bottom of the outliner. Click and drag it to expose a second instance of the outliner. This is useful because it allows you to work on a location but see another location.



NOTE: You can choose to see the individual shape nodes by going to Display -> Shapes. People usually keep this off because it makes things a bit confusing by making the outliner a bit overpopulated.

Node Editor

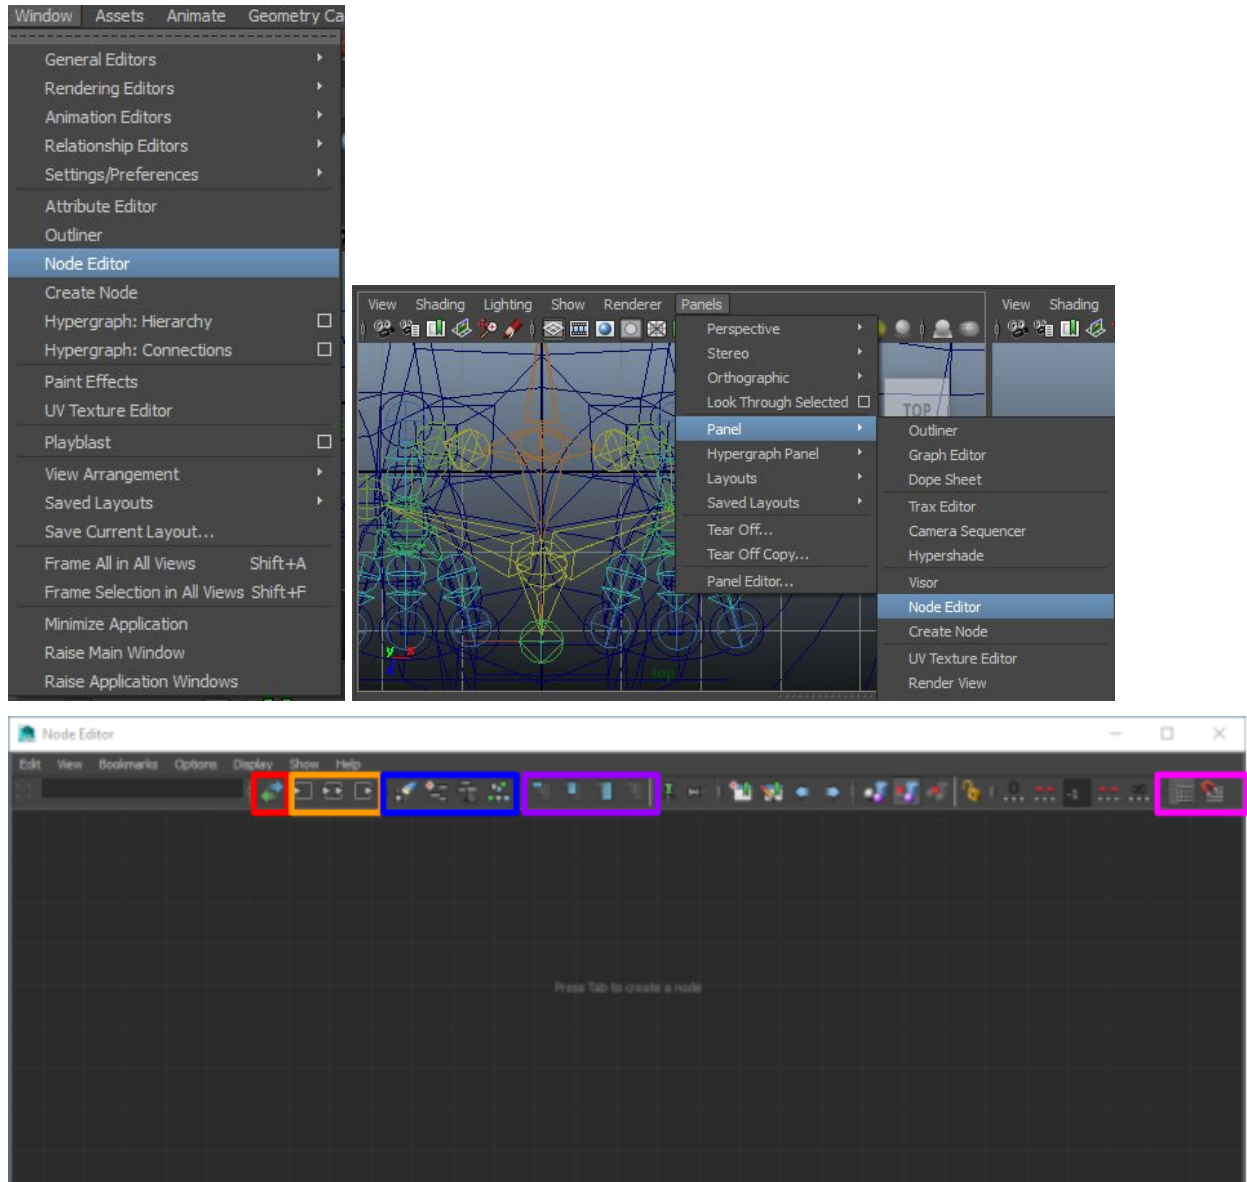
The node editor is a lot like the hypergraph/outliner. You can view the nodes in your scene, expand them out to see their properties, and even link (constrain) certain properties together.

For example, you can make it so that the translation of an object is copied to the translation of another object.

There are 2 ways of opening up the node editor.

1. In a new window: Window -> Node Editor
2. In a panel: Panels -> Panel -> Node Editor

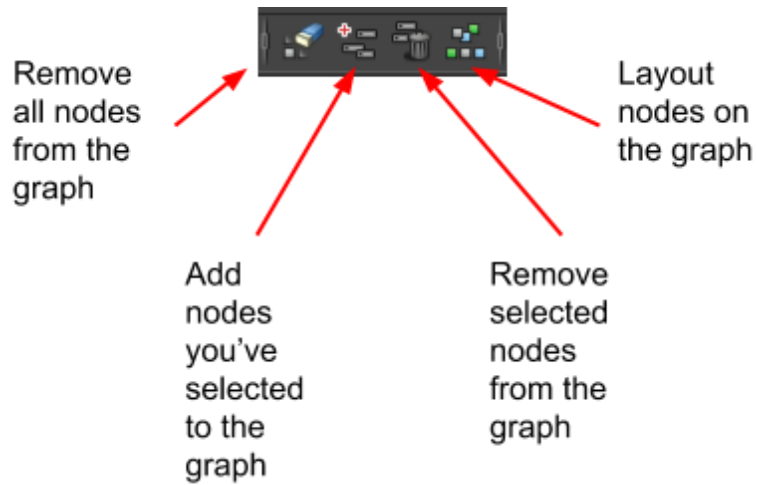
NOTE: You can have multiple node editor instances open at once if you want.



NOTE: You can treat the node editor like a viewport. F to frame up on a node, A to frame up on all, move around like you normally would in the viewport (Alt+MMB), zoom in like you would normally do (Alt+RMB or scroll wheel)

Nodes

Use this to add/remove nodes to the node editor, as well as layout nodes. Buttons should be self-explanatory from the icons / tool tips.




NOTE: To add a node, you have to have that node selected somewhere (e.g. in your viewport/hypergraph/outliner/attribute editor/etc...). If you select it and hit the add button, it'll show up in the node editor.

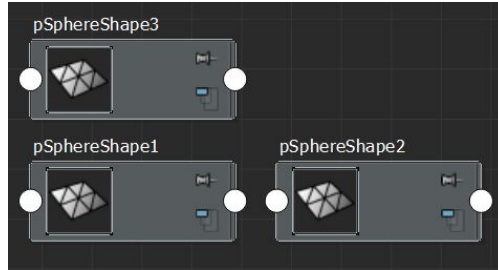
NOTE: When you're removing nodes from the graph, you're not actually deleting them from your scene. You're only removing them from the node editor. The nodes still exist.

Connections


Use this to include the inputs/outputs for the selected node(s) in the node editor.



For example, imagine I had the following objects in my node editor...



When I select pSphereShape1 and click the show outputs button, everything other than pSphereShape1 and its outputs get cleared.




Sync

Synchronizes what's selected in the node editor with what's selected in the outside world. For example, if you select a box in one of your viewports that same box will be selected in the node editor (assuming that the node for the box has been added to the node editor).

You can turn this off if you want to inspect your objects in the viewport without messing up your selections in the node editor.


Expand

Use these buttons to expand out nodes in the graph to see their properties/attributes. There are 3 levels available.

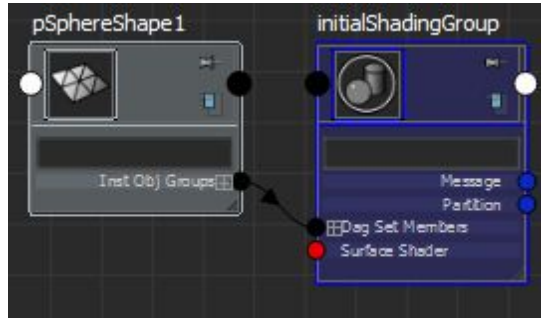


NOTE: Notice how the same button is available on individual nodes. You can expand/collapse by individual node instead of using the toolbar version. The toolbar version will collapse/expand ALL nodes in the node editor.

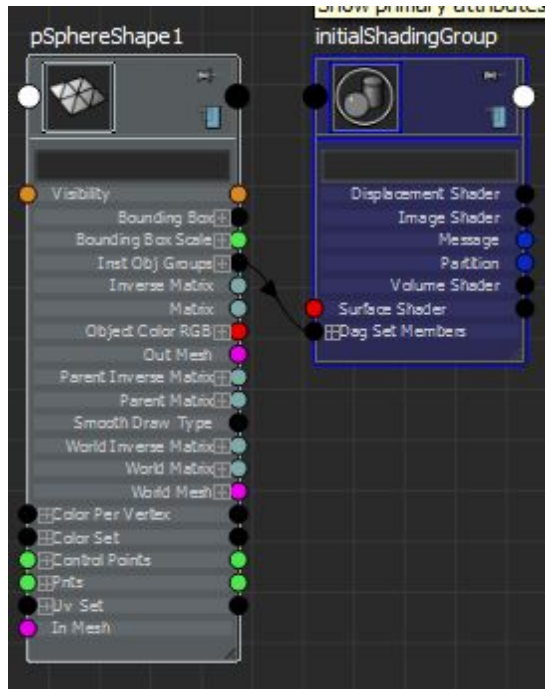
1st level completely collapses the node(s)...



2nd level shows only attributes that are connected to other nodes...



3rd level shows all primary attributes...




Snap

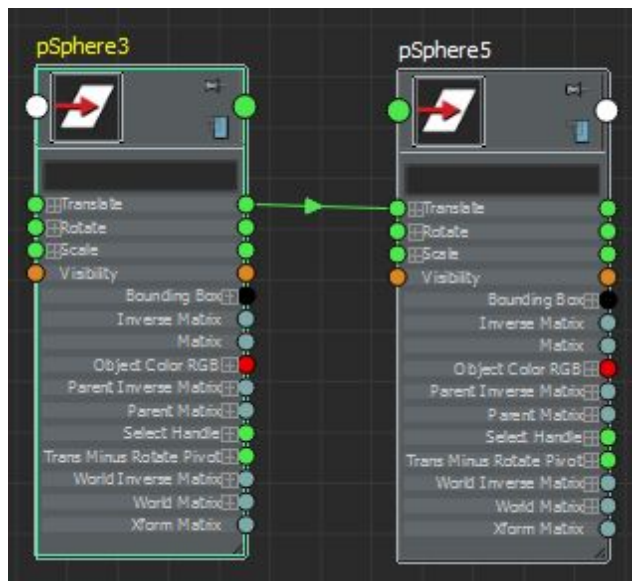
Use these buttons to toggle grid visibility / grid snapping. Pretty self-explanatory.

Linking/Constraining Attributes


You can link attributes from one node to others, such that the other nodes will take on those values for the linked attribute. For example, imagine I had 2 spheres...



I want to make it so that every time I translate the first sphere, the other sphere picks up those same translation values. I can do this by adding these spheres to my node editor and linking the translate attributes together.



Now when I move the right-most sphere, the left-most sphere moves the same.



NOTE: The spheres here had their transforms freezed. If they didn't have their transforms freezed, the right-most sphere would get superimposed on the left-most sphere.

If you open up the right-most sphere in the channel box and look at the translate attributes, you'll see that they're yellow. Yellow means that the values are coming from elsewhere / you can't directly change them.






Image Planes

Image planes allow us to bring in a background image to serve as an environment, background, or reference for modelling.

To create an image plane... Create -> Free Image Plane



This generates the image plane node, which we can then assign our texture/image to via the attributes editor (Ctrl+A). In the following example, we assign a reference drawing of a crow, which we can then view in our various cameras and use as a reference to model our crow.



NOTE: Note that this node is transformable -- you can move/rotate/scale it as desired.

NOTE: You can load multiple image planes. For example, it's normal to have a different image plane node for each orthographic projection (top/side/front). In the attributes editor, you can make it so that an image plane node is only visible to a certain camera (Display option).

NOTE: If you're using an image plane as a reference for modelling, note that moving it forward/back won't cause any distortion in the orthographic cameras (remember how orthographic cameras work -- projected objects don't get smaller as they're moved farther back). It's normal to move reference images back because sometimes what you're modelling may intersect with the image plane. Another thing you can try is to let it intersect but change the alpha value of the image plane (in attribute editor).

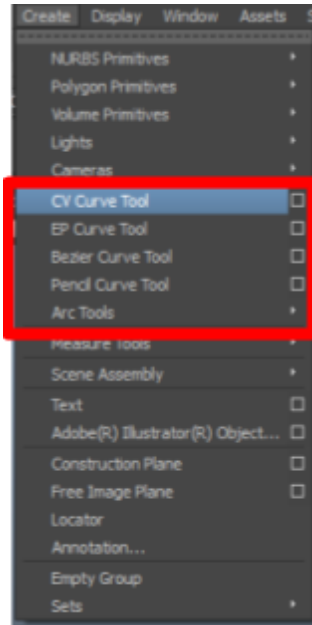
NOTE: In our example, we're using a drawing (of a crow) so the image lines up pretty well. However, if we were using photo instead, we'd get camera distortion due to the lens of the camera that photo was taken in. In that case we kind of just accept it and pick a view that works best for us / improvise with our other views as we model.

Curves

Curves are non-renderable objects used throughout Maya for a number of different applications...

- Hair
- Surfaces
- Guides for other nodes/objects (animation?)


You can find the different types of curves that Maya has under the Create menu




CV/EP Curve Tool

CV/EP Curves are the most commonly used curves. They're essentially the same type of curve, just with a slightly different manipulation mechanic:

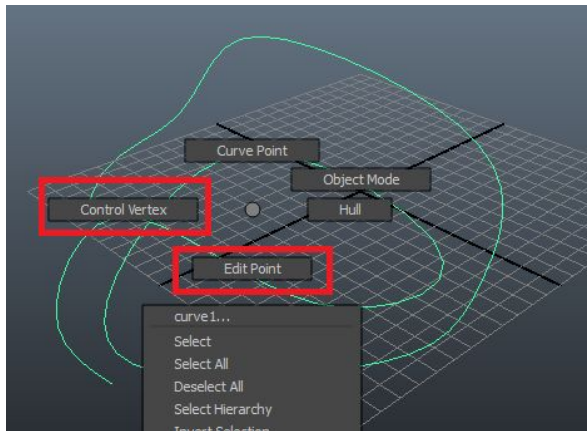
- **CV = Control vertex.** This means that you control how your curve bends using control points (remember how splines/bsplines/nurbs worked from school).



- **EP = Edit point.** This means that you control how your curve bends using edit points -- the nodes that shape the curve are actually on the curve itself.



Once the curve has been created, you can edit the points via EP or CV (regardless of which mode it was initially created in). Just select the curve and right-click to show up a marking menu.



NOTE: Before creating a CV/EP curve, double-click on the curve tool icon in the toolbox (it'll show up as the most recently used tool when you go to Create -> EP/CV Curve) and double-click. In the Tool Settings, you can choose the degree of the curve. YOU CANNOT CHANGE THIS AFTER THE CURVE IS CREATED.

Pencil Curve Tool

The pencil curve tool allows you to freehand a curve and it'll attempt to turn that into a curve (very similar to what inkscape's pencil tool does). Most of the time this isn't really used because of the jaggy nature of drawing by hand -- you end up with a curve with lots of components.

Arc Curve Tool

Sets up a basic 2-point or 3-point arc. Left-click on the viewport to set the points for the arc. For the last point, you can click-and-drag to see how the end result will look before you release.

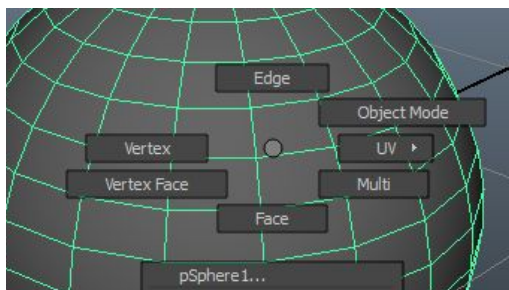
Polygon Selection

Polygons are n-sided faces where $n > 2$ (3=triangle, 4=quad, 5+=n-sided polygon). If you remember from school, there are other properties here as well. For example, for polygons to render properly they have to be convex (maya may implicitly handle this).

There are 2 ways to create polygons.

1. As a primitive shape... Create -> Polygon Primitives -> (primitive type)
2. As a single polygon... Mesh Tools -> Create Polygon Tool, then click around viewport and enter when finished

Once you have your polygon object, you can select it right-click to show a marking menu that allow you to select by...



- Vertex -- obvious
- Face -- obvious
- Edge -- obvious
- Vertex face -- this is an exploded view of the object, which is useful for debugging (e.g. if you have duplicate vertices at the same position)
- UV -- like vertex selection, but this is used for projecting textures on our object (you can't transform the vertex here)


A useful tool to have displayed when working with polygon objects is the Polygon HUD. You can enable this by going to...

Display -> Heads Up Display -> Polycount

Your viewports should now contain information about the scene, selected object, and selected components. Column number ...

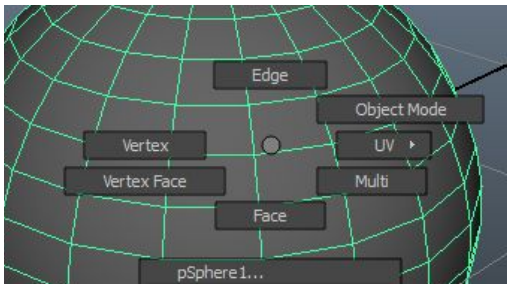
1. contains polygon stats for the entire scene
2. contains polygon stats for the selected object(s)

3. contains polygon stats for the selected component(s)



Selecting Components via Mouse

Right-click on an object to give you the typical polygon marking menu.



From here you can select by the component edge you want.

Edge Selection

Typical text selection rules apply In edge mode. To select ...

- a single edge: Click an edge
- row/ring of edges: Double click edge in that row/ring
- multiple selection: Hold shift and click/double-click one or more edges

Vertex Selection

Vertex mode is like edge mode, but with some caveats. To select ...

- a single vertex: Click a vertex

- row/ring of edges: Click a vertex, then hold shift and double-click a second vertex
- multiple selection: Hold shift and click/double-click one or more edges


Face Selection

Face mode is like edge mode, but with some caveats. To select ...

- a single face: Click a face
- all faces: Double-click a face
- row/ring of faces: Click a face, then hold shift and double-click a second face
- multiple selection: Hold shift and click/double-click one or more faces

Selecting Components via Select Menu

In the main menu, the select menu provides even more functionality/tools for selecting polygons. Note that you have to be in the polygon menu set or else you won't see this menu!



Edge/Vertex/Face Selection

These are exactly the same as their marking menu counterparts, but be aware that when you select via this menu it's entering you into a global selection mode. Using the marking menu puts you into a local selection mode (local to the selected object).

Selection details have been replicated here from the mouse section...

Edge Selection

Typical text selection rules apply In edge mode. To select ...

- a single edge: Click an edge
- row/ring of edges: Double click edge in that row/ring

- multiple selection: Hold shift and click/double-click one or more edges

Vertex Selection

Vertex mode is like edge mode, but with some caveats. To select ...

- a single vertex: Click a vertex
- row/ring of edges: Click a vertex, then hold shift and double-click a second vertex
- multiple selection: Hold shift and click/double-click one or more edges


Face Selection

Face mode is like edge mode, but with some caveats. To select ...

- a single face: Click a face
- all faces: Double-click a face
- row/ring of faces: Click a face, then hold shift and double-click a second face
- multiple selection: Hold shift and click/double-click one or more faces


Select Edge Loop Tool

This is similar to how you select rows/rings with the edge selection tool. The difference is that you can now select edge loops even if you aren't in edge selection mode (e.g. you can select and edge loop in vertex mode instead of edge mode).



Select Edge Ring Tool


This is similar to the select edge loop tool, but instead it gives you the ring of edges that are adjacent to your loop.



Select Border Edge Tool

This tool allows you to select edges that are on the borders of the object. Border means that the edge isn't shared with another polygon.


In the example below, we took a sphere and removed the bottom faces. Then we clicked the chose one of the bordering edges with the tool and all the bordering edges got selected.



Select Shortest Edge Path Tool


Selects the shortest path (edges) between two vertices. Really basic. Continually single click vertices to build the shortest path between them.

NOTE: This works in vertex mode and UV mode only.

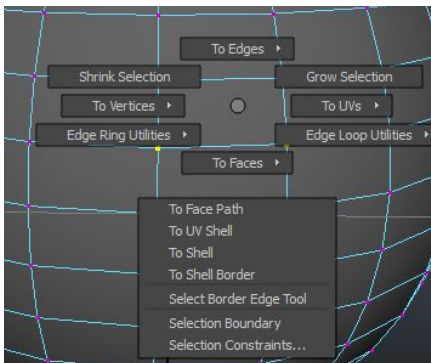


Convert Selection

Convert selection contains a set of really useful tools that allow you to switch your selection between different component types (e.g. faces to vertices).



You can also get this as a marking menu by selecting your components and hitting Ctrl+RMB.

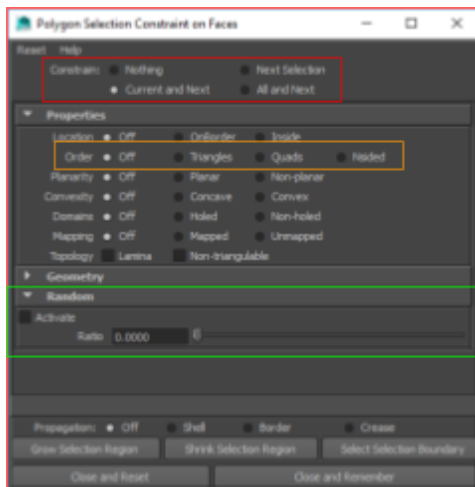


The basic conversions (this are all obvious)...

Vertex -> Edge: selects edges vertex is a part of
 Vertex -> Face: selects faces vertex is a part of
 Vertex -> Contained Face: selects faces when all vertices that are a part of it are also selected
 Edge -> Vertex: selects vertices edge is a part of
 Edge -> Face: selects faces edge is a part of
 Edge -> Contained Face: selects faces when all edges that are a part of it are also selected
 Face -> Vertex: selects face as vertices
 Face -> Vertex Perimeter: selects vertices on the outside (perimeter) of selected faces
 Face -> Edge: selects face as edges
 Face -> Edge Perimeter: selects edges on the outside (perimeter) of selected faces

Select Using Constraints

This tool allows you to constrain what gets selected. It only seems to work on components (not objects).



I'm unsure what all these options do. The main option is called **constrain**. It defines which selection your changes will have an effect on: current selection, current and next selections, next selections, etc... The one you want is probably current and next.

The other 2 options that seem to be somewhat useful...

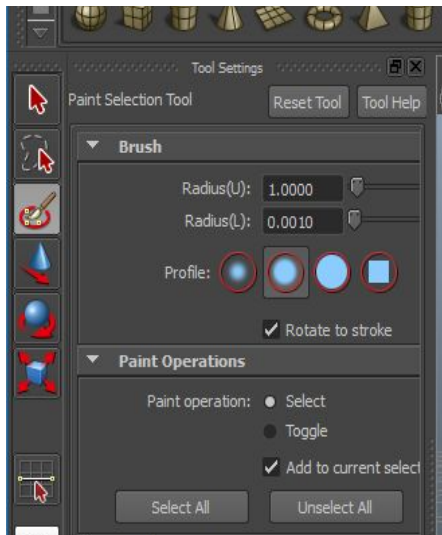
- If you only want to select faces that are triangles (not quads or n-sided), you would set the **order** option to triangles.
- If you want to select faces at random, activate the **random** option and set the ratio. Then, click and drag over the area for randomized selection.

NOTE: This dialog is for faces, but there's an equivalent with somewhat similar options when you're in vertex/edge mode.

NOTE: When closing the dialog, always choose Close and Reset, never Close And Remember. Otherwise your selection constraints will hang around.

Paint Select Tool

The paint select tool can be accessed through the toolbox. It allows you to select wide regions of faces/edges/vertices, as if you're spray painting.



NOTE: A quick way to resize the radius is to hold B, then click and drag to the right or left.

NOTE: Hold CTRL when painting to de-select components.

Polygon Extrude Faces

There are two ways to use the extrude tool


1. Mesh Tools -> Modeling Toolkit -> Extrude (in the popup panel/window)
2. Mesh Tools -> Extrude Tool
3. Edit Mesh -> Extrude

Options #1 and #2 lead to the same thing. Option #3 gives you a few extra features in the UI that you can use to tweak (although these are also available in 1 and 2 in the modeling toolkit).

NOTE: Extrusion is a really powerful tool for box modeling. Box modeling is when you start with a primitive box, apply symmetry/mirroring (explain how to do this in the next


note), and extrude out the quads to get a rough model of what you want. This seems to work really well if you have drawings you want to load up via image planes.

NOTE: To apply symmetry, first open up the modeling toolkit. Then go to your object and select an edge that directly cuts across the object. Once selected, go back to the modeling toolkit and click the big 'Symmetry' button. Now every component selection you make will be mirrored.




One thing to note when you extrude: if you extrude faces that are direct neighbours, they'll extrude as if they're stuck together.

- If you're using the modeling toolkit, you can turn this off by UNSELECTING the 'Keep Faces Together' option.
- If you're using the basic Edit Mesh -> Extrude option, you can go to the newly created extrude node in the attribute editor and turn off 'Keep Faces Together' option.




HINT: If you want to keep extruding in the same direction, use Ctrl+Shift+LMB to end the current extrude and start a new extrude from that point



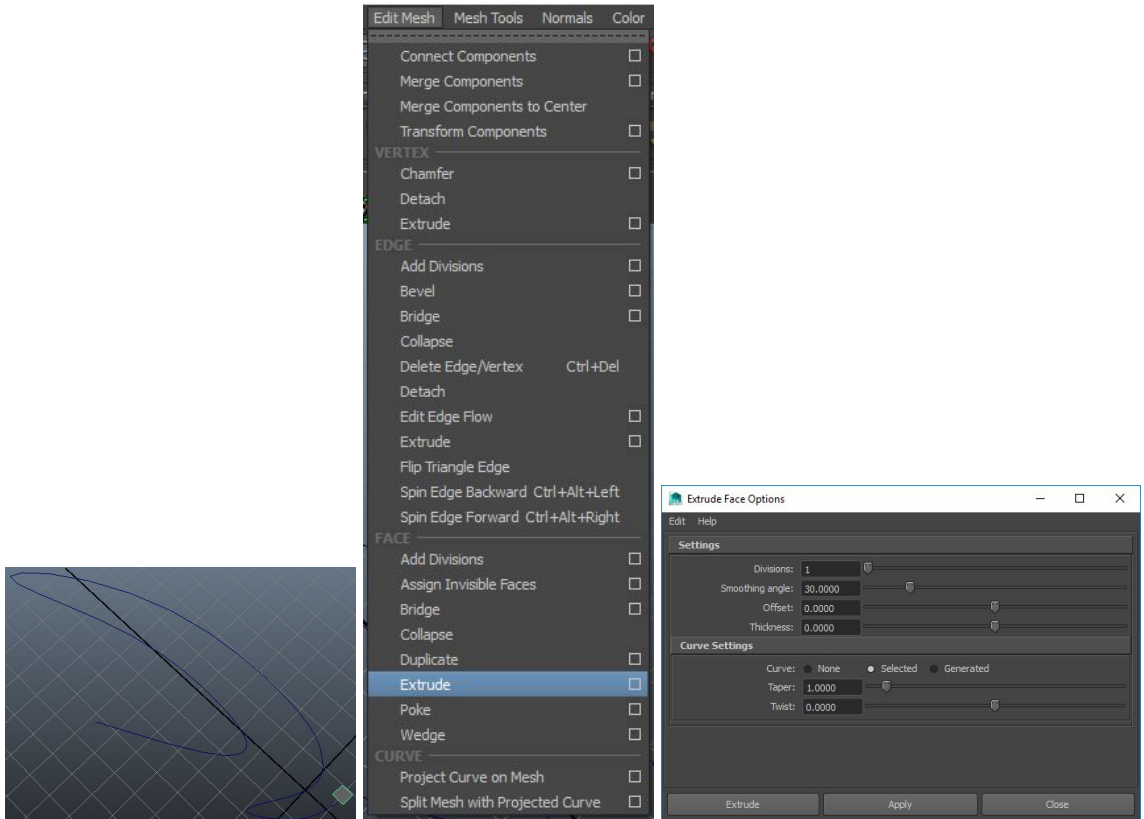
Polygon Extrude Face Along a Curve

If you want, you can choose to extrude along a curve rather than manually orienting and extruding yourself. To do this...

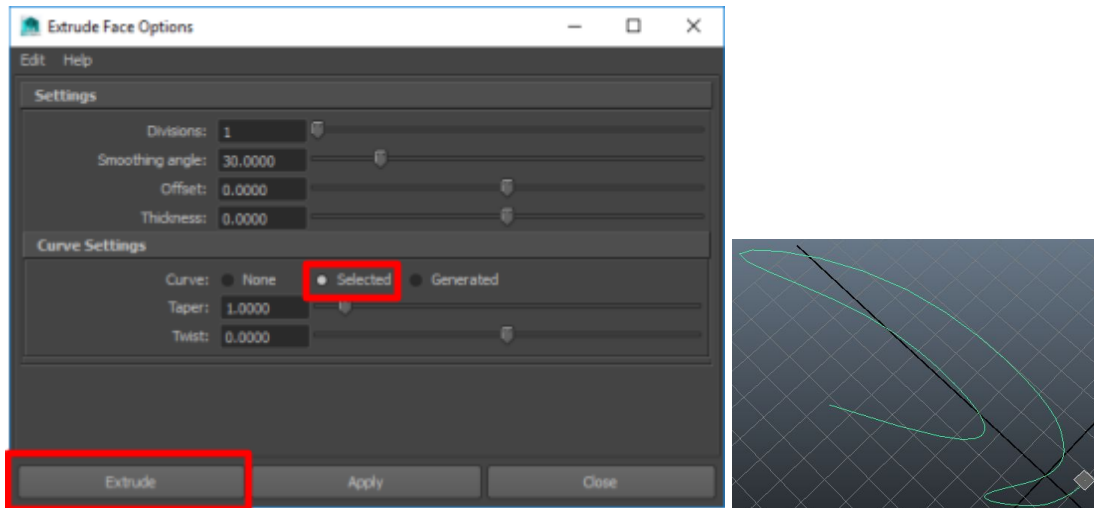
1. Create your curve.




2. Orient the face you'll be extruding on either edge of the curve, then go to Edit Mesh -> Face -> Extrude and choose the square on the right-hand side to open the options.



3. In the extrude face options dialog, set Curve to selected and select the curve you want the extrude to follow in your viewport. Hit Extrude when you're done.




4. In the channel box/attribute editor, adjust the number of divisions.




5. If you notice some weird distortion with the extrude near the end of the curve, such as the one shown above, that's likely because you placed the face at the end of your curve rather than the beginning.


To fix this, simply select the curve, change the menu set to Surfaces, and select Edit Curves -> Reverse Curve Direction.




Now the curve looks normal, but all the normals have been flipped. To fix this, go to Normals -> Reverse.



6. Note that edits to the curve after the fact will affect the extrude and everything else you did. That's because the edits to the curve are being made on the history node. As such, everything else you did follows along.



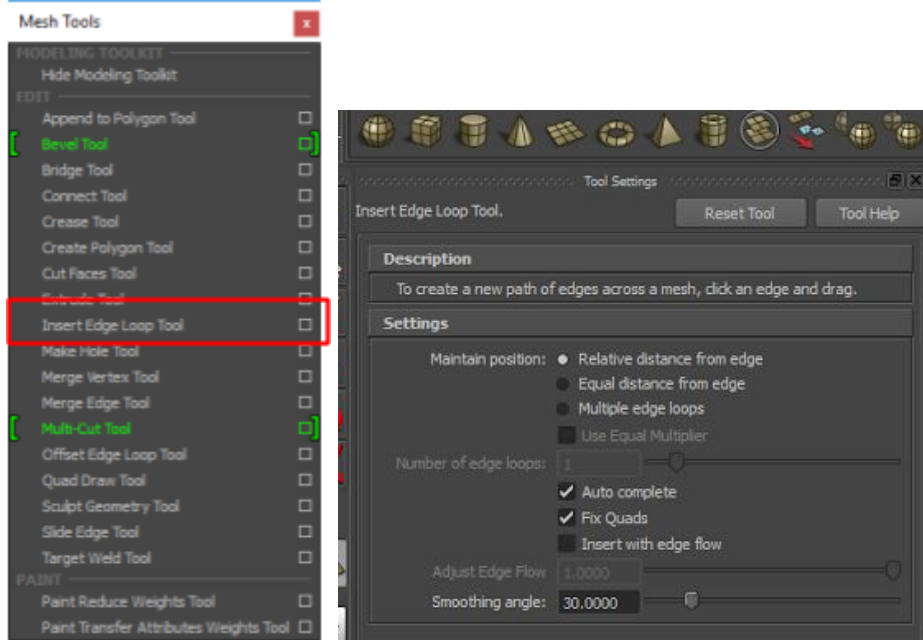
7. Apply twist and taper on the extrude node if you'd like. In this example, taper was set to 0.5, twist to 1440, and the normals were softened + the horizontal edge normals were hardened.




NOTE: Attribute editor will likely have finer graph-based tapering controls.

Polygon Edge Loops

You can create edge loops by going to Mesh Tools -> Insert Edge Loop Tool (click on the box to the right of the menu item text to open up the tool setting panel)




With the tool enabled, left click and hold to on an edge to insert an edge loop



This creates an edge loop that's exactly planar to the faces it's being created on. If you want the edge to be extruded/shrunk a bit to match the shape of whatever it is you're modeling, select the Insert with edgeflow checkbox in the tool settings. You can adjust the edge flow parameter in the slider below (can make it extend past 1.0 by typing in a number).

Here's that same edge with a 5.0 edgeflow.



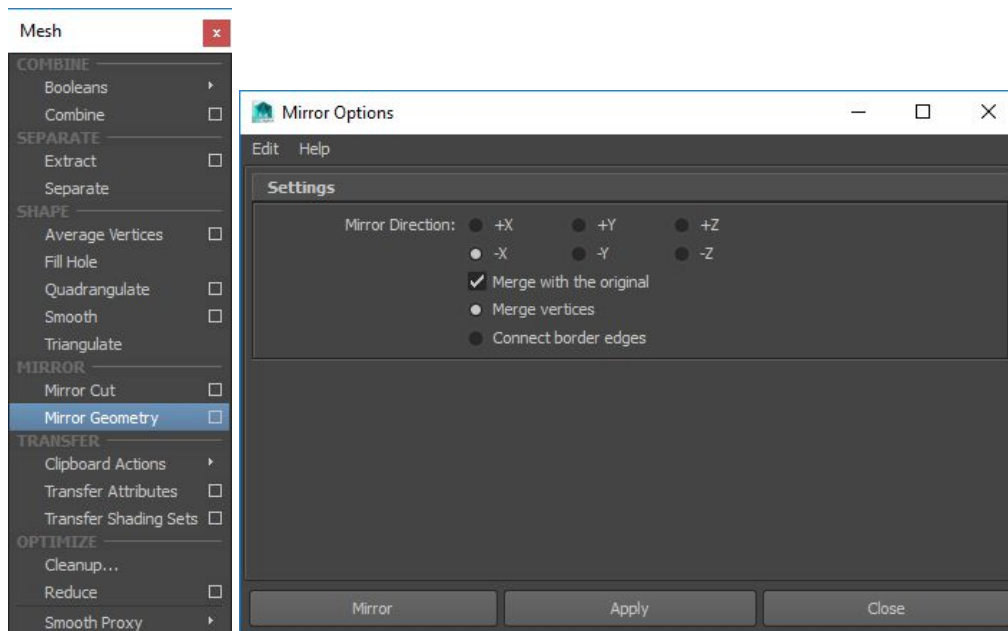
If you want to create an edge loop that is exactly in between the exactly in between the edges of the faces it's being created on, here's a neat trick:

In the tool settings, select the Multiple edge loops option, but set the Number of edge loops option to 1. Now when you create an edge loop it'll be created dead center.

NOTE: You can also do this after the fact by going to the history node for the edge loop (named polySplitRing#) and setting the weight to 0.5


Polygon Mirror Geometry

You can mirror geometry by going to Mesh -> Mirror Geometry (click on the box to the right of the menu item text to open up the tool setting panel)



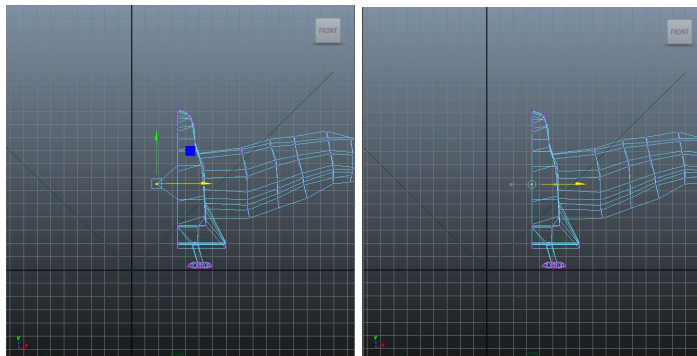
- Mirror Direction defines which axis the mirror should show up on
- Merge with the original will attempt to merge with what's being mirrored

If you're merging with the original, make sure that all the vertices that are going to be merged are sitting on the plane of the direction that you're mirroring on. For example, if we're mirroring on -X, we want the following...

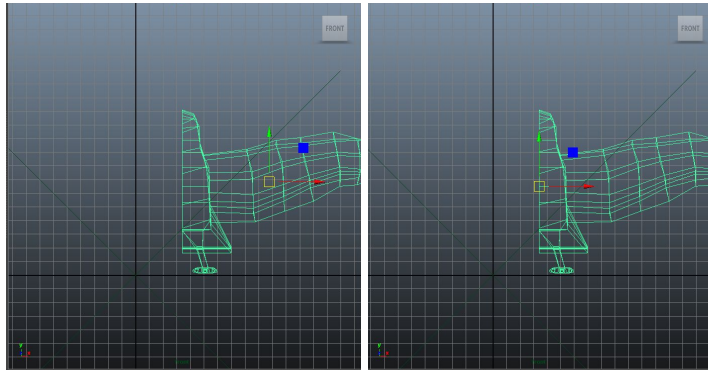


HINT: If you're merging but some of the vertices being merged aren't perfectly aligned to a plane (there's either a gap or a protrusion). Sometimes you can try going to the history node after the merge and fiddling with the Merge Threshold to see if you can get things to work. A lot of the times though, this will result in messed up geometry. If the vertices you're merging aren't properly aligned to a plane, here's what you can do to make sure you merge properly...

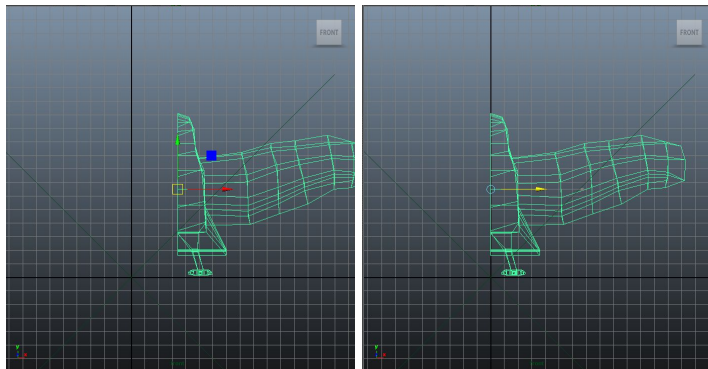
1. Go into vertex mode, select the outlier vertex, then hold V and use the axis mover to mover of the move tool to align that vertex with the rest of the vertices.



- Go into object mode, hold V and D and move the pivot point to the edge being mirrored.



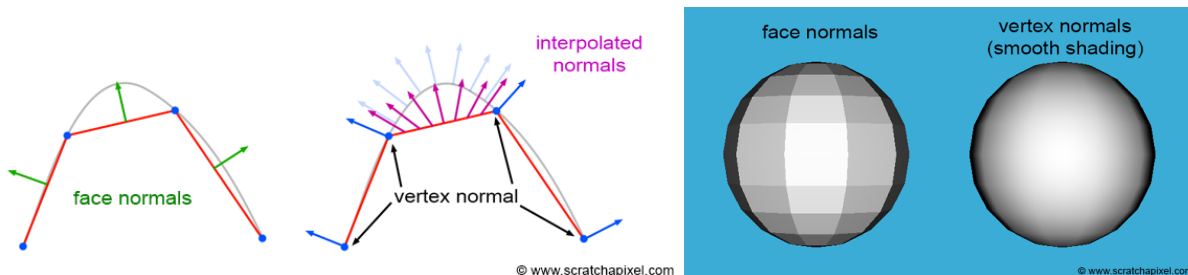
- Hold X and use the axis mover to move the object right up to the plane



Polygon Normals


Polygon normals are vectors that determine how light reflects off of a surface. Normals are typically perpendicular to the surface that they're reflecting light off of.

There are 2 types of normals: face normals and vertex normals. A face normal is a single normal that's applied to all the points on a face, while vertex normals are interpolated between the vertices that make up a face. The end result is that face normals make each face look flat while vertex normals (depending on how you orient them) allow you to have a smoother look, a harder look, or a flat look.



NOTE: In Maya, it looks like vertex normals are NOT shared between the faces they're attached to -- meaning that the a vertex has 1 normal for each face that it's a part of. As such, face normals are probably just a special case of vertex normals: setting the face normal just orients all the vertex normals for that face the same way.


For example, here's a picture showing the vertex normals of a cube that has flat shading. Note that there are 3 normals for each vertex (each vertex in a cube attaches to 3 different faces)...



NOTE: You went through the math for all of this during school. Go here for a refresher on the various concepts:
<https://www.scratchapixel.com/lessons/3d-basic-rendering/introduction-to-shading/shading-normals>

Face Normals


To turn on face normals, go to Display -> Polygon -> Face Normals.



NOTE: Face normals are very likely just the average of the vertex normals. When you manipulate the face normals, it'll manipulate all the vertex normals the same way (probably).


Vertex Normals

To turn on vertex normals, go to Display -> Polygon -> Vertex Normals.



Remember that each vertex has a normal for each face that it's attached to. A vertex normal is interpolated between the other vertex normals of that face in order to give smooth out the look of the face.

Compare the following 2 spheres...



In the first picture, each vertex normal for a face is oriented such that it's normal (perpendicular) to the face. In the second picture, each vertex normal is a blend (average?) of the normals of the faces it's attached to. As such, the second picture looks smoother while in the first picture each surface is hard/flat.

If you want to get a hard look (1st picture), select object and go to Normals -> Harden Edge

If you want to get a soft look (2nd picture), select object and go to Normals -> Soften Edge


NOTE: You can apply Harden/Soften Edge to vertices, edges, faces, or objects.

NOTE: The history nodes for these (polySoftEdge#) will give you an angle option which you can use to control how hard/soft you want it.

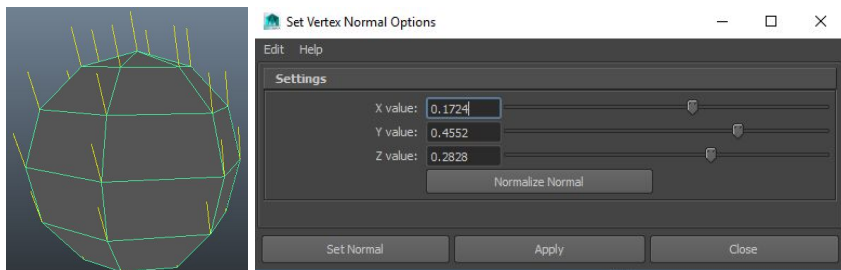
Manually Editing Normals

You can manually edit normals by going to Normals -> Vertex Normal Edit Tool

The tool will give you a rotation tool prompt and you can use it to adjust the vertex normal for vertices, edges, faces, or objects. In the following picture I selected the vertex normals for multiple vertices and rotated them.



If you want to change all the vertex normals to some specific value, you can do that by going to Normals -> Set Vertex Normals (click on the box on the right-hand of the menu to open the options).




NOTE: When you manipulate a vertex normals like this, you're locking the normal in to some position. To unlock the normals, go to Normals -> Unlock Normals.

Backfacing

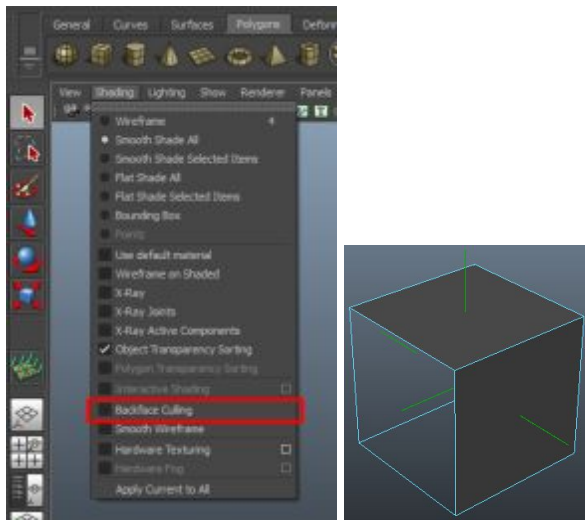
A normal is either facing forward or is backfacing. In the cube example above, the normals are all facing forward. If you wanted to, you could select a face and reverse the normal.

To reverse the normal, select the face and go to Normals -> Reverse.




Note that the face on the left has had its normal reversed, and there's now a difference how that face is reflecting light vs the original (subtle).

A good way to find out if a surface is backfacing without having to show normals is to enable backface culling in the viewport. To do this, you can go to the viewport's menu and choose Shading -> Backface Culling. Once you do this, if the normal for a surface is backfacing, it'll be transparent (not shaded).



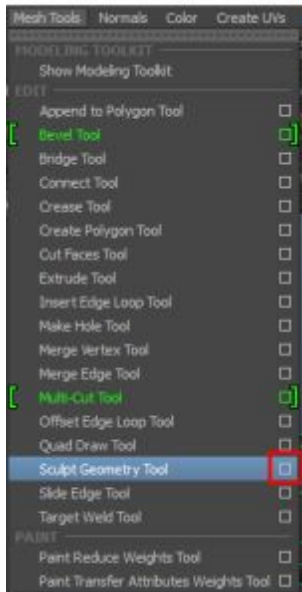
Another good way to find out if a surface is backfacing is by turning off two-sided lighting in the viewport. To do this, you can go to the viewport's menu and go to Lighting -> Two Sided Lighting to turn it off. Once you do this, if the normal for a surface is backfacing, it'll be black (no lighting applied to it).



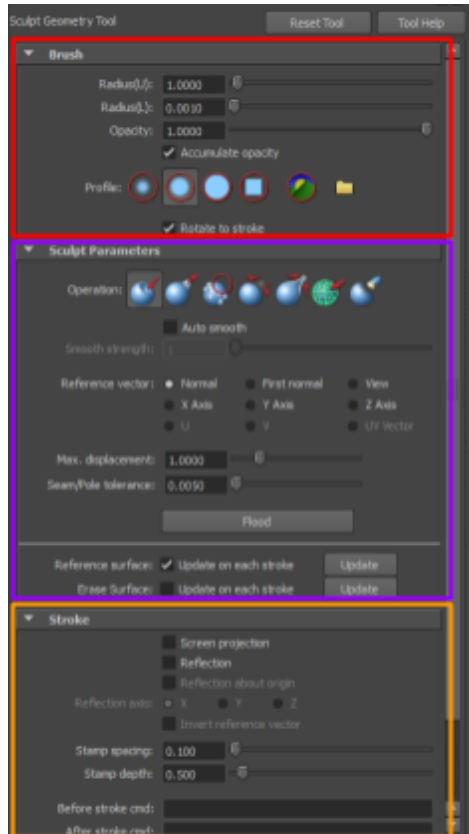
Polygon Sculpt Geometry

The polygon sculpt geometry tool is a tool that'll let you manipulate geometry with a brush. The brush is something analogous to an inkscape brush. You get a brush size and brush opacity setting. The size effects how wide an array you can spray and the opacity controls how much of the effect is applied.

To open the polygon sculpting tool, go to Mesh Tools -> Sculpt Geometry Tool (make sure to click the little box on the right hand side so you open up the tool settings).



In the tool settings that open up, the **Brush** and **Stroke** options apply to pretty much all functions of the tool, while the **Sculpt Parameters** define what function is actually performed.



For brush options...

- Radius(U) is the brush size -- typically this is set in the viewport by holding B + dragging the LMB.
- Radius(L) is the minimum size the brush can be when you resize in the viewport.
- Opacity is how strongly the selected function the brush is applied.

For sculpt parameter options...

- Reference vector is the vector in which the sculpt tool is operating against
- Max displacement is the maximum amount of change the sculpt tool will make per application (per application = as long as the mouse button is held down)

NOTE: These options are a bit confusing. As an example, if you're using the push/pull tool with the reference vector of X-axis, it'll only push/pull in the X-axis directions. The units talked about in max displacement are in relation to the grid in your viewport.

For stroke options...

- Reflection lets you mirror whatever you do on one side to the other


- Invert reference vector is used to define how the effect of the brush is reflected

These options will be described further below in the tools they're important for.

Push/Pull Tool




The push and pull tools let you essentially do the same thing. You're either pushing out or pulling in faces. If you select one, you can temporarily switch the functionality to the other by holding down the Ctrl key.




The max displacement setting describes the maximum amount you want the face to go down/up, while the opacity setting describes how much of that max displacement is applied.

For example, this is one stroke with a max displacement of 1 and a opacity of 1.



Note that it only pushed down 1 unit. No matter how many times I go over this groove with the mouse button pressed, it'll only go down to the max displacement (1 unit). If I release the mouse button and do another stroke, then it'll go down another unit.

Here's another example. This time the max displacement is 1 and the opacity is 0.1.



Note that this also only pushed down a max of 1 unit, but it's a more boat-like gouge. It was one continuous stroke, but I went over the outside lightly and went over the inside more heavily.

NOTE: By default this pushes/pulls in relation to the normal vector of the face. If you want to push/pull in a certain direction, try setting the reference vector. However, if you have reflection turned on, you might need to use invert reference vector to get it things to be applied symmetrically.

Smooth/Relax Tool




The smooth and relax tool do similar things.

The smooth tool averages the vertices that you're sculpting on, in essence smoothing them out. A problem with the smooth tool is that tends to deflate the volume of the object. Sometimes you want to smooth things out but you don't want the volume to heavily change, which is what the relax tool tries to do.

You would use these tools to try to smooth out jagged edges in your object.

In the example below, the left image has had the smooth tool applied to an area of the sphere while the right image has had the relax tool applied. Note that the left image looks deflated, while the right image is mostly retains the shape.



As with the other tools, max displacements and opacity effect how much of the effect is applied.


NOTE: Changing the reference vector to something other than normal also seems to have an effect on how the tool operates but I can't think of a usecase where this would actually be needed.

Pinch Tool



The pinch tool allows you to tighten the vertices in a certain area. Note that the pinch tool tries to maintain the shape of the area that it's tightening (it tries not to distort the overall shape).

In the example below, I used the pinch tool to smush the front of the sphere.




Slide Tool



The slide tool allows you to slide the vertices around in a certain area. The direction you move the mouse is the direction the vertices will slide in. Note that the slide tool tries to maintain the shape of the area that it's tightening (it tries not to distort the overall shape). As such, it's very useful for evening out your geometry.

In the example below, I used the slide tool and continually went in circles to try to create a swirl on the outside of the sphere...




Erase Tool



The erase tool reverses the sculpting changes you made. Why would you use this instead of Ctrl+Z (undo)? Because this lets you target specific areas to undo via the brush.

In the example below, I undid the bottom half of the swirl distortion from the slide tool example...

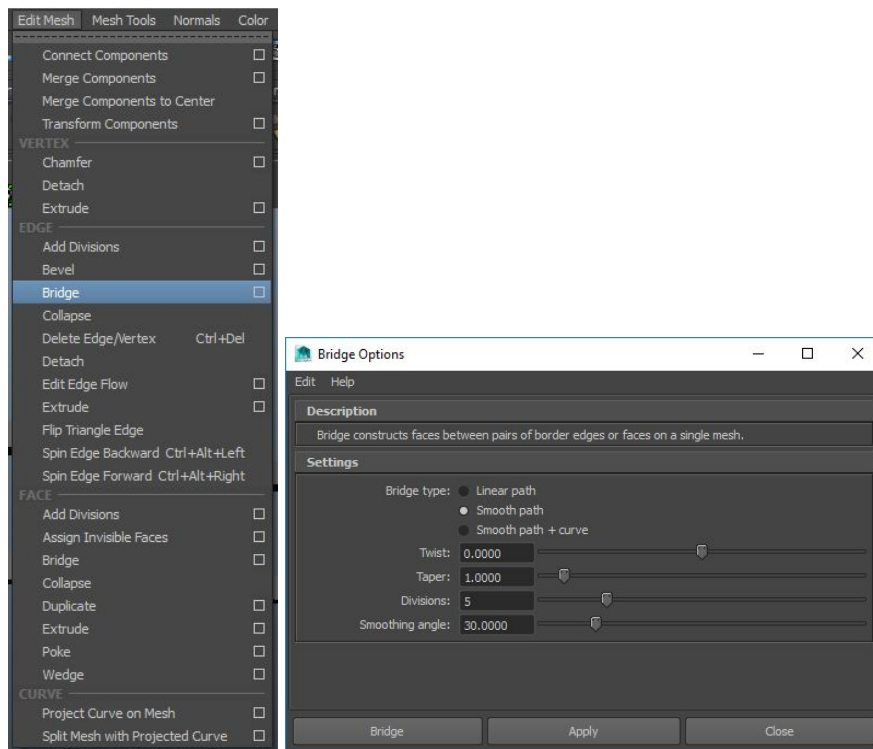


NOTE: This erase tool only works during the particular sculpting session that you're in. So if you navigate away from the sculpting session (e.g. navigate away from the sculpt tool by going to the select tool or to the modelling toolkit or something), you won't be able to use the erase tool to undo the previous sculpting changes you made. The moment you navigate away from the sculpting tool is the moment the erase tool loses its context.


Polygon Bridge

Bridging allows you to connect n open edges from one side of an object to n open edges on another side of the same object. Note that the edges must be open (not connected to a face) and the number of edges being connected must be the same.

To do bridge edges together, first select your edges and then go to Mesh Tools -> Bridge (click the box on the right-hand side of the menu to open up the tool settings).



NOTE: Not all options are exposed in the options dialog. Some options are only available in the attribute editor and channel box. For example... If you want a high number of divisions (more than 25), you need to do it via the channel box / attribute editor.




NOTE: Sometimes when you change the number of divisions, you'll notice that the shading won't show up on the newly created faces. This is a bug in Maya, and it resolves itself as you continue to work. If you need to fix it right away, go into object mode + right click on the object + choose Assign Existing Material -> lambert1 (or whatever material you're using).


NOTE: At least some of this bridging functionality is exposed through the modeling toolkit.

Basic Bridge


1. Create your cube and duplicate it. Make sure you have 3 subdivisions on each side.





2. Combine these 2 objects by going to Mesh -> Combine.



3. Remove the face attached to the edges you want to connect + select those edges.



4. With the edges selected, go to Edit Mesh -> Bridge (click the box on the right-hand side of the menu item to open up the options). Choose a bridge type of linear.

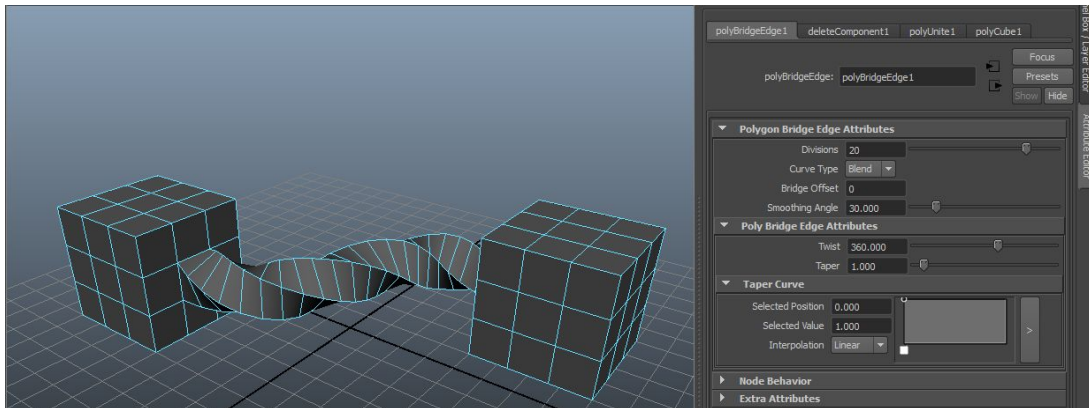


NOTE: In the attribute editor, change the parameters are you see fit. The image below sets the draw type to Blend and the twist to 360. Additional changes you can make include...

1. You can use smoothing angle to adjust the normals (you can always manually set/adjust the normals through other means)

2. You can use taper to bulge/slim the bridge either by setting the taper value or using the taper curve graph.


One thing you can't do is set the draw type to Curve. This is explained more in the curved bridge section.




Curved (Smooth + Curve) Bridge

You can create a bridge that has an underlying curve that shapes the bridge. The caveat here is that you need to specify this when creating the bridge. You can't change to a Curve type after the fact via the channel box / attribute editor (because no curve will exist).


1. Create your cube and duplicate it. Make sure you have 3 subdivisions on each side.




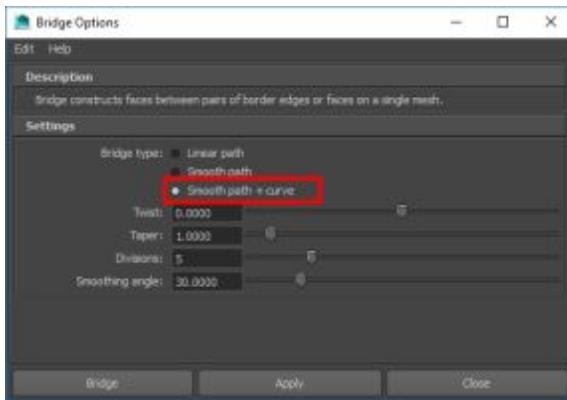
2. Combine these 2 objects by going to Mesh -> Combine.




3. Remove the face attached to the edges you want to connect + select those edges.



4. With the edges selected, go to Edit Mesh -> Bridge (click the box on the right-hand side of the menu item to open up the options). Choose a bridge type of smooth path + curve.




If you switch to wireframe view, you'll now notice that there's a curve running through the bridge. You can adjust the control vertices on this curve to curve the bridge to your liking.

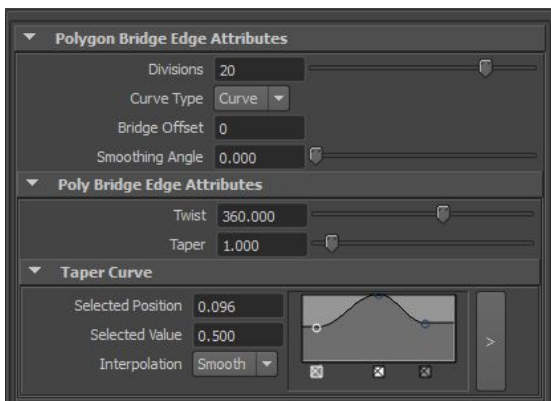




5. Just for kicks, I messed with the attributes here a bit here as well.

I increased the divisions to 20 and then curved the bridge up.



I then twisted the bridge and I used the taper to try to get rid of that pinch at the top. I also smoothed the normals but hardened the normals along the horizontal edges.




Blended (Smooth) Bridge


This is pretty much the same thing as a curved (smooth + curve) bridge, but without the underlying curve that lets you manipulate the bridge after the fact. You can think of it the same way except that the underlying curve gets removed after the bridge is created.

Unlike the curve bridge, you CAN switch to a blended bridge after the fact via the channel box / attribute editor.


1. Create your cube and duplicate it. Make sure you have 3 subdivisions on each side.




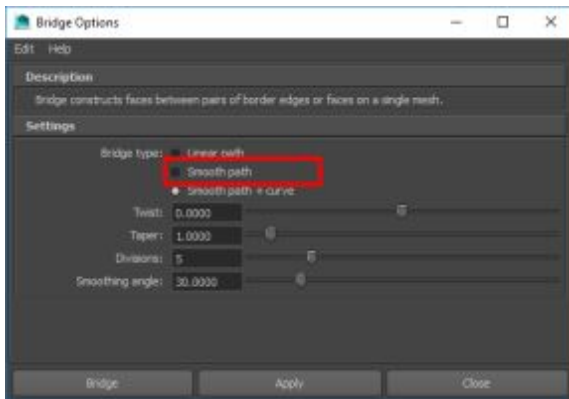
2. Combine these 2 objects by going to Mesh -> Combine.



3. Remove the face attached to the edges you want to connect + select those edges.



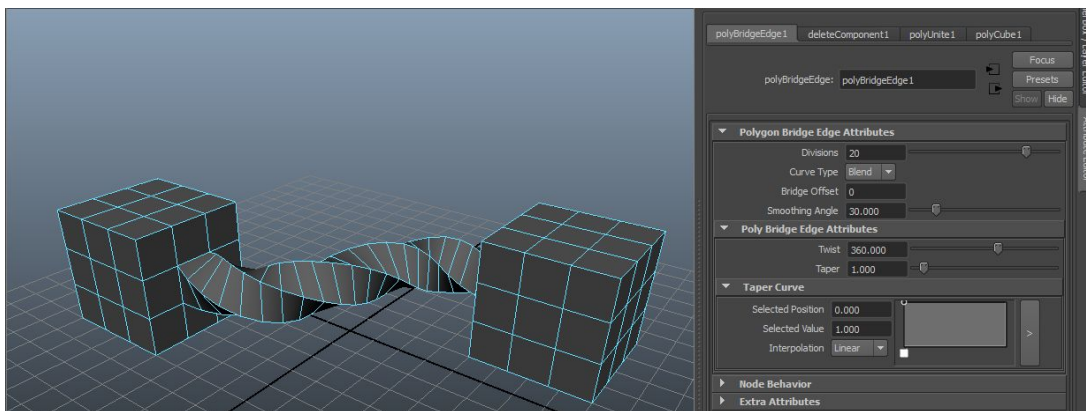
4. With the edges selected, go to Edit Mesh -> Bridge.



NOTE: In the attribute editor, change the parameters are you see fit. I set the twist to 360. Additional changes you can make include...


1. You can use smoothing angle to adjust the normals (you can always manually set/adjust the normals through other means)
2. You can use taper to bulge/slim the bridge either by setting the taper value or using the taper curve graph.

One thing you can't do is set the draw type to Curve. This is explained more in the curved bridge section.

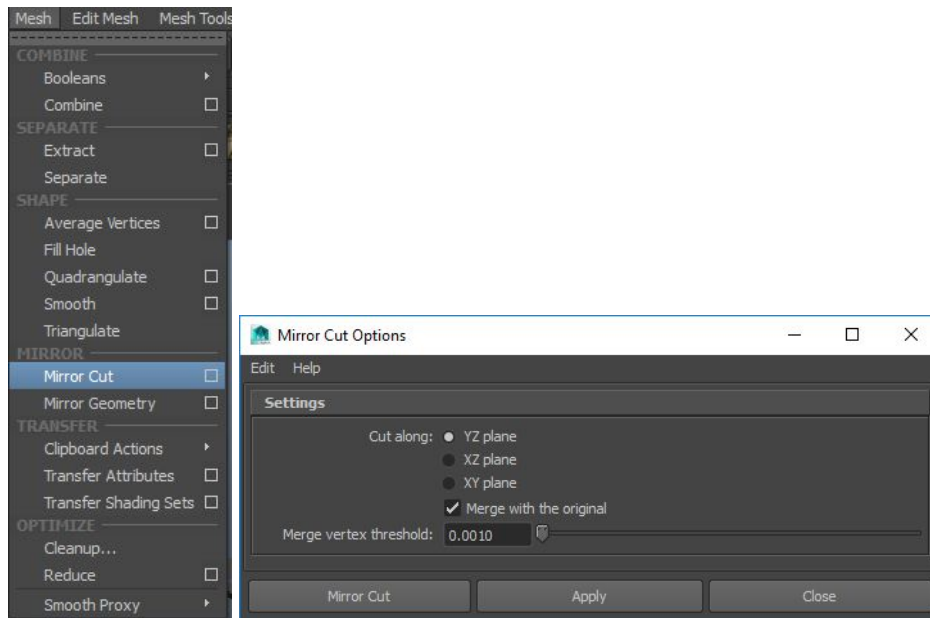


Polygon Mirror Cut

The mirror cut tool allows you to mirror one side of your object across one of the planes that the basis vectors make up (XY, XZ, or YZ). You can use the mirror cut tool to easily extend a repeatable pattern across different orientations. For example, you can use it to extend a model of a fence / give it a 90 degree angle.




To use the mirror cut tool, first select your object then go to Mesh -> Mirror Cut and select the box on the right-hand side to open up the tool options.





Once you cut, an invisible surface will show up on the axis you specified. This plane is used to cut and/or mirror your geometry, and you can orient it different ways to get different effects.

In the example below, we're mirroring the right-side of the object to the left side.





If we wanted to mirror the other way around (from the left-side of the object to the right-side), just rotate the plane 180 degrees.





NOTE: If you're wondering why portions are showing up as green, I have zero clue why this is happening. This is how things look in viewport 2.0. In the legacy viewport, things seem to not get shaded at all (although in the legacy viewport this seems to be a temporary thing that eventually resolves itself). To fix this, go into object mode and right-click on the object and choose Apply existing material -> lambert1.

You can orient this image plane however you like. For example, here's a weird rotation setting.




Or, you can even extend the plane past bounds of your object.




Once you're done with your mirror cutting, you need to delete the history so it can stay this way. To do this, with the object selected go to Edit -> Delete by type -> History.

You'll notice that the mirror plane is still there (transforming will no longer have any effect). That's because it wasn't part of the node history for that object (it's a separate object of its own). To get rid of it, open up the Outliner (or hypergraph) and delete the mirrorCutPlane# and mirroredCutMesh# items.

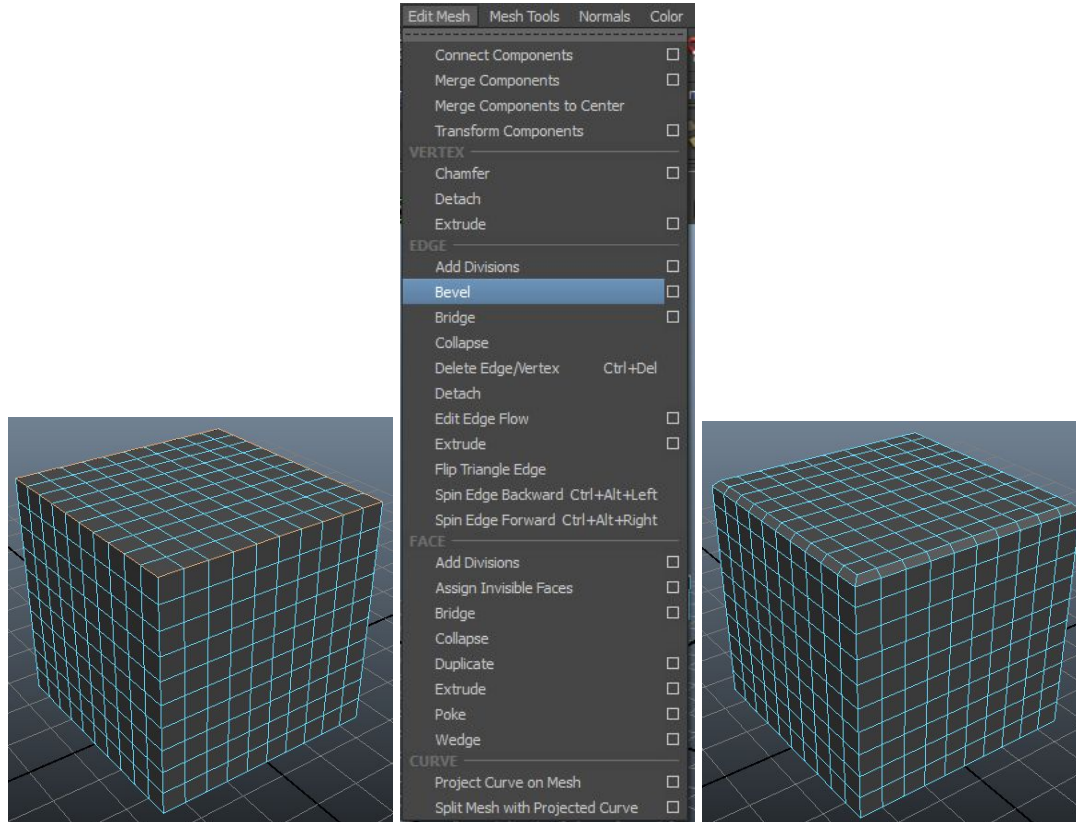


Polygon Bevel

A bevelled edge refers to an edge on a structure that isn't perpendicular to the faces of the faces of that piece.




To create a bevelled edge, first select the edge(s) and then go to Edit Mesh -> Edge -> Bevel. Modifications are best made based in the channel box / attribute editor, so you don't need to click the box on the right-hand side to open up the tool options.



Useful options that control how the bevel looks can be found in the channel box / attribute editor.


You can increase the number of segments that make up the bevel by changing the segments property. In the example below, segments is set to 5.



If you zoom in and look close at the bevel, you can see that it's rounded.


Trying to change the roundness property does nothing. That's because the autofit property is set to on. Turn off the autofit property and you can now set the roundness property. Setting a roundness > 0 will cause the bevel to protrude outward, while setting a roundness of < 0 will cause it to sink inward.

The left picture is a positive roundness while the right is a negative roundness (if the pics below are too small for it to be noticeable, zoom in)



The fraction property controls how large the bevel is. If you use the slider, it'll only extend to 1, which will take up half of each polygon. You can manually set this to 2 to take up the entire polygon. Anything more than that and the object ends up distorting.


The left picture is a fraction of 0.1 while the right is a fraction of 2.



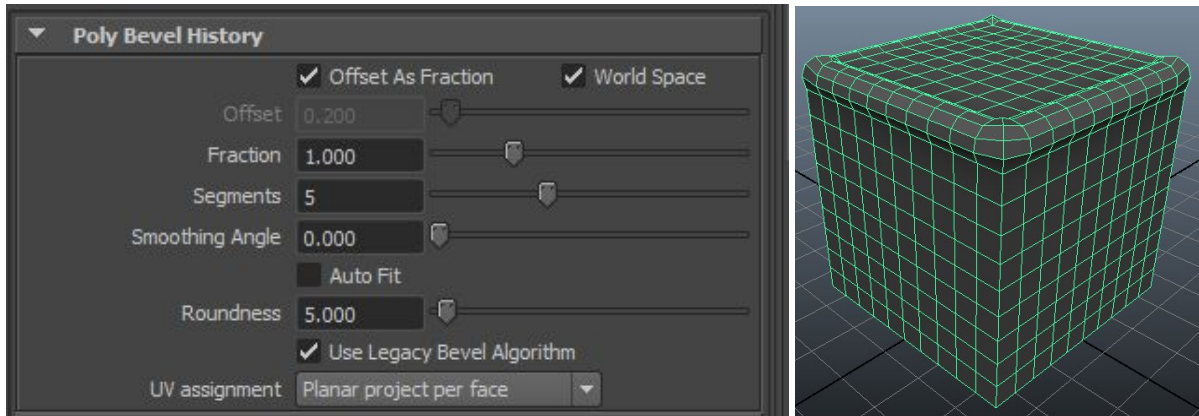
NOTE: It looks like if you set a fraction greater than 1 and a roundness > 0 , sharp corners may end up bulging out slightly.

The smoothing angle property controls the normals. Set to 0 for hard edges, set to 180 for soft edges. Or, control the normals manually via hand-editing or the normals menu. Either way works.

In the example below, the smoothing angle has been set to 0 so the faces of the bevel look hard.



If you want the bevel to protrude out, kind of like a pipe. You can go to the attribute editor, choose the use legacy algorithm option, and manually set the roundness to something really high. This causes the bevel to extend out of the object.

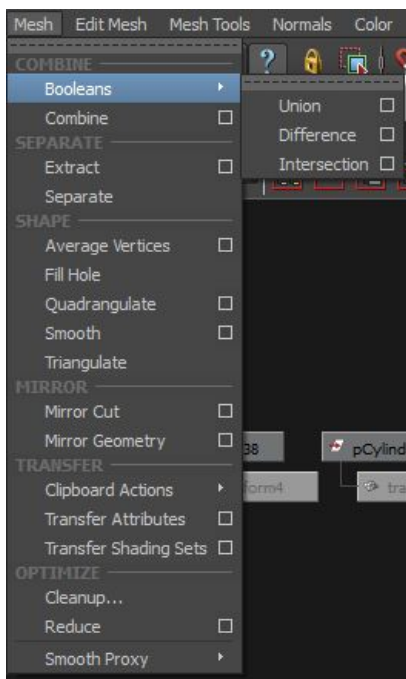


Polygon Booleans

Boolean allow you take 2 objects and ...

- add (combine) them together
- intersect them such that you get a new object that only occupies the space where they collide
- subtract (difference) them such that one object's volume gets carved out from the other


To access these options, go to Mesh -> Combine -> Booleans -> (option)




Add / Union

Union combines one or more objects together. For example, this cube and this cylinder are separate objects, but they become the same object once combined (notice how the intersecting portions have merged).

The left image is before the add, the right is after. The 2 shapes become 1.




NOTE: Fun fact... if you go to the hypergraph and view the connections for the new object, you can still see and access the previous objects the new object was made of. If you select one of the shape nodes for the previous objects, you can transform it (move it/scale it/rotate it) and the new shape will pick up the changes.




NOTE: I'm unsure how doing this operation (Mesh -> Combine -> Booleans -> Union) is different from doing Mesh -> Combine -> Combine.

Difference / Subtract


Difference subtracts the intersecting portion of one object from another. The order in which objects are selected is important here. The first selected object is the object that's subtracted from. For example, consider the following two objects.



If I select the cube first and the cylinder second, here's what the difference would look like...




If I select the cylinder first and the cube second, here's what the difference would look like...



Intersect

Intersection just keeps the part where all objects meet. Unlike difference, the selection order for intersection doesn't matter. For example, this is the intersection of the following objects.



Polygon Quad Draw (cleanup geometry / scan data)

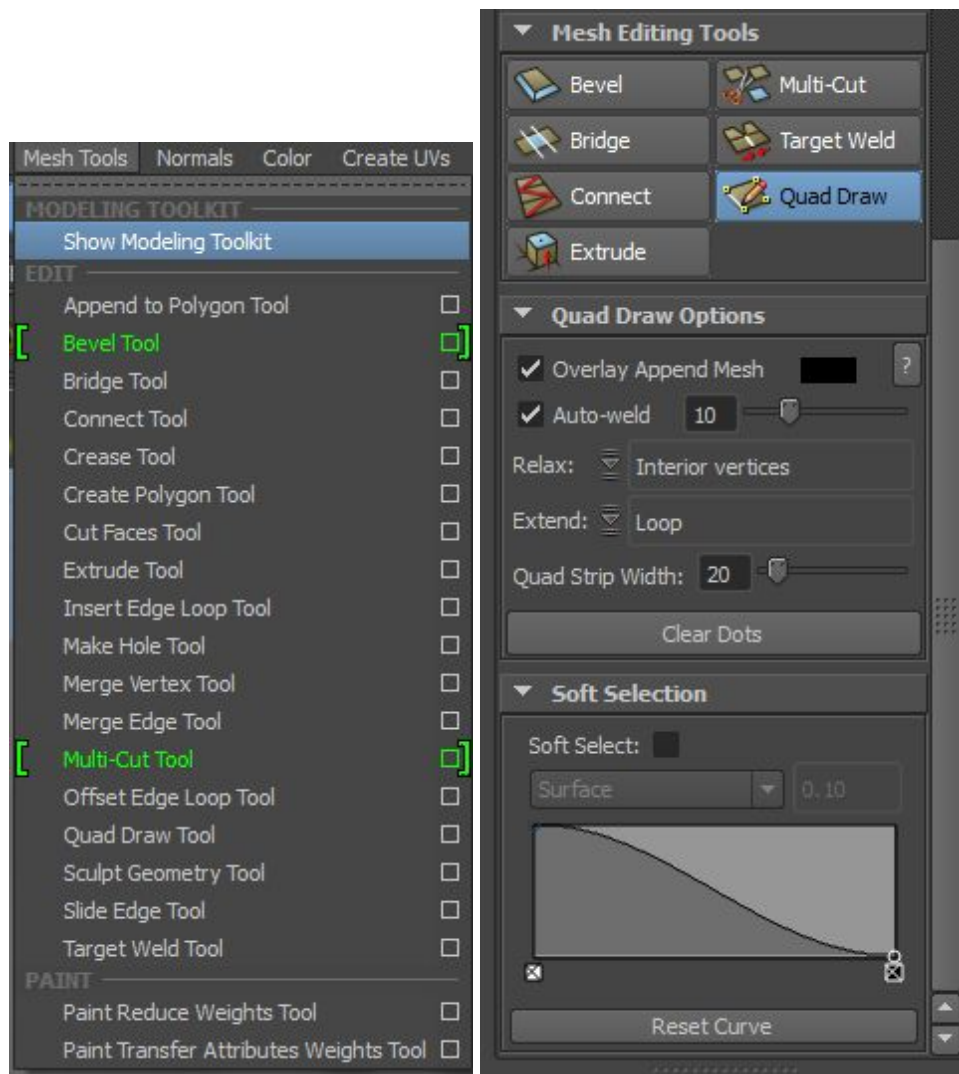
The quad draw tool allows you to draw new geometry on top of old geometry. It's used a lot to re-do geometry or cleanup bad scan data.

To use the quad draw tool, you first need to select an object as the live surface. To do this, select the object and click on the magnet. You should see the name of the object you selected show up next to the magnet.



NOTE: If you go to wireframe, you'll notice that an object that's set to live surface will be represented in green.


Then, go to Mesh Tools -> Show Modeling Toolkit. The quad draw tool should be available as a button in the panel that opens up.




HINT: Handy shortcuts ...

- Context menu = ctrl + shift + rmb
- Move = lmb/mmb
- Split = ctrl + lmb
- Extend = tab + lmb drag
- Delete = ctrl + shift + lmb
- Soft select = b (hold b + drag to change selection size)


Once you do that, you can click LMB on your object to create points, then shift+LMB to create a polygon based on those points.



If you're unhappy with where the vertices/quads are sitting, you can click+hold MMB to move them around (LMB seems to work as well).



If you want to delete an edge loop/border or a face, hold ctrl+shift and click LMB to delete.



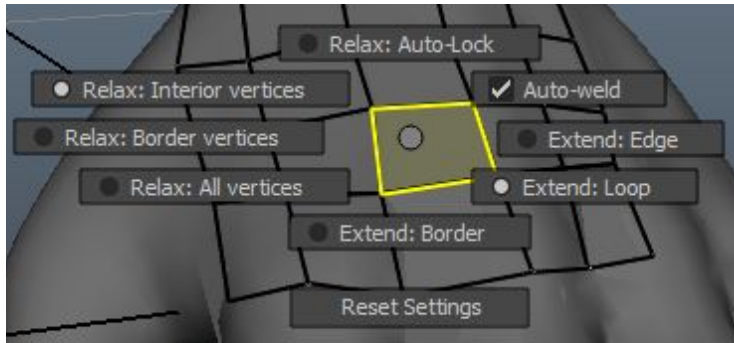
Extending Polygons

If you want to create a new polygons from previous polygons, you can use the extend feature.


Extensions come in 3 types... single or loop or border...

- Single will only extend a single polygon by an edge.
- Loop will extend the entire row/column of edges.
- Border will extend the border edges (those not going to another polygon)


Choose between the three by opening up the marking menu. Open the marking menu by holding Ctrl+Shift and press+hold the RMB. Then choose Extend: ????.



Once you're done that, you can hold Tab then LMB+drag an edge to extend it.




If you want to split your row/column of polygons in 2, you can hold down ctrl and click LMB on where you want the split to be.

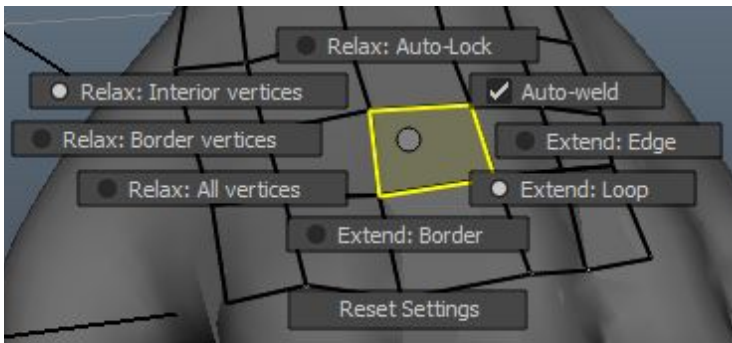


Relaxing Polygons


Sometimes the polygons you create may be jacked up. You can relax your polygons by holding shift and click+hold LMB while you go over them with your mouse.



By default the tool only relaxes interior vertices, but you can change to border vertices or all vertices through the marking menu. Open the marking menu by holding Ctrl+Shift and press+hold the RMB. Then choose Relax: ???? vertices.




Here is a version with the borders relaxed as well...



Soft Select


Sometimes you don't want to just work on one edge/face/dot, you want the surrounding components to be affected as well. We can do this by using the soft select tool.

To enable soft select, either press b or enable it in the panel.

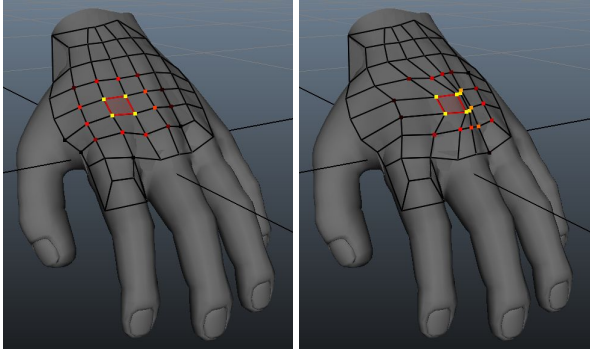


NOTE: Surface selection seems to be the one you want.

Now when you go to make a selection, you'll see a heat map. The brighter the dot, the more it'll be impacted by whatever happens.



Notice how the surrounding polygons change when I move this face right a bit. The surrounding vertices got dragged a bit long with it.



Maya 2015 Materials

[Materials/Shader Primer](#)

[Lambert Shader](#)

[Blinn Shader](#)

[Anisotropic Shader](#)

[Phong Shader](#)

[UV Primer](#)

[Creating UVs](#)

[Planar Mapping](#)

[Cylindrical Mapping](#)

[Spherical Mapping](#)

[Automatic Mapping](#)

[Modifying UVs](#)

[Cutting Edges/Shells](#)

[Combining Edges/Shells](#)

[Flipping UVs/Shells](#)

[Fixing Distorted UVs](#)

[Saving UVs as Texture \(Snapshotting\)](#)

[Texture onto UVs](#)

[Hypershade](#)

[Creating a Material](#)

[Modifying a Material](#)

[Deleting a Material](#)

Materials/Shader Primer

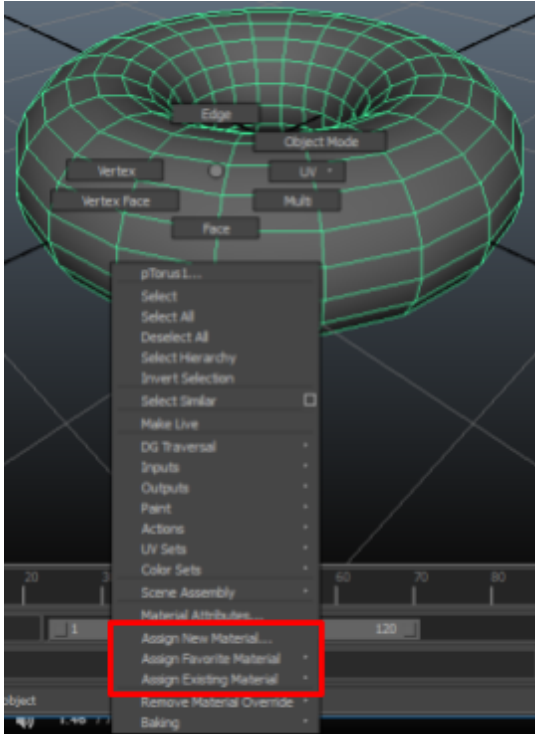
NOTE: This is a poor discussion on materials/shaders. I need to look more into this. It should be good enough as an introduction.

The word materials is synonyms with the word shaders -- they both basically mean the same thing. A shader/material describes the makeup of a surface. It controls how light is absorbed or bounced off of that surface. It also helps define what the surface "is made out of."

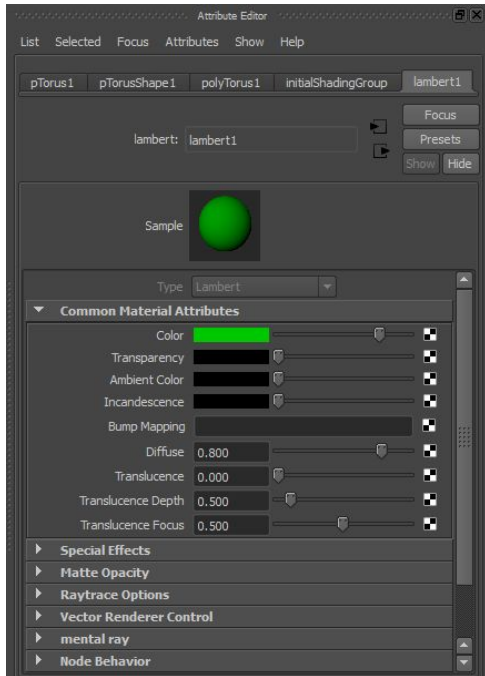
When we talk about what the surface is made out of, we're referring to the real-time/real-world counterpart. For example, if we were trying to replicate concrete, we would want our material to take up concrete-like properties.

NOTE: Without a shader/material, a surface cannot be rendered. If you have a surface without any type of material on it, it'll basically be invisible. You won't be able to see it.

To assign a new material to an object, right-click and either go to Assign New Material / Assign Favourite Material / Assign Existing Material.

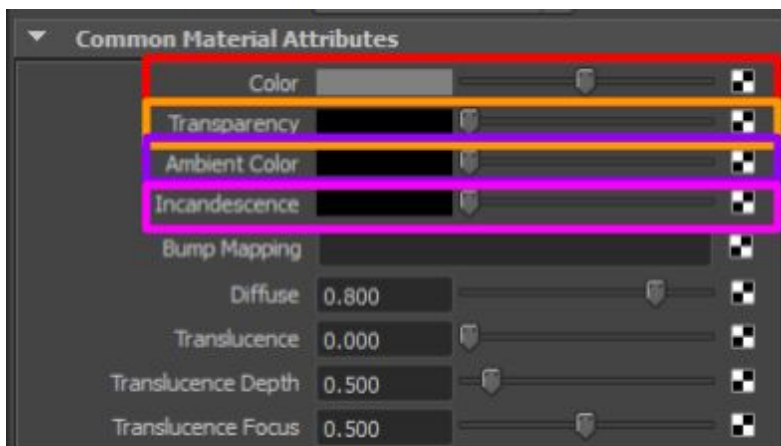


To modify a material's properties, select the object in the attribute editor (Ctrl+A) and navigate to the material tab.



NOTE: When you create an object in Maya, it's automatically assigned a default material. That default material is a lambert shader. Typically you can switch between material types by going to the Type dropdown, but you'll notice that for the default material the Type dropdown is grayed out (see pic above). This means that the default material will always be a lambert shader.


The common types of materials are lambert, blinn, anisotropic and phong. There's a subset of attributes that are common to all of these shaders.



Color

The color of the material.


A torus with a red color.



Transparency

The visibility of the object (how see-through it is).


A torus with a red color and 50% transparency.



Ambient Color

“Set to black by default, which means it does not affect the material’s Color. As the Ambient Color becomes lighter, it affects the material’s Color by lightening it and blending the two colors. If there are ambient lights in the scene, the color and brightness of those lights is used to control how much the ambient color contributes to the final color of the material.”

A torus with a dark red color and a light gray ambient color.




Incandescence

“The color and brightness of light that a material appears to be emitting. (Incandescent objects do not illuminate other objects.) For example, to simulate lava, use a bright red Incandescence. The default color value is 0 (black). Although incandescence makes a surface appear to glow, it does not actually act as a source of light in the scene.”

A torus with a dark red color and a gray incandescence.




Lambert Shader



A Lambert shader only has the common attributes listed above. It does not come with a hot spot/specular highlight (talked about more in the other shaders).

Blinn Shader




Blinn shaders come with a specular highlight / hotspot. This shader is good for replicating metals. You can access the attributes for the specular shading in the Specular Shading rollout portion of the attributes.



Eccentricity

Controls the size of the specular highlight/hotspot.


A torus with a 0.1, 0.5, and 0.9 eccentricity.



Specular Rolloff

Ability of the surface to reflect its surroundings. What's the difference between this and eccentricity? Eccentricity controls the size of hotspot (white area), while specular rolloff controls the amount of light reflected off the hotspot.


A torus with a 0.3 eccentricity + 0.1, 0.5, and 0.9 specular rolloff.



Specular Color

The color of the hot spot.


A torus with a red specular color.



Reflectivity

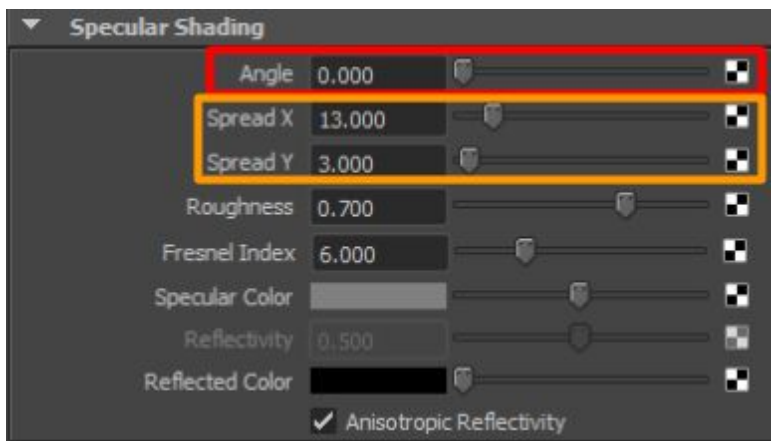
Controls how much stuff reflects off of the material? You won't be able to see the reflectivity in real-time + you need to have other objects in the same scene that the surface can reflect.

Anisotropic Shader



An anisotropic shader is like a blinn shader, except that it allows us to control the specular highlight a little bit differently. Anisotropic shaders are good for simulating lighting effects on surfaces with very small detail -- this causes the specular highlight/hotspot to breakup or otherwise react differently.


Some of the same properties here are available in the blinn shader. Check out the blinn shader section to find out what they do.



Angle

Angle of the specular highlight


A torus with a 0, 45, and 90 angle.



Spread X/Y



How “spread out” the specular highlight / hotspot is.

A torus with a spread of 13/3, 3/13, 1/15.



Phong Shader


An phong shader is like a blinn shader, except that it has fewer attributes and is better suited to replicate plastic materials.



The only option that’s new here is cosine power. The larger the cosine power, the smaller the hotspot / specular highlight becomes.

UV Primer

UV’s are points on a surface that describe a mappable area. They provide a “holdout” area for images and other related material-esque data to be placed. You can think of it as a clean sheet of paper being wrapped around parts of your polygon object.

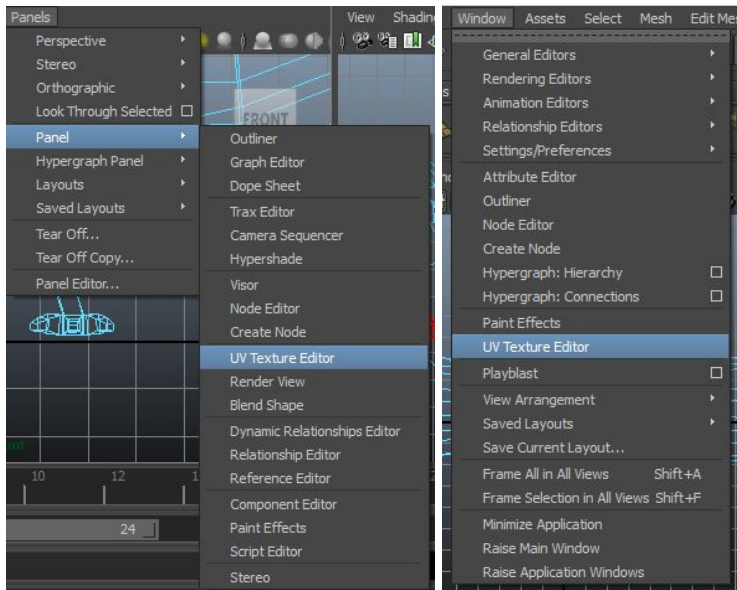


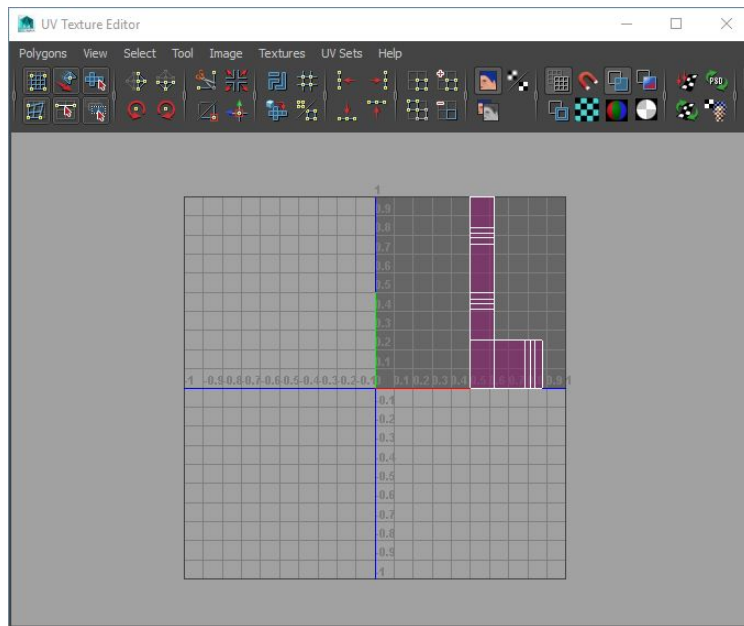
The goal with UVs is to make them human readable (and proportional to the real geometry). That is, when you open your “sheet of paper” in something like photoshop, you should instantly know which regions match to which parts of your polygon object. In our example above, the top and bottom of the cylinder maps to the top and bottom circle (respectively) and the body of the cylinder is in the middle. With more complex objects such as models of animals, the head, arms, legs, etc.. may be broken off into their own identifiable sections.

To create/modify UVs, open the UV Texture Editor. You can do this by ...

1. Changing a viewport panel by going to the menu in the panel and selecting...
Panels -> Panel -> UV Texture Editor

2. Opening a new window by going to...
Window -> UV Texture Editor





NOTE: If you've done something similar to box modeling to create your object, you'll notice that your UVs look nothing like your object (instead they look like the original object you started from). This is essentially junk and you can leave it as-is. As you create new UVs, these old mappings will disappear.

Creating UVs

NOTE: Before doing any of this, make sure you get rid of any history nodes on your object + freeze the transform on your object (make the current transform permanent) + open the UV Texture Editor

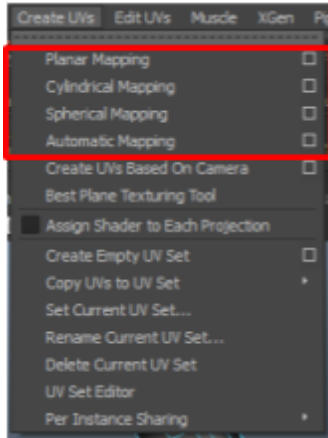
1. To delete history, select the object and go to Edit -> Delete By Type -> History.
2. To freeze transform, select the object and go to Modify -> Freeze Transform.
3. Go to Window -> UV Texture Editor

PROTIP: If the model you're working on is symmetrical, you can create UV mappings only for one side and then use the mirror-cut tool to mirror the model and automatically create UVs for the other side. You'll need to go to the UV Texture Editor and flip/move the shells around, but other than that there's not much else to do.

You can create UVs by selecting an object (or object faces) and selecting one of the options under the Create UVs menu. The type of mapping you choose depends on the the shape of what you're creating UVs for. When you select one of these mapping types, you'll be presented with some geometry around the selection you made.


- Planar Mapping -- when faces are roughly facing some plane (e.g. side of building)

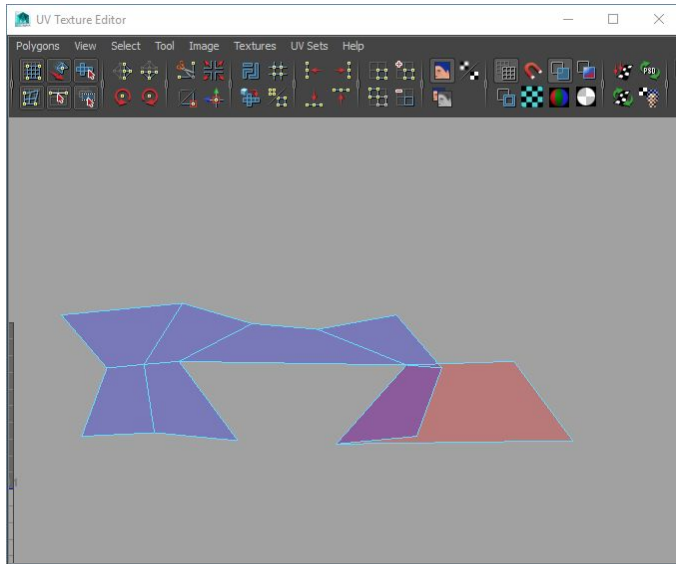
- Cylindrical Mapping -- when faces roughly form a cylinder (e.g. leg)
- Spherical Mapping -- when faces roughly form a sphere (e.g. apple)
- Automatic Mapping -- multiple planar projections, when something has lots of grooves and valleys (e.g. human face?)



NOTE: For best results, option the options instead of just running the tool (click the little box on the right-hand side of the menu)

You'll need to orient that geometry to best suite your selection. You can do this by clicking the T symbol next to the placeholder that shows up (pressing T will expose the manipulator). You'll get a view of how the UVs look in the UV editor pane as you orient.






A lot of times, you'll need to manually fix these UVs by hand. More on this in the "Modifying UVs" section.

Planar Mapping


Planar mapping maps UVs based on a plane. This is good for mapping flat-ish surfaces that are all generally facing the same way. A good example of this may be a bird's wing that's been panned out.

To create a planar map...

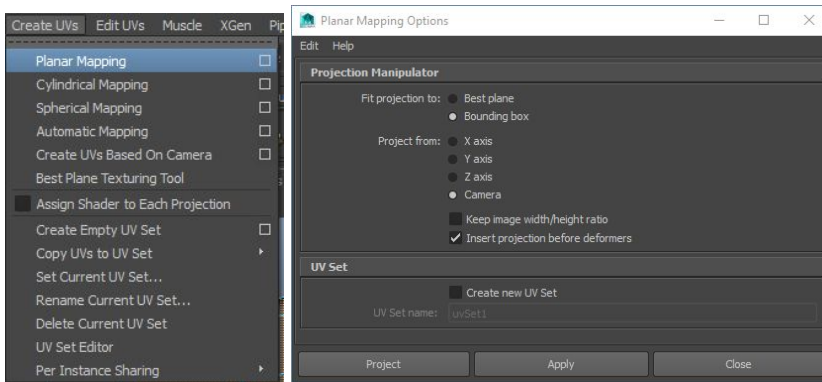
1. Select the faces/object you want to create a mapping for.



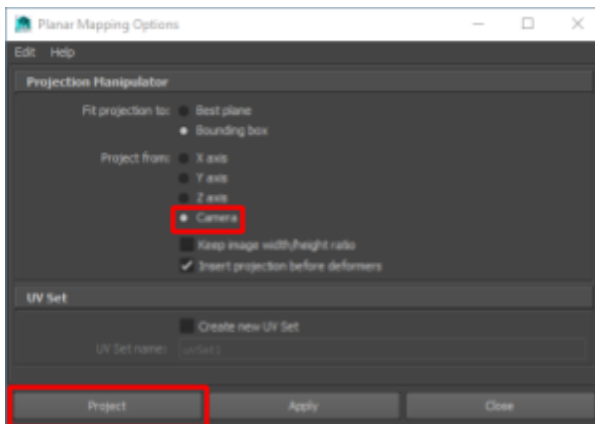
- Orient the perspective camera such that it's looking at most of the faces straight-on.



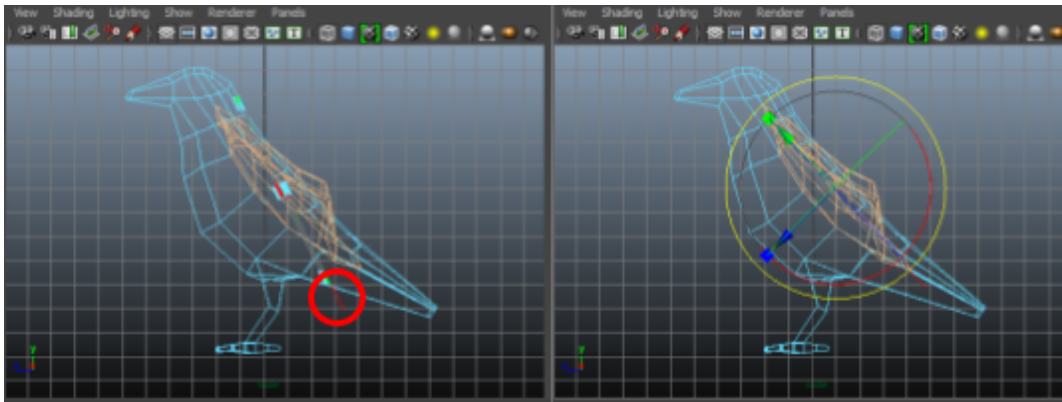
- Go to Create UVs -> Planar Mapping and open the the options (little box on the right).



- In the options, select Camera for the "Project from" option + hit Project.




5. Hit T on the manipulator and orient it as best you can.



As you orient, you'll see the mapping show up in the UV Texture Editor (more complete guide on the UV editor below). One specific thing to note with this example is that the polygons on the wing wrap around, so what happens is that we get UVs that...

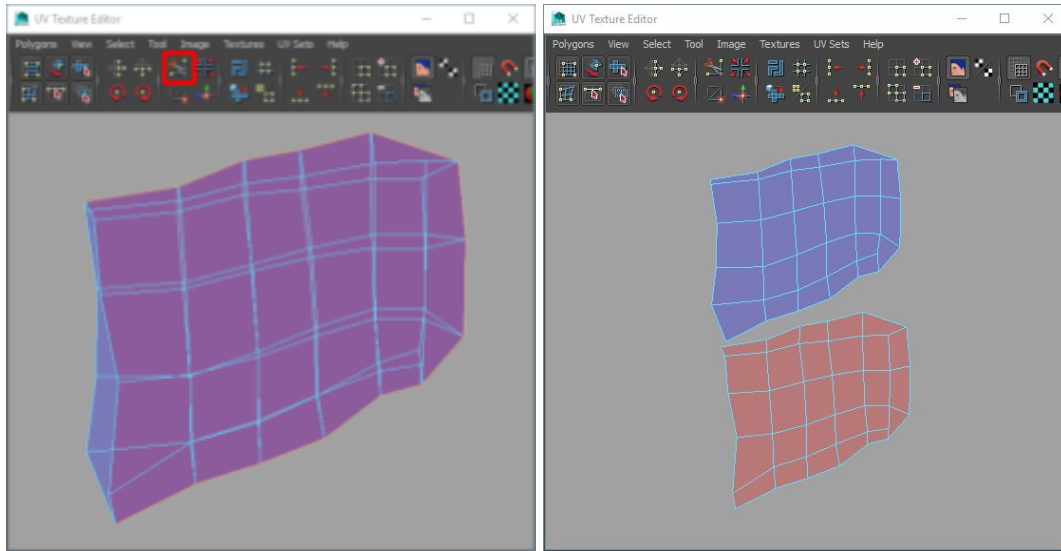
1. overlap each other
2. are flipped around (the top half is flipped)



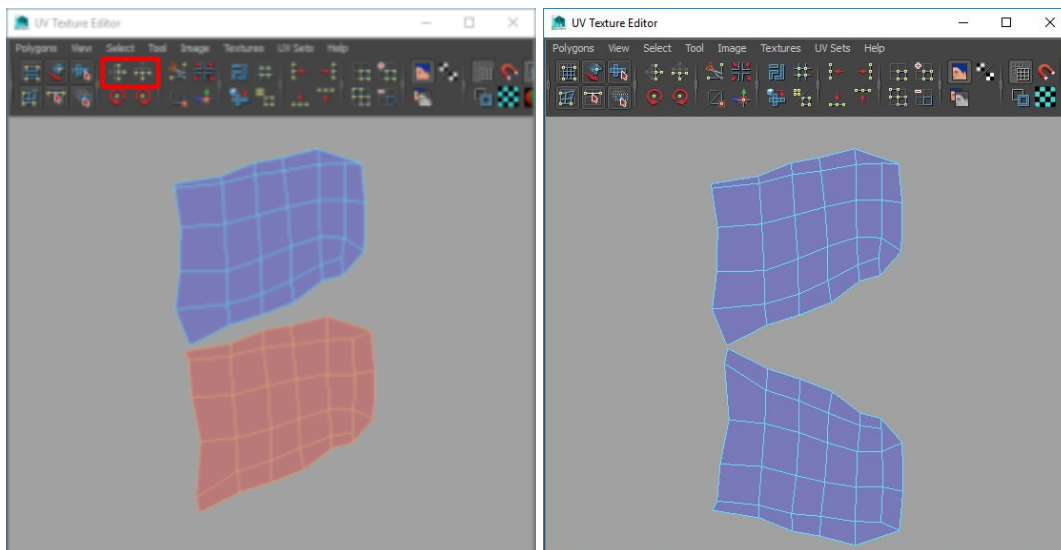
NOTE: This may not always be the case. Sometimes your polygons won't wrap around. For example, The face of a building doesn't wrap around. And, you could have avoided this problem entirely by mapping one side of the wing at a time.

We can fix this in the UV Texture Editor pretty easily...

To fix the overlap, you select the edges that split the top part of the wing from the bottom part of the wing and choose the cut tool. Once cut, simply select the shell and move it such that it no longer overlaps.



To unflip the top-half of the wing (now that it's been split)...



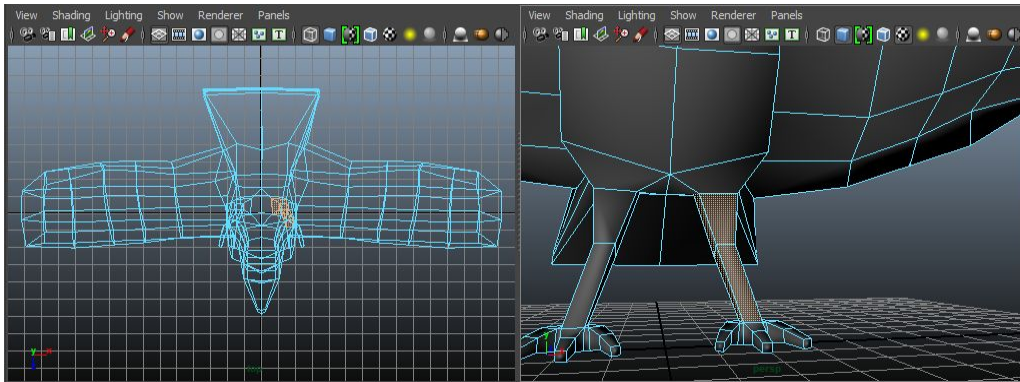
NOTE: There's a bug here with the flipping tools. You need to explicitly select all the UVs instead of just selecting the shell. This is described further in the section on flipping UVs.

Cylindrical Mapping

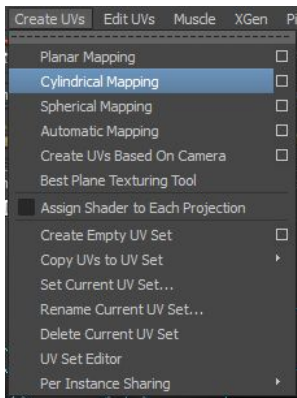
Cylindrical mapping maps UVs based on a cylinder. This is good for mapping parts of meshes that roughly map to a cylinder. A good example of this may be a bird's leg.

To create a cylinder map...

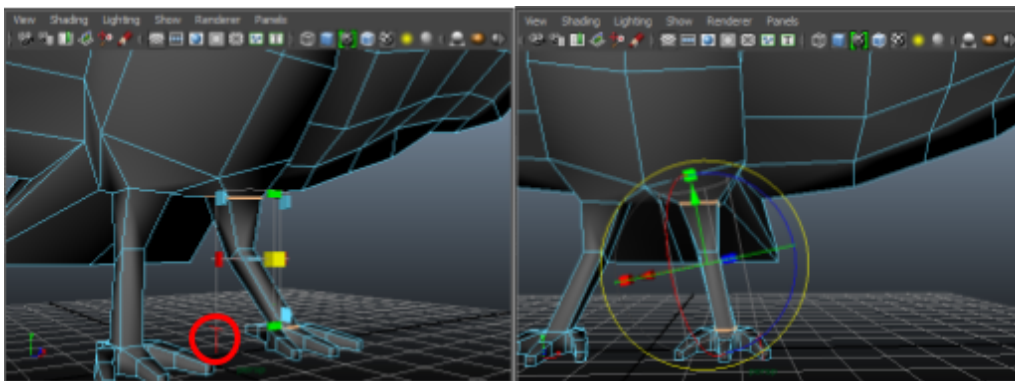
1. Select the faces/object you want to create a mapping for.



2. Go to Create UVs -> Cylindrical Mapping.



3. Hit T on the manipulator and orient it as best you can.




As your orient, you'll see the mapping show up in the UV Texture Editor (more complete guide on the UV editor below). One specific thing to note with the cylinder mapping tool is that a lot of times, no matter how you orient the the cylinder, the UVs will require manual edits to correct.


You can do some things to avoid overly poor mappings (e.g. leave out the top/bottom of the cylinder if there is one), but you'll almost always still require manual edits.

NOTE: Don't try to do two separate pieces with one mapping. Otherwise, you'll end up with UVs where it looks like everything matches up but none of the edges will be bound together.

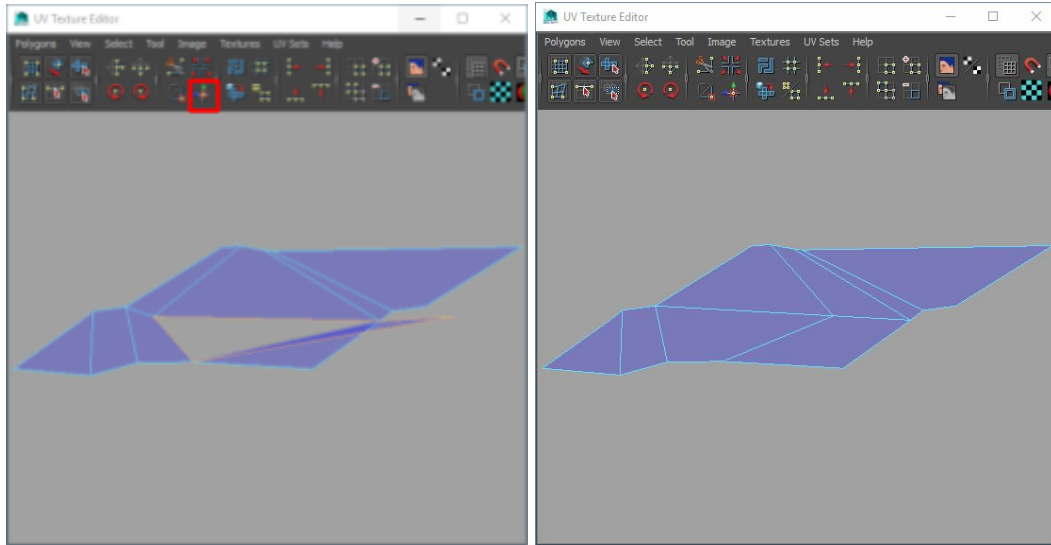
In this example, the UVs actually needs to have some the edges fused to get a proper-ish mapping...



You'll notice that the UVs for one of the polygons is warped out / not properly attached to the rest. You can select these 2 edges and their corresponding edges will automatically highlight.




If fused these 2 edges, you'd get a nice contiguous piece of geometry for the leg that we're mapping. To do this, select the edges and choose the Move and sew tool.



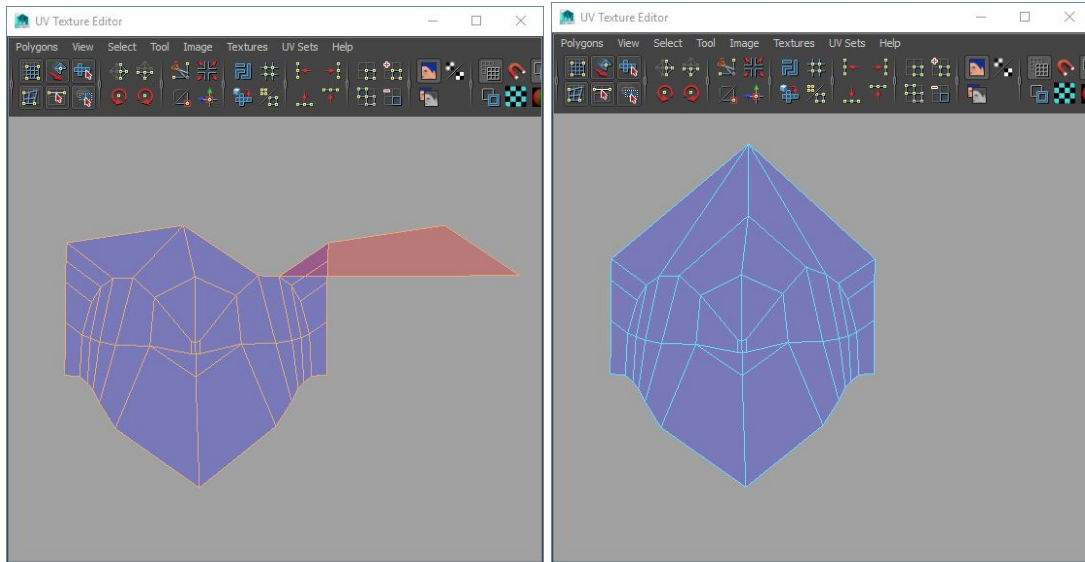
Spherical Mapping

Spherical mapping is similar to cylindrical mapping, but your manipulator is a sphere isn't of a cylinder. You can still do all the same stuff (e.g. press T and transform as needed).



Just like with the cylinder mapping tool, the sphere mapping tool will most likely require manual edges to UVs to correct. Depending on how you orient things, you can minimize the manual effort.

In this example, it's (somewhat) obvious what's out of place / which of the edges have to be fused together to get a nice contiguous mapping for the geometry we tried to UV map.

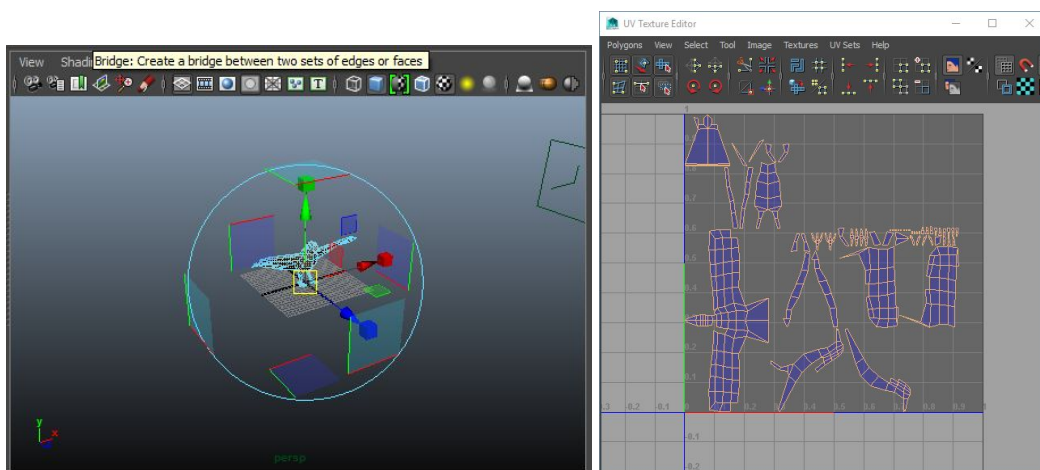


Won't go into much more detail here. See cylindrical mapping section -- most of that applies to spherical mapping as well.

Automatic Mapping

Automatic mapping is similar to planar mapping, but instead of just having 1 plane maya will project many planes to try to get a good mapping. I haven't played with this enough to know how it works, but just by looking at the mappings it generates it looks to do an okay job. Some manual edits will be required.

Note that this mapping option doesn't have a T for you to click on to enable transforming. Transformation is enabled by default.




NOTE: There are tons of settings in the tool options.

Modifying UVs

In the UV Texture Editor, you can modify UVs by hand. You can select which components you want to modify just like you would in a normal viewport + modify them around with the same manipulator tools (move/scale/etc..). The selections you make will show up in your viewports as well.

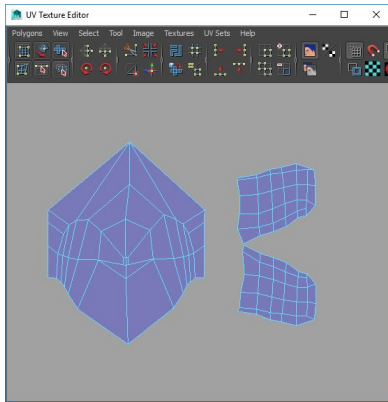
Right-click in the window to open a marking menu and choose the component selection type you want....



- **Shell** -- allows you to select one contiguous UV mapping.

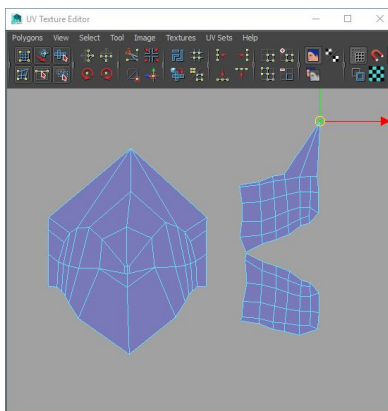
In the example below, we have 3 shells: one for the face, one for the top of the left wing, and one for the bottom of the left wing.

NOTE: There are some issues when selecting by shell and trying to flip. See section of flipping UVs/shells for more information.



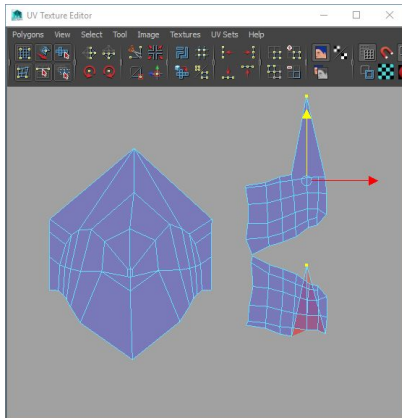
- **UV** -- allows you to manipulate specific UV point(s).

In the example below, we've stretched out one of the UVs for the bottom wing.



- **Vertex** -- allows you to manipulate the UV(s) for one or more vertices.

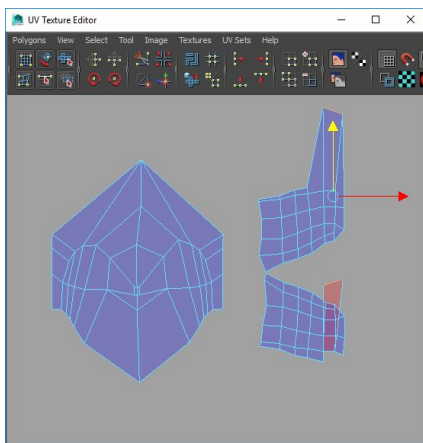
In the example below, we select the same point as we did for the UV and stretch it out. But, note how we automatically get a selection on the shell for the other side of the wing. This is because the vertex is shared between the 2 shells.



NOTE: You could get exactly the same effect by choosing to select by UV (instead of vertex) and manually selecting these 2 points.

- **Edge** -- allows you to manipulate the UV(s) for one edges.

This is similar to selecting by vertex. When you select an edge, the corresponding edge will also get selected (assuming that there's another shell with a corresponding edge). In the example below, we select an edge that's a part of 2 different shells and move it.



NOTE: You could get exactly the same effect by choosing to select by UV (instead of edge) and manually selecting the relevant points.

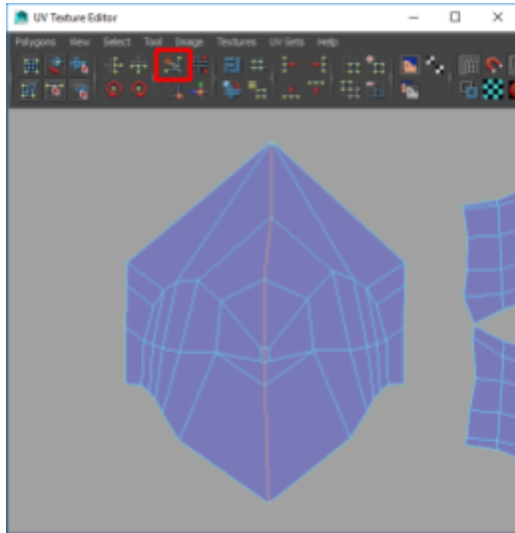
Cutting Edges/Shells

A lot of times you'll come across cases that where you'll need to cut up your shell into smaller pieces. Reasons for this include...


- You'll get distortion if you keep it in one piece
- The UV shell wraps around in 3d space, causing faces to overlap on each other

- Probably other reasons as well?

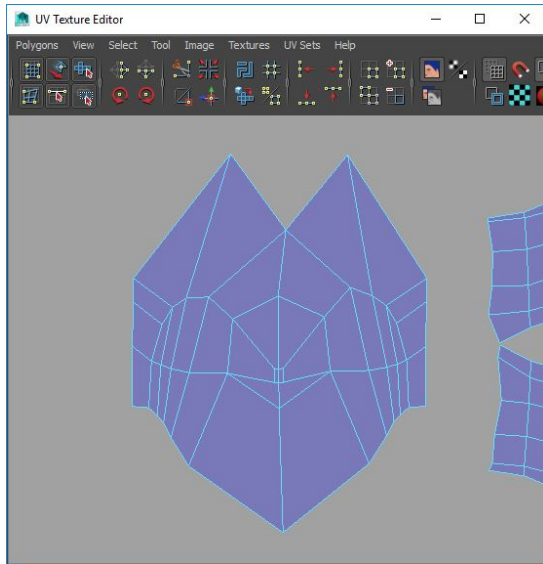
To cut up a contiguous mapping, go into the UV Texture Editor and select the edges you want cut and select the cut by edges tool...



Depending on your edge selection, you might end up with 2 separate shells. In the example above, we do indeed end up with 2 separate shells.



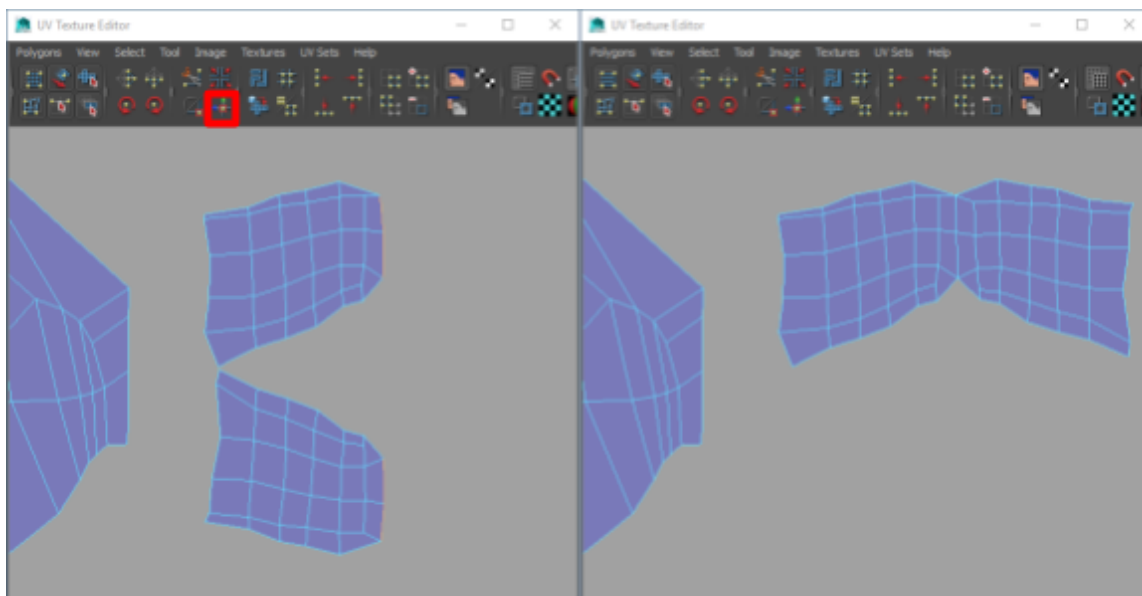
But, if you only selected some of the edges, maybe the topmost edge in our selection, you'd still have 1 shell but you could separate the mapping for those faces from each other.



Combining Edges/Shells

If you want to merge edges together, you can use the move and sew tool. Just like cutting, merging can be done on the same shell or it can be used to combine separate shells.

If you want to merge, first select the edge(s) you want to merge and then select the move and sew tool...

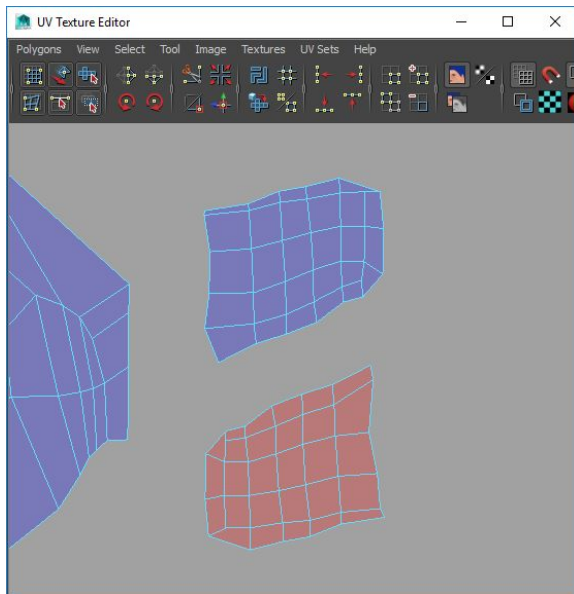


In the example above, we combined the shell of the bottom bird wing to the shell of the top of the bird wing.

Flipping UVs/Shells

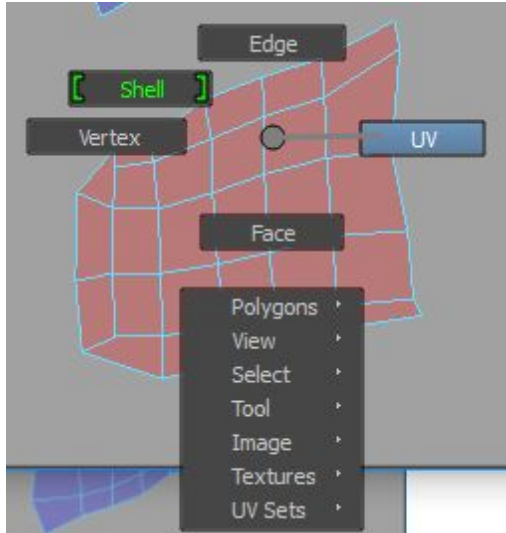
In certain cases, depending on which mapping tool you used / how you used the mapping tool / where you used the mapping tool, some of your UVs may come out flipped. To flip them around, you can use the UV flip tool.

You can tell when UVs aren't oriented correctly when they appear pink in the UV Texture Editor...

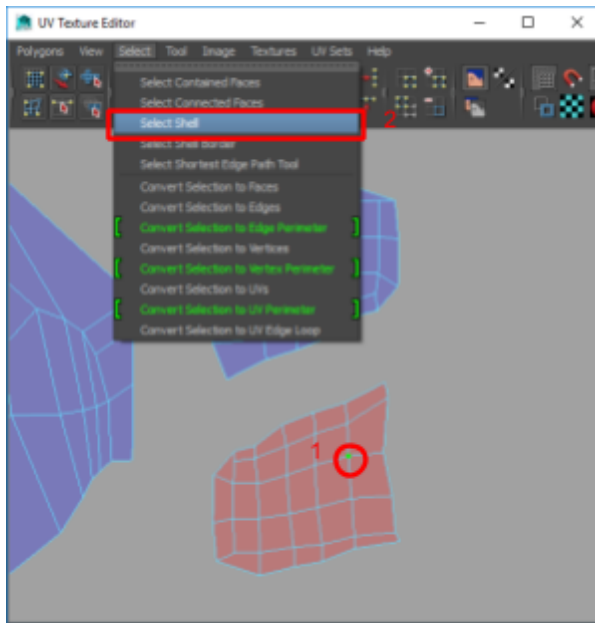


To correct this, you need to flip the shell around with the flip tool. The key thing with this is you cannot do this in shell selection mode. If you do, it'll look like it worked but all the edges will be cut from each other (but still touching).

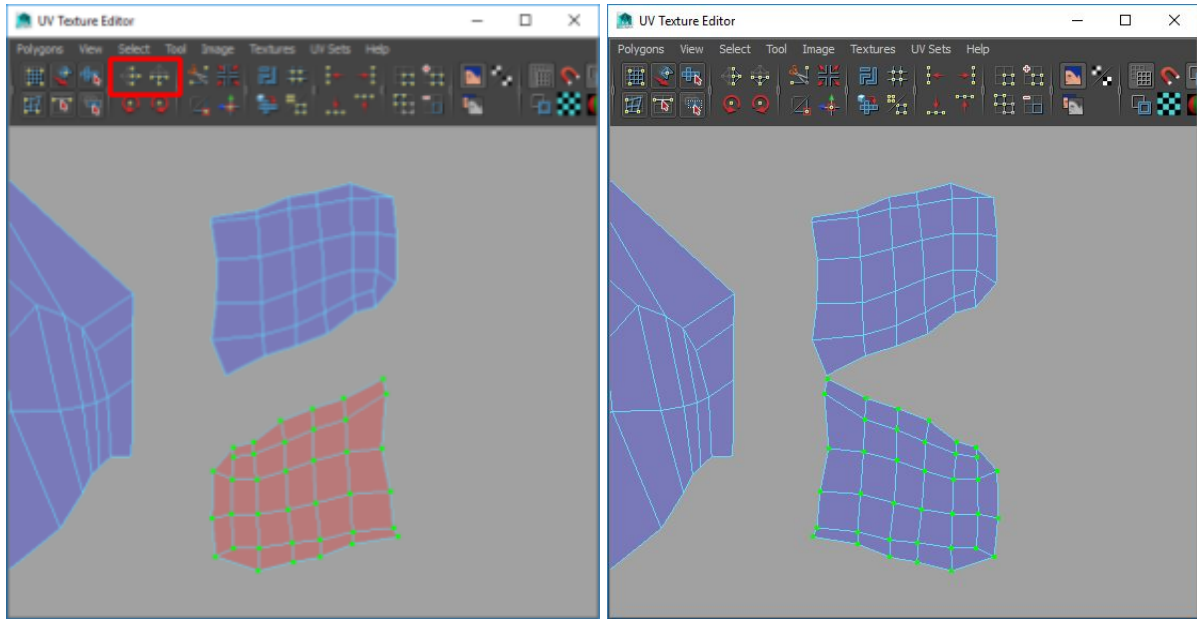
To flip a shell, first go into UV selection mode...



Then select a single UV on the shell you're trying to flip and go to Select -> Select Shell...

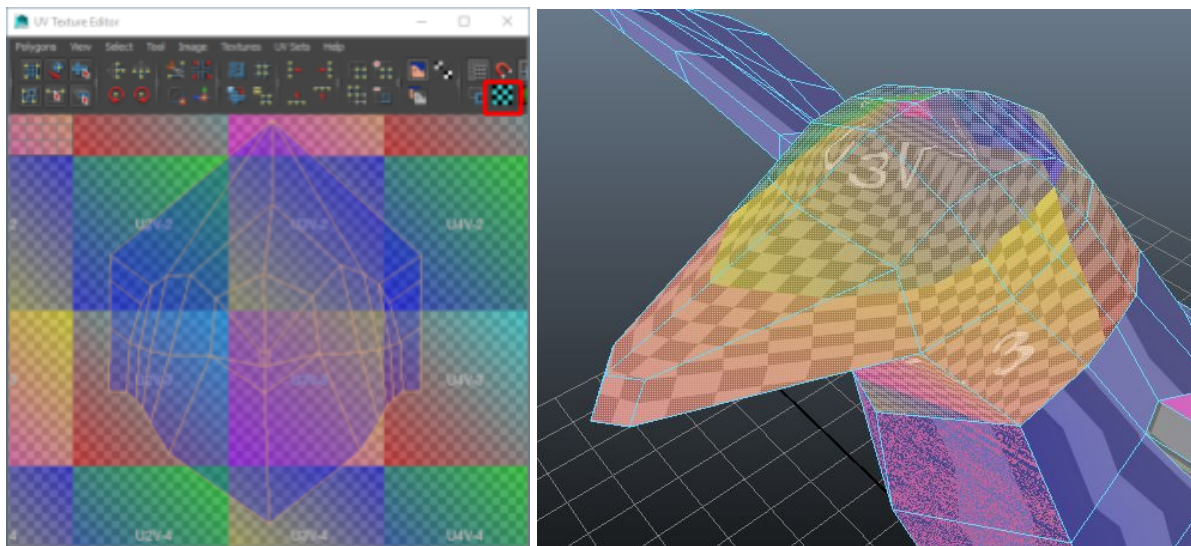


All the UV points in the shell should now be selected. Click one of the flip tools to flip.



Fixing Distorted UVs

Once you've mapped out your UVs, you'll want to check to make sure that there's as little distortion as possible. You can do this by enabling the Display checkered tiles option in the UV Texture Editor. Enabling this will cause the background of the UV Texture Editor to show checkered tiles, and those checkered tiles will get mapped to your viewports.




In the example above, you can see that the checkered tiles on each face of the head look heavily distorted. You can see some faces have the tiles smooshed and others have the tiles stretched.

It doesn't really matter what direction the checkers are facing, but it matters that they're all the same size. These clearly aren't all the same size.

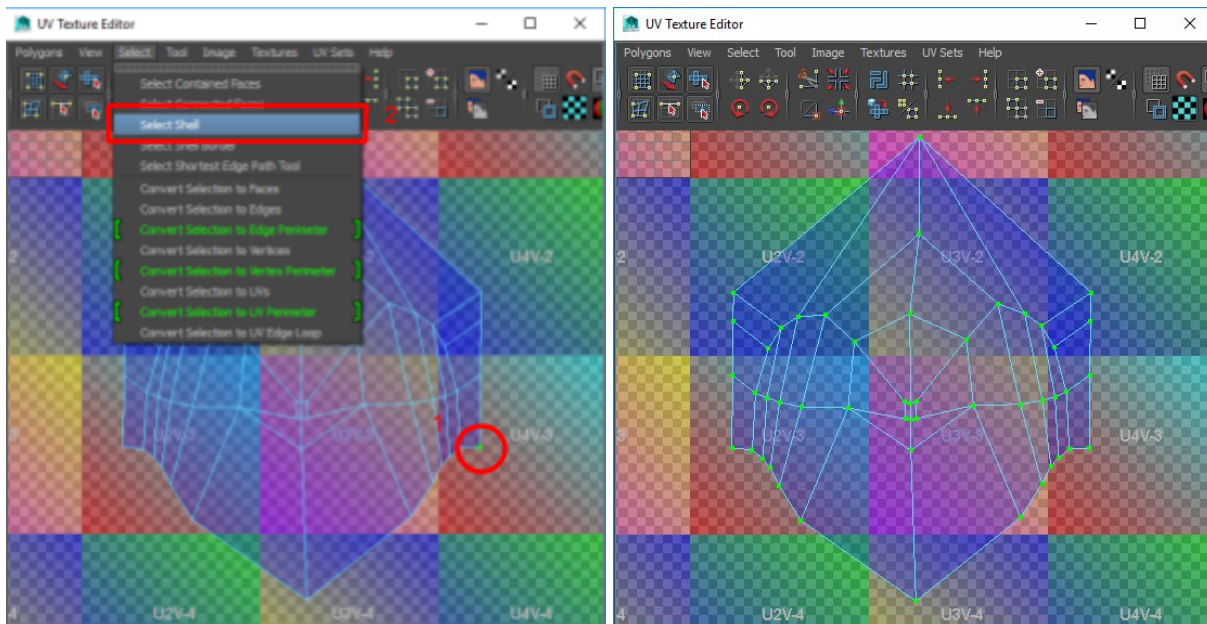
NOTE: The size of your UVs will only matter when you're exporting as a texture image for an artist to draw on. For now, it's fine to have different shells be different sizes, so long as distortion is kept to a minimum.

Maya has a tool to fix this called the unfold tool.

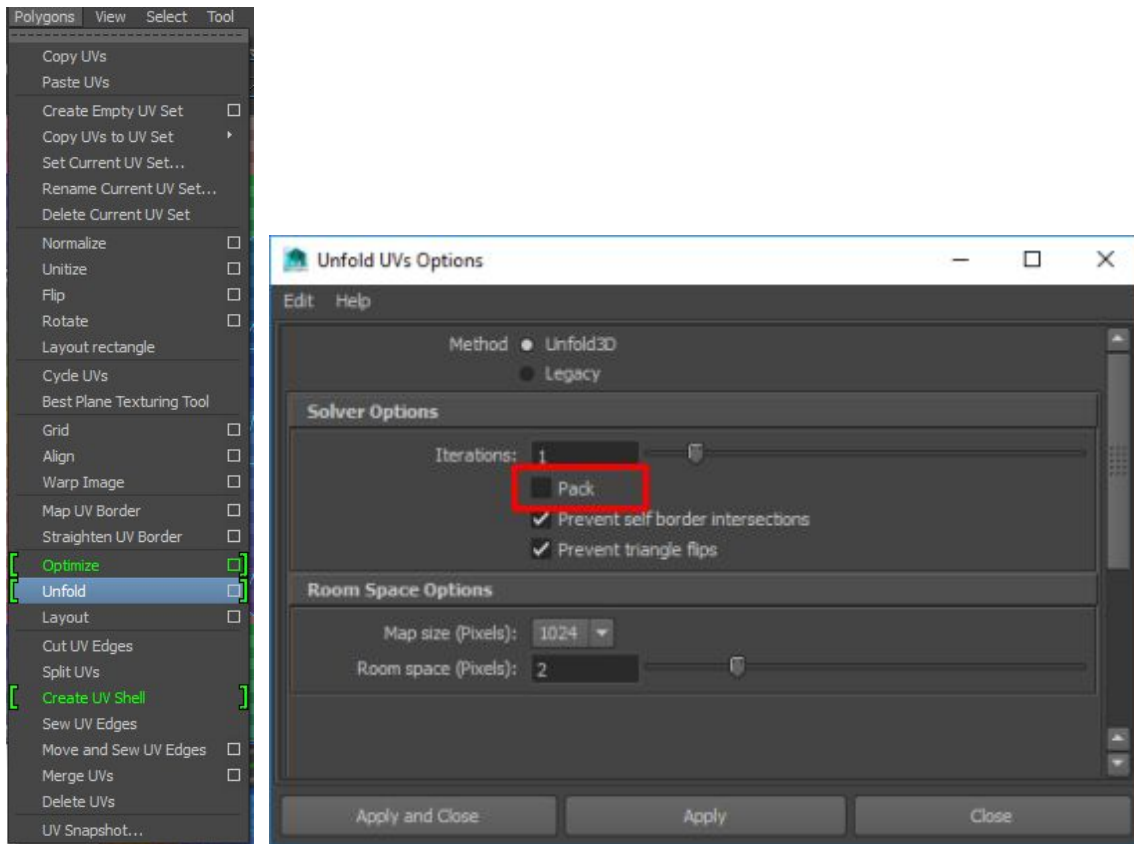
To use the unfold tool, first go into UV selection mode...



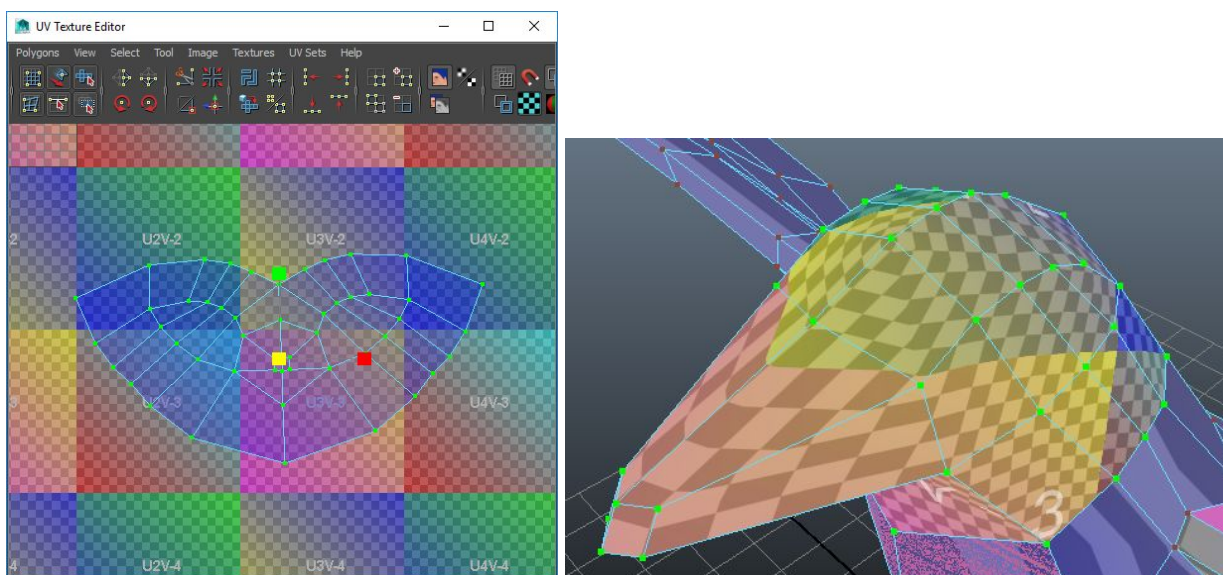
Then, select a single UV on the shell and go to Select -> Select Shell...



With the UVs selected, go to Polygons -> Unfold and open the tool options (little box on the right). In the window that opens up, make sure the Pack option is unchecked and hit Apply.



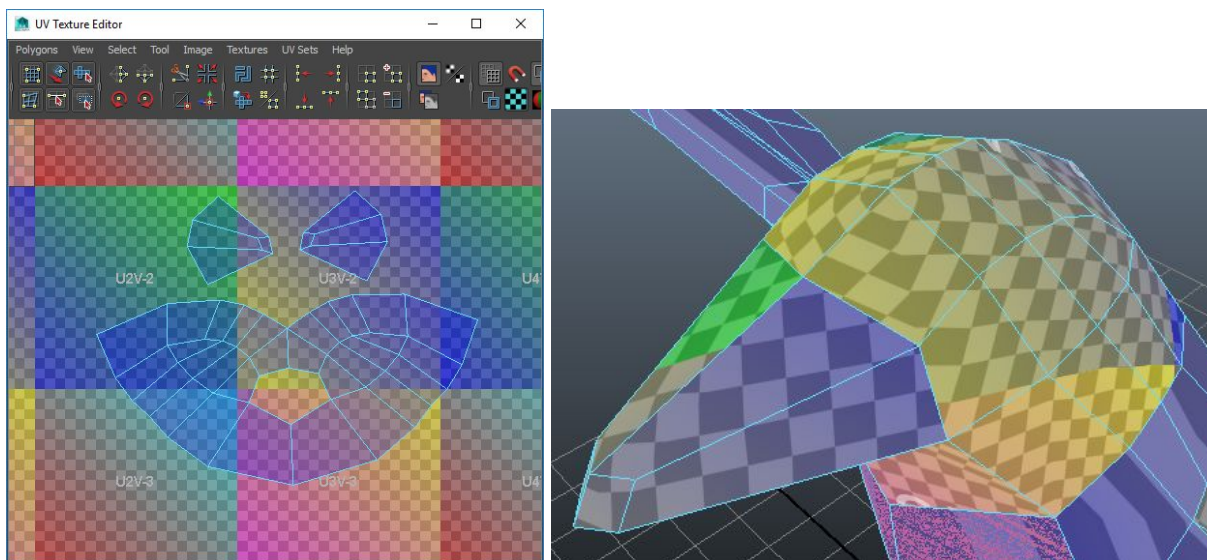
The selection will shrink down. Scale it back up using the scale tool.



NOTE: The scaling tool is bugged. If you want to uniformly scale, you need to select from the **BOTTOM LEFT** edge of the yellow box. Selecting the top/right edges will cause the scaling to not be uniform.

Compared to before, this looks a lot better. But, it still has problems (the beak looks super skewed). Turns out that in a lot of cases where you have faces converging into a point-like shape (like this bird's beak), it's hard to keep everything as a single UV shell.

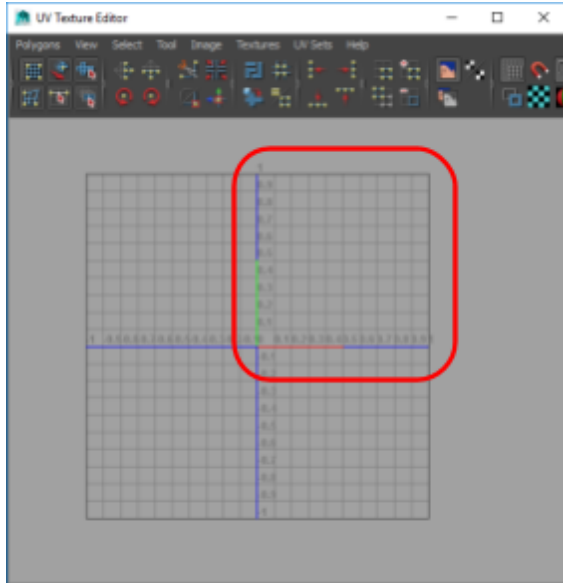
Everything other than the beak looks fine, so we can simply remove the portions of the beak as into their own shell and use the unfold tool again on each of the shells.



In the example above, we took out the beak and split it into 2 pieces. After unfolding the head and the pieces of the beak, a lot of the distortion has gone away.

Saving UVs as Texture (Snapshotting)

Once you have all your UVs mapped out, you can use the unfold tool to pack them into the 0.0 to 1.0 texture space.




NOTE: For many larger models, we don't pack all the UVs into the 0 to 1 texture space. Instead, we'll pack pieces of the model into different regions. For example, if you had a human body, maybe the head would go into $x=0-1/y=0-1$, the torso in $x=1-2/y=0-1$, the left leg in $x=0-1/y=1-2$, etc.. This way we don't end up with mega textures that can't all be loaded into memory at once.

Select your object in the viewport.


In the UV Texture Editor, go to Select -> Convert Selection to UVs. Then, go to Polygons -> Unfold and open the options.

Set the following options and hit apply...

- **Pack** is selected
- **Map size** is set to how large you want your 0.0 to 1.0 texture area to be (2048 is a good value)
- **Room space** is the amount of space you want between your UV shells (10 is a good value, any lower and you might get 'texture bleeding')



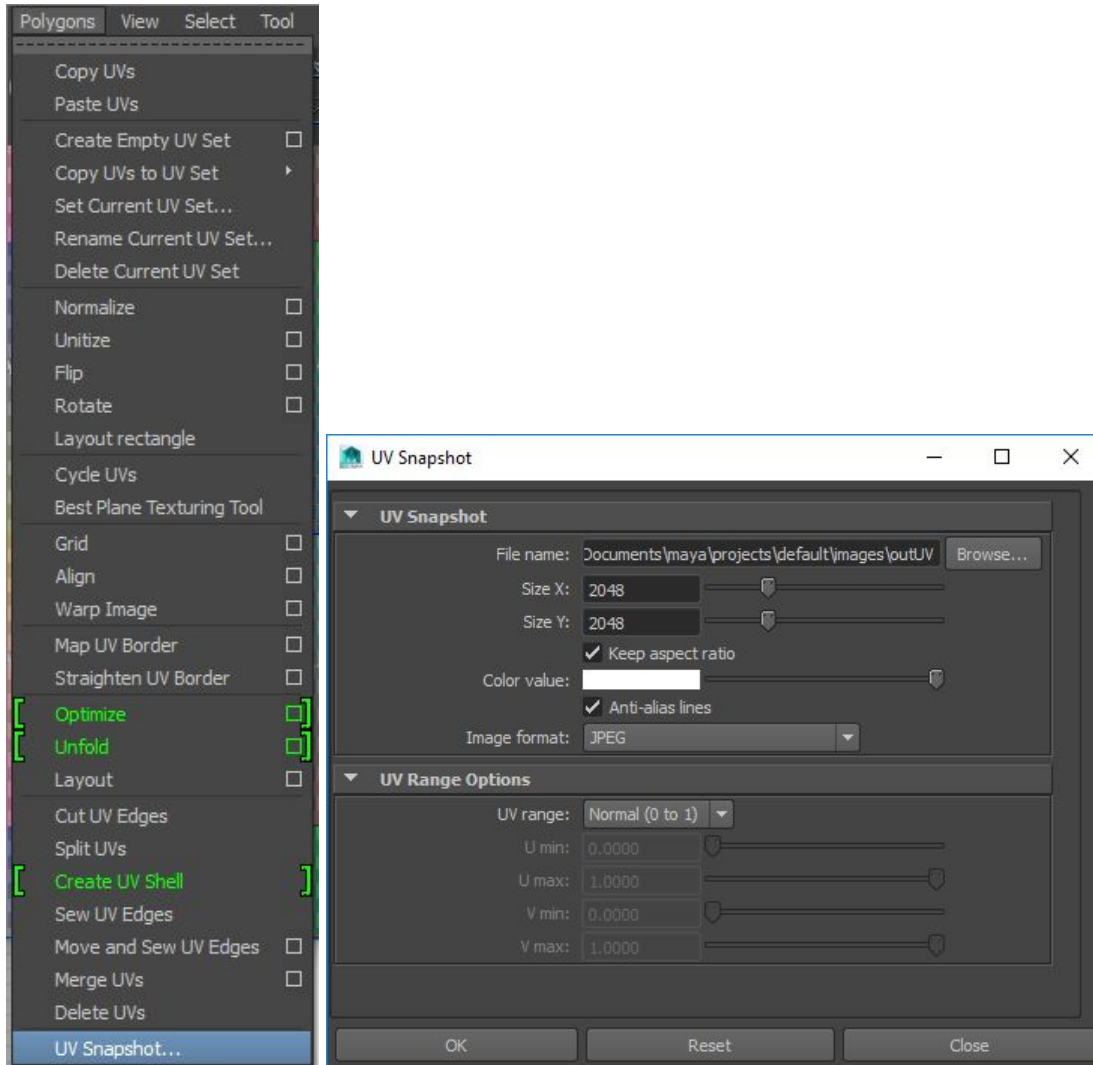
After this, all your UVs should be sized evenly and packed into the 0 to 1 texture space.



While this is okay, most times you'll want to move these around by hand a bit to make sure that related regions are close to each other. You can use the Move UV Shell tool to move shells and make sure they don't collide.

Just remember: At this point you only want to move and you never want to resize. No matter where you move things, you'll want all your shells to ultimately be in the 0 to 1 range before you export a UV snapshot.

At this point, you can export your UV snapshot by going to Polygon -> UV Snapshot...

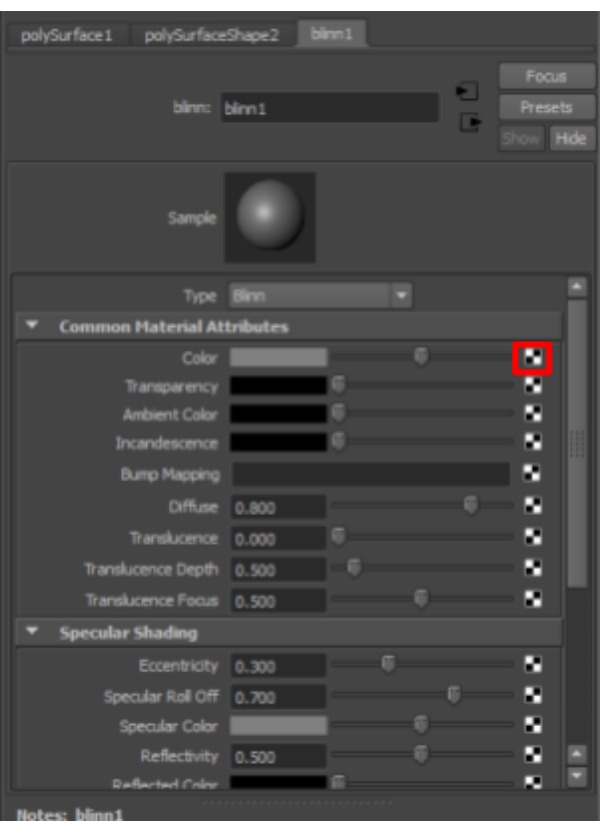
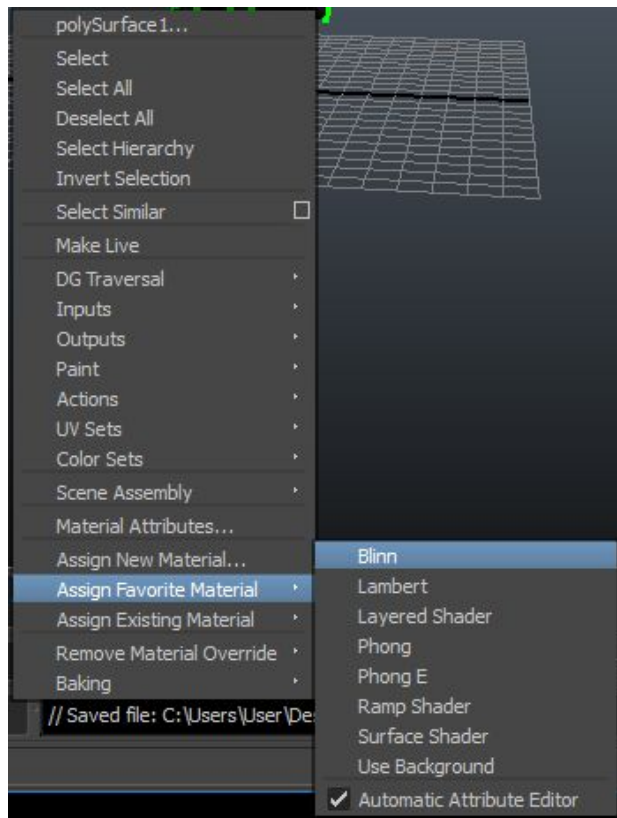



Set the location of where to dump the snapshot + set the size X / size Y to the map size value you chose in the pack tool + select your image format. Hit Ok to save the snapshot image.

Texture onto UVs

Once you've painted your UV snapshot, you can apply it onto your model by applying a new shader/material (e.g. Blinn) and setting the color attribute on that shader to the image.

NOTE: It looks like you need to Delete By History before you do this.



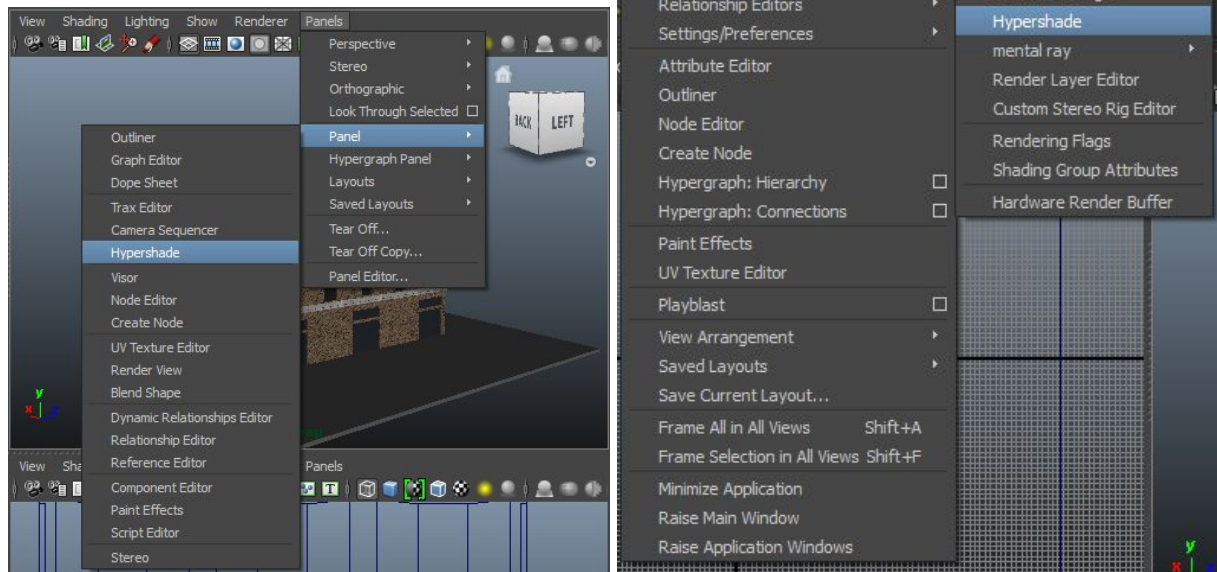


Hypershade

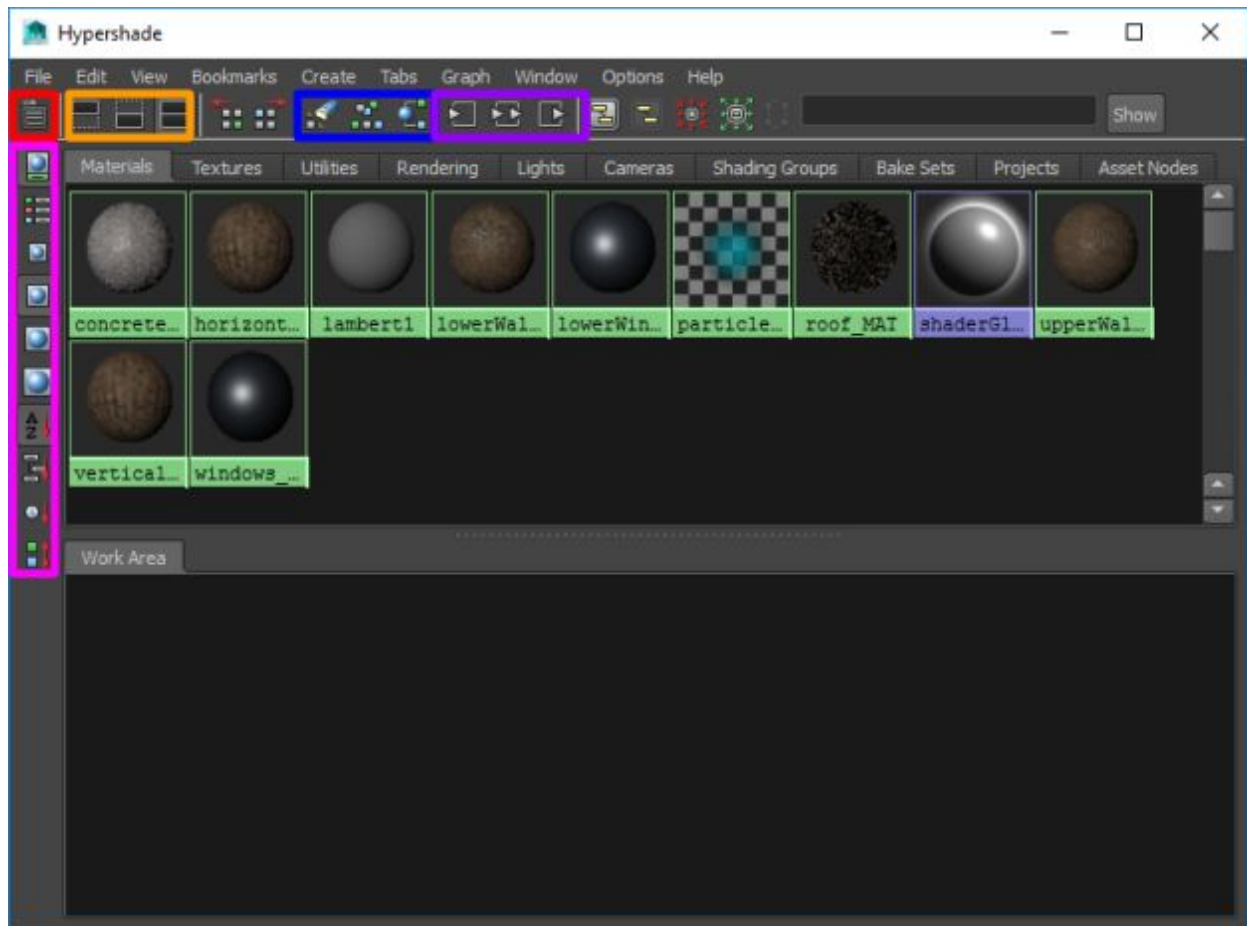
Most shaders are made up of a DAG structure where nodes feed into each other. The Hypershade is a tool inside of Maya that exposes this and allows the user to create/modify these nodes and connections.

There are 2 ways of opening up the hypershade...

1. In a panel, you can go to Panels -> Panel -> Hypershade (opens in panel)
2. Go to Window -> Rendering Editors -> Hypershade (opens as new window)



NOTE: You can only have one instance of the hypershade open.



Create control

Use this button new create new shaders / textures / <node type here>. Opens up a panel that allows you to select different types of nodes

Panel controls

By default, the hypershade has 2 panels open...

1. The top panel exposes various nodes: materials, textures, etc...
2. The bottom panel is a work area where you can view/create/modify connections between nodes

These buttons control which panels you want visible. It should be obvious what button does what.

Work area controls

Helpful buttons that control the work area.

- 1st button clears the work area
- 2nd button re-arranges the nodes in the work area
- 3rd button graphs out the shader(s) for the currently selected object

Input/output controls For any selected node (doesn't matter if it's in the top pane or the work area), the ...

- 1st button shows inputs to the node
- 2nd button shows inputs to AND outputs from the node
- 3rd button shows outputs from the node

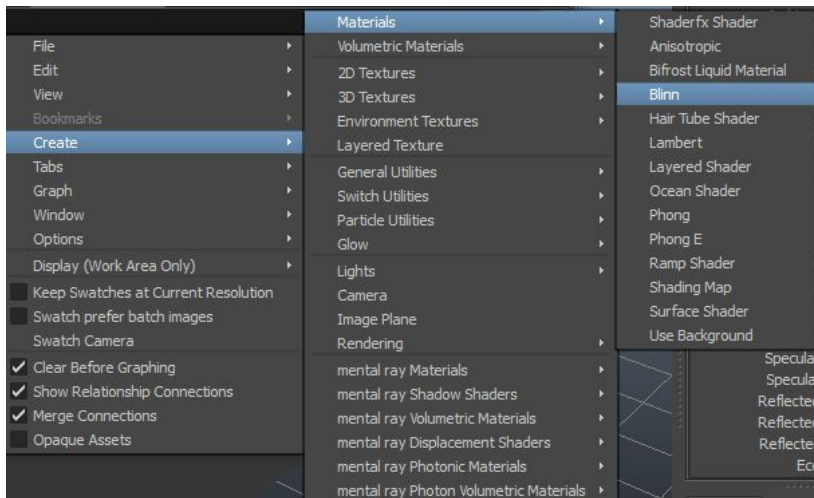
View controls Controls how the swatches are displayed. Doesn't seem to do anything at the moment.

Creating a Material

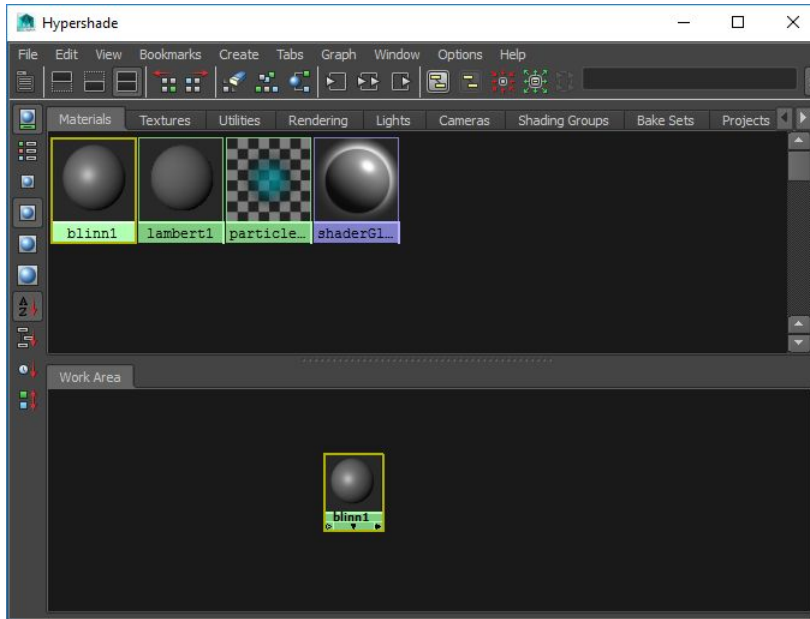
In the following example, we're going to create a new material that combines a couple of different textures with a blinn shader.

Begin by right-clicking either pane and going to Create -> Materials -> Blinn.

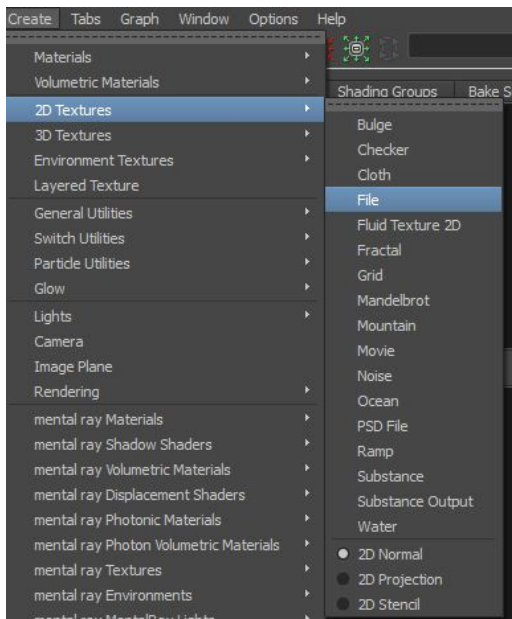
NOTE: You can also do this via the main menu or by using the Create button on the toolbar.



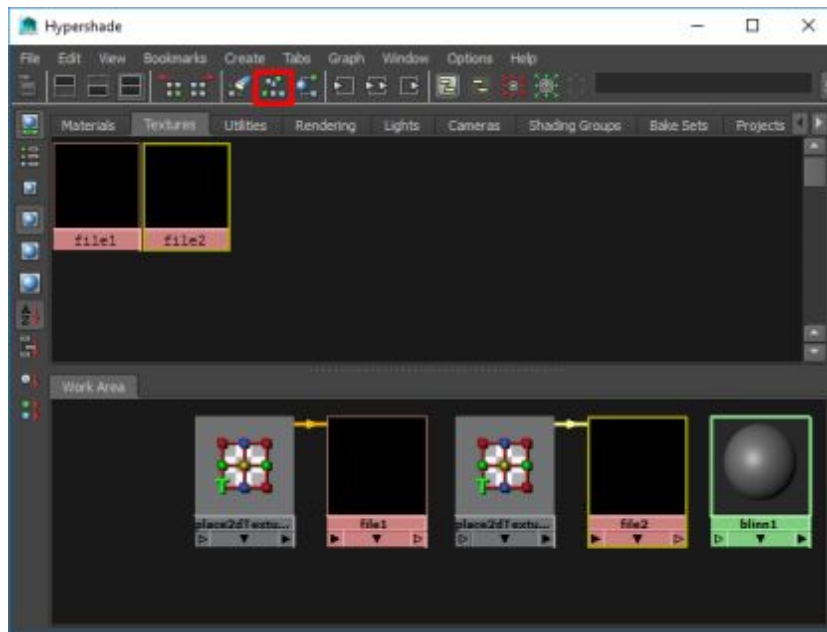
A new node for your blinn shader should show up in your work area + Materials tab + Shading Groups tab.



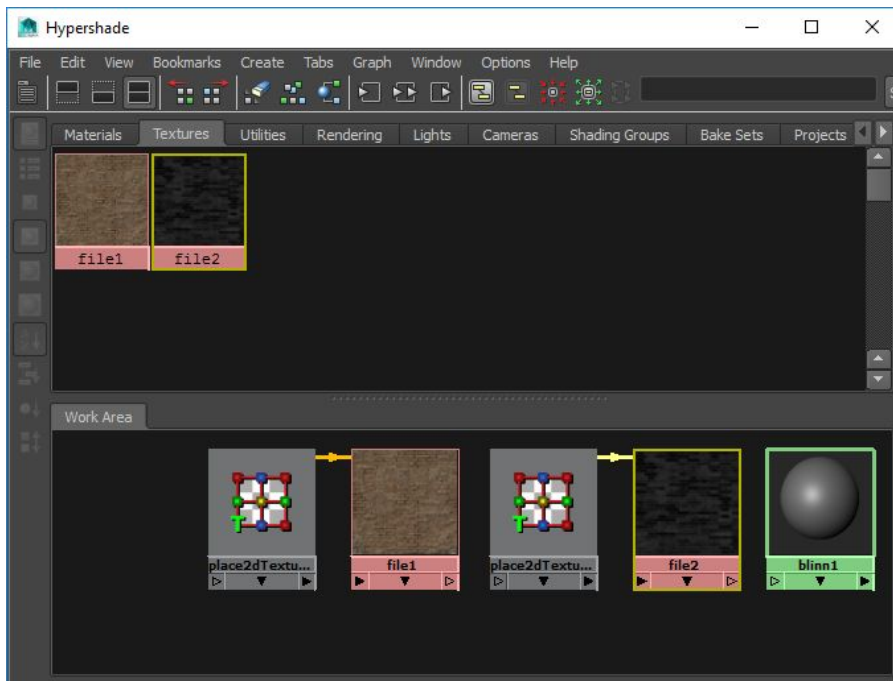
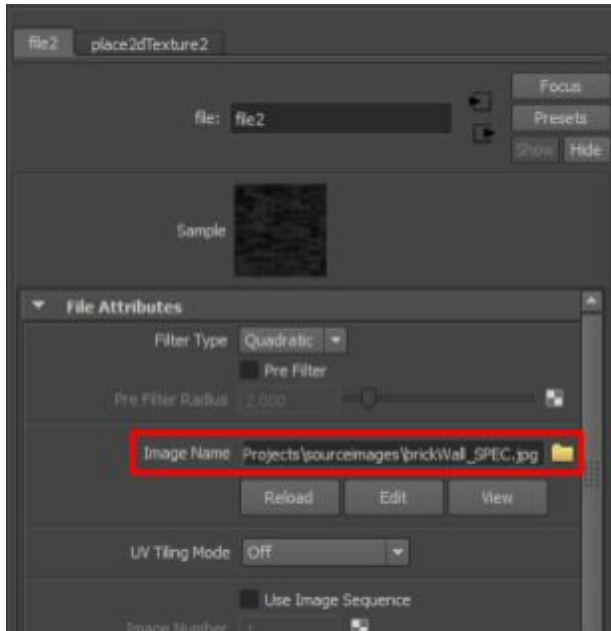
Now, create 2 file texture nodes by going to Create -> 2D Textures -> File.



The 2 nodes should be added to your work area + the Textures tab. Hit the Rearrange graph button to properly organize the work area after you do this.

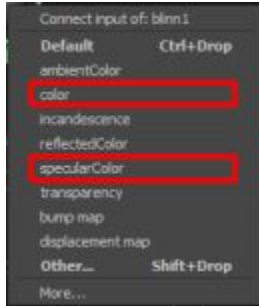


When you select a node in the hypershade window, you can view/change it's properties in the channel box and attribute editor. For each texture node -- select the node, go into the attribute editor, and open an image file.



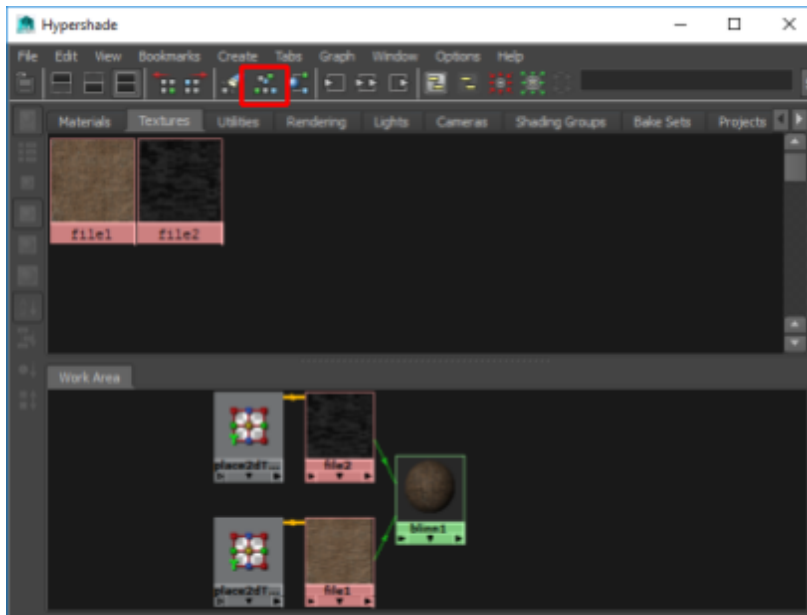
Now you need to connect file1 and file2 to blinn1. File1 should go to the color channel. File2 should go to the specular color channel.

1. Click-and-drag MMB on file1 and drag it to blinn1. In the menu that opens, select color.
2. Click-and-drag MMB on file2 and drag it to blinn1. In the menu that opens, select specular color.




NOTE: Another way to do this is to select blinn1, open up the attribute editor, then click-and-drag the file nodes directly onto the attributes in the attribute editor.

Hit the Rearrange graph button to properly organize the work area after you do this.



Apply your material to your texture as you normally would... Right-click on object, Assign Material -> blinn1.



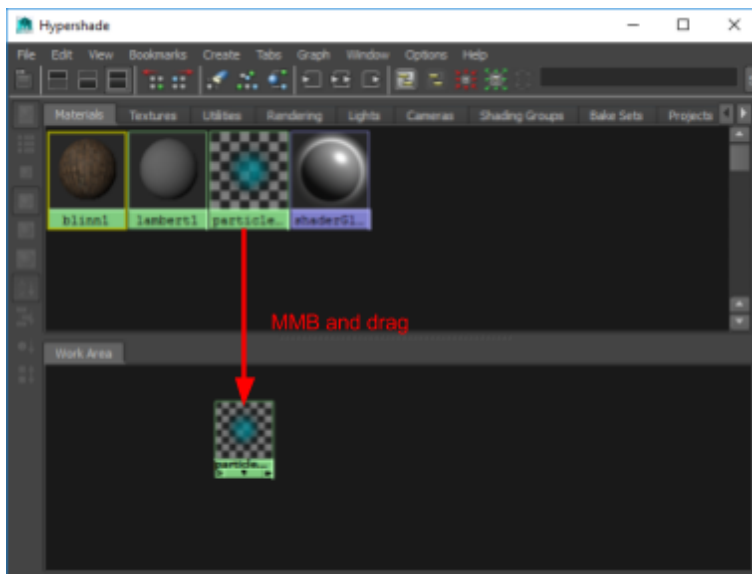
Modifying a Material

Modifying a material is similar to creating a material.

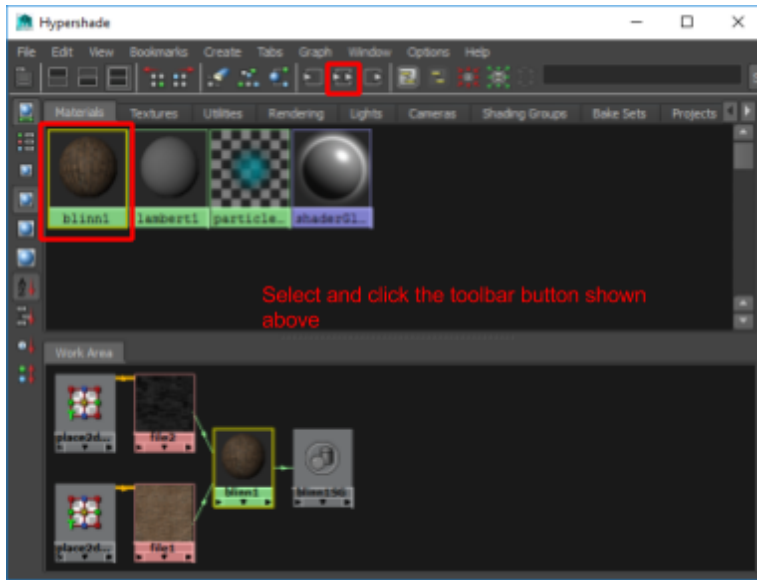
You can create nodes just like you did in the section on creating materials (previous section, just above this one).

You can connect nodes just like you do when you're creating materials (previous section, just above this one).

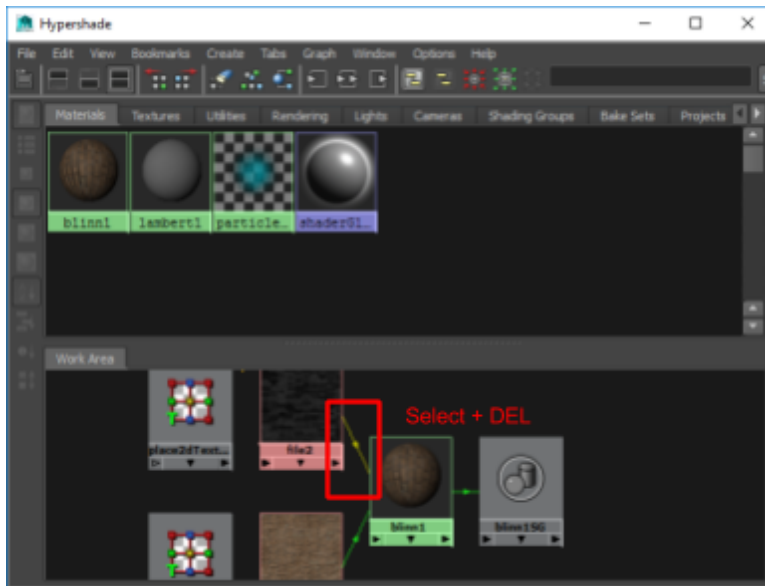
You can drag in existing nodes to the work area from any of the top panes (e.g. textures, materials, etc..) by click-and-dragging MMB on the node and moving it into the work area.



To view the inputs and outputs of a node, simply select it (doesn't have to be in the work area) and select the show inputs and outputs button in the toolbar. If the node wasn't already in the work area, it will be now along with its input/output connections.

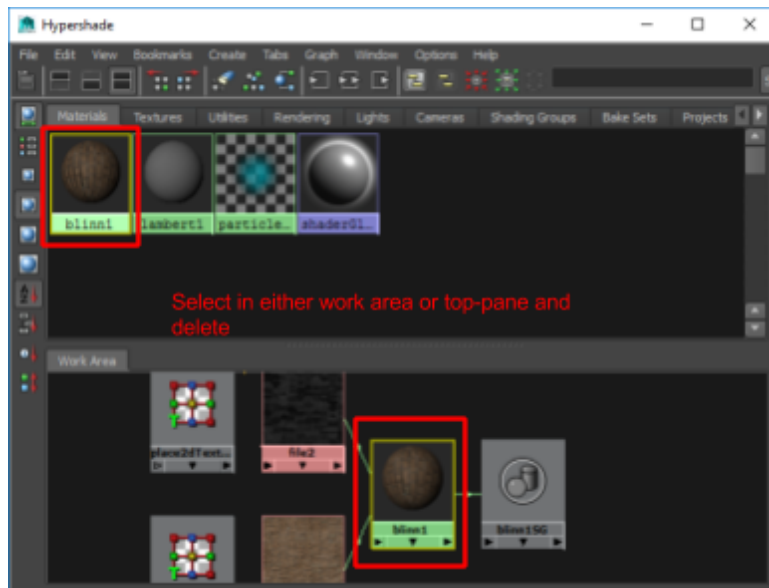


To delete a connection, select the connection and hit the Delete key on the keyboard. Note that this won't delete the node, just the connection.



Deleting a Material

If you want to permanently delete a material/texture/etc... node, select it and hit the Delete key.



Maya 2015 Rigging

[Joints](#)

[Creating Joints](#)

[Joint Placement](#)

[Joint Mirroring](#)

[Modifying Joints](#)

[Connecting Joints](#)

[Disconnecting Joints](#)

[Renaming Joints](#)

[Finalizing Joints](#)

[Creating Skinned Geometry](#)

[Updating Skinned Geometry](#)

[Painting Weights](#)

[Adding Weights](#)

[Removing Weights](#)

[Hammering Weights](#)

[Copy/Paste Weights](#)

[Mirroring Weights](#)

[Updating Joints Post-skinning](#)


[IK \(Inverse Kinematics\) Handles](#)

[Rig Controls](#)

Joints

Joints are the fundamental unit of a rig. They hold on to rotation + position and are used to mimic the functionality of a skeleton.


While joints can be standalone, bones are often used to connect two or more joints together.



In Maya, bones are NOT an explicit object (like joints are). They automatically show up between joints once you parent one joint to another.

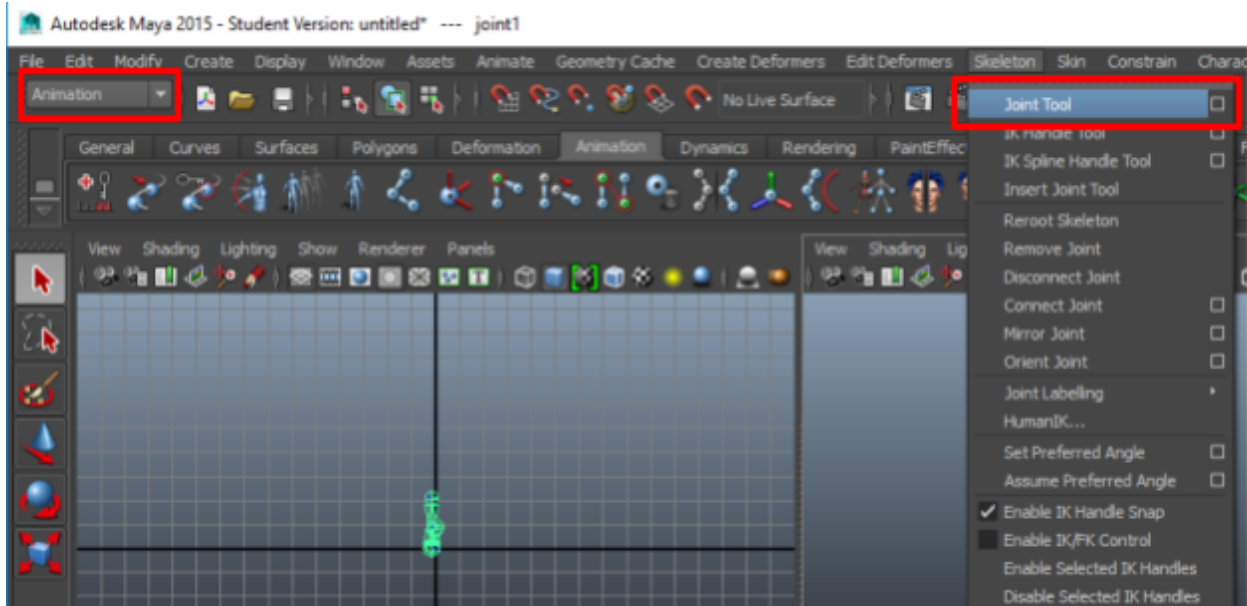
Depending on the type of rig you're doing you may have lots of bones or barely any. For example, it's unlikely that joints for a rig of a face will have many bones, while the joints for a rig of a body will have lots of joints.

HINT: Before doing anything with joints, make sure to turn on the joint xray feature in the viewports and turn off texturing (Press 4 to go into smooth shading mode only). Texturing sometimes makes it tough to see joints.




Creating Joints

To create joints, begin by going to Skeleton -> Joint Tool (must be in animation menu set).



Once the tool is active, you can click in your viewport to create a set of connected joints and hit Enter once you've placed your last joint.

HINT: Before doing anything with joints, make sure to turn on the joint xray feature in the viewports and turn off texturing (Press 4 to go into smooth shading mode only). Texturing sometimes makes it tough to see joints.




Joint Placement


NOTE: If you want to tweak the position of your nodes while you're still in creation mode, you can MMB click-and-drag the node to revise it.

You'll notice that the joints get placed on the grid plane of your viewport. So for example...


- if you created joints in perspective view, they'll end up on the ground.



- if you created joints in side view, they'll end up on the YZ plane.




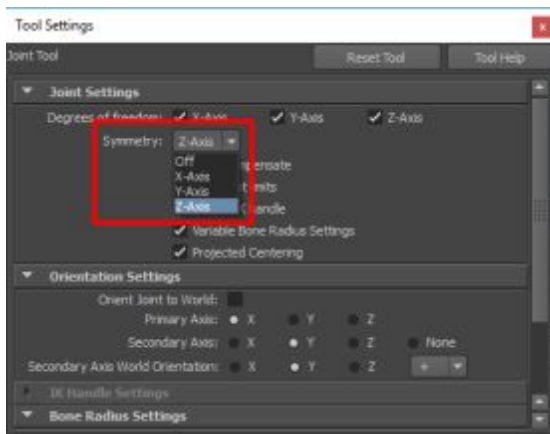
There are a couple of ways to get around this.



1. Select an object and set it as a live surface. Created joints will sit on the surface of that object.
2. Enable snap to projected center. Created joints will sit in the center of whatever part of whichever object you hit when you create the joint.

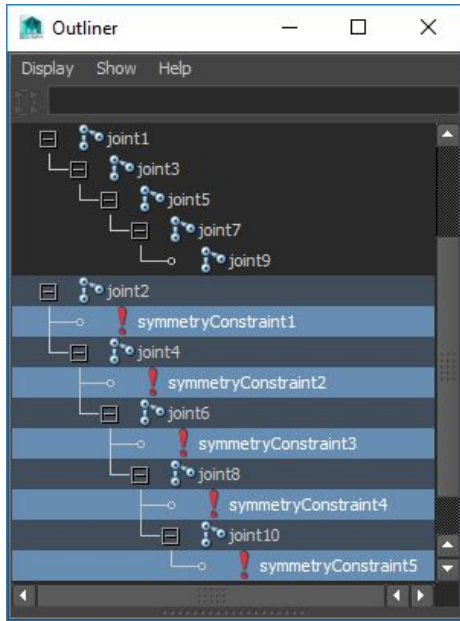
Joint Mirroring

If your model is symmetrical, you can go into the joint tool options and turn on symmetry before you begin creating your joints. Symmetry will mirror joint creation across to the on whatever axis you specify.




When you create joints with symmetry enabled, the mirror joints will take on the transform properties of the original. As such, any scale/rotation/translation you apply to one will get applied to the other.

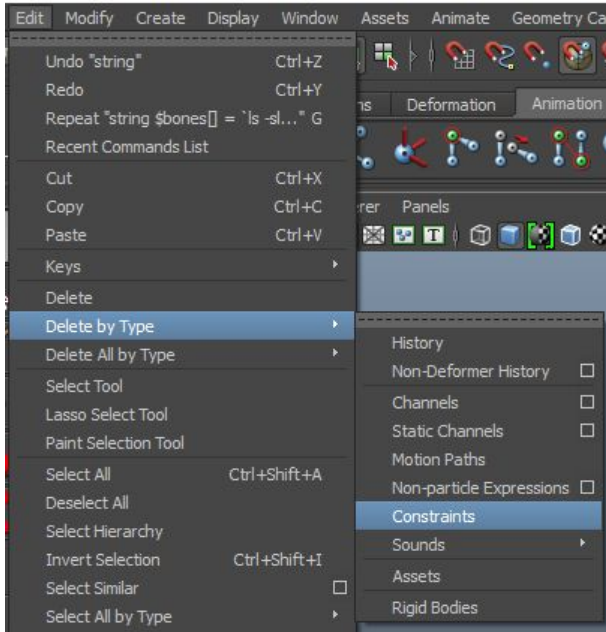
To stop this from happening, you can remove the symmetry constraints from the mirrored joints. Once you do this the mirrored joints will become independent.



NOTE: You can also right manually select each joint (not just the parent, but every single joint -- use the outliner if you have to), right-click to get the marking menu, and choose Break Symmetry.



NOTE: If the only constraints you have are from mirrored joints, you can go to Edit -> Delete By Type -> Constraints.



Modifying Joints


The top most joint in a chain/tree of joints is called the root joint. Remember that bones automatically show up between joints that a parent/child relationship.

NOTE: Reminder that to parent one object to another, select and press P. To unparent one object from another, select and press Shift+P. See parenting/grouping section for more information on parenting.


The fact that there's a parent/child relationship means that there are a couple of important properties in relation to joint manipulation. Child joints have local coordinates for rotation and translation. That means that...

1. Any rotation applied to a parent joint will get passed down to child joints.

In the following example, only the root joint was rotated by 45 degrees. Note how all the children down the chain are also rotated 45 degrees around the root.

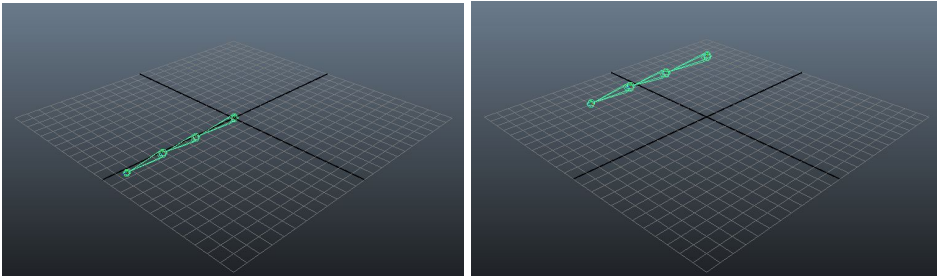


If I go down 1 child from the root and rotate by -90 degrees, we see all the children coming after it also rotated -90 degrees around the joint that had the rotation applied to it.



2. Any translation applied to a parent joint will get passed down to child joints.


In the following example, only the root joint was translated up by 5. Note how all the children down the chain also pick up that same translation.



Note that scale isn't automatically applied to child joints. Neither is radius.

HINT: Before doing anything with joints, make sure to turn on the joint xray feature in the viewports and turn off texturing (Press 4 to go into smooth shading mode only).

Texturing sometimes makes it tough to see joints.




NOTE: If you're wondering what radius is, it controls how large the joint is when viewed in a viewport. It has no effect on the joint -- it's just a visual thing. You can find it in the channel box when a joint is selected.

NOTE: The rotation and translation shown in the channel box / attribute editor for child joints contain LOCAL coordinates. That means that what's displayed is an offset from the parent, not the total change from the root.


NOTE: If you want to apply the same amount of rotational offset / translational offset / scale to each joint, just select ALL joints in the chain and rotate (can't use the marquee

tool for this -- either use the outliner or hold shift and manually select each joint in the viewport). For example, if you select ALL joints and rotate them (instead of just selecting the parent and rotating), the skeleton will look like it's curling in on itself. This is because each joint is taking all the rotations up the chain and adding its own rotation as well.




Connecting Joints

You can add bones between joints by parenting them together. For example, there are 3 separate skeletons....



If you want to select join them the diagonal skeletons to the skeleton going straight up, you need to parent them. To do this, first select the root joints of the diagonal skeletons, then select the root joint of the straight skeleton, then hit P.




This will cause the skeletons the root joint of the diagonal skeletons to parent with root joint of the straight skeleton, combining them into a single skeleton with the root joint being the original root joint of the straight skeleton (the joint at the origin).


NOTE: Remember how parenting works... The LAST node you select will be the parent. Everything you selected before it becomes a child of that parent.

Disconnecting Joints

You can remove bones between joints by unparenting the. For example, There's a single skeleton that branches out 3 ways...



The root joint for the skeleton is in the middle. If we want to remove the root joint from the 3 branches, we can select the joints just after the root joint and hit Shift+P.



NOTE: Remember how unparenting works... You select the CHILD and unparent it. You can't unchild the parent, you have to unparent the child. You can also unparent by breaking links in the hypergraph or the outliner.

Renaming Joints

When you create joints, Maya will use the default joint name for every joint you create (e.g. joint1, joint2, ...). What we want is to have a more intuitive naming scheme for our skeletons. For example, if we had a skeleton of a human body, we want the joints for the right arm to be prefixed with "r_arm", the left arm to be prefixed with "l_arm", the torso with "torso", etc..


Rather than manually renaming these joints one-by-one, we can use Maya's bulk renaming functionality.

If you haven't already done so, expand the toolbar to expose the rename tool...



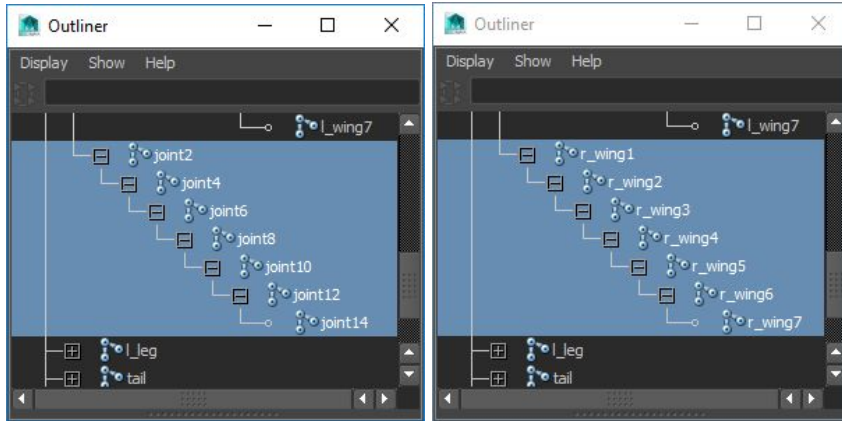
Once you've done that, individually select the joints you want to rename **IN THE PROPER ORDER**. That is, if you're selecting joints in an arm, you want to start by selecting the top-most joint first, then the second most top-most, all the way down to the leaf/child joint.

It doesn't matter if you're selecting in the viewport or the outliner or the hypergraph. **ALWAYS SELECT IN THE PROPER ORDER** because when you rename, it will append a number to each joint (the first one you selected will be appended with 1, the second one will be appended with 2, and so on and so forth).




Once you're done selecting, go into the rename toolbox, type in your new name and hit Enter. The joints will now be renamed.

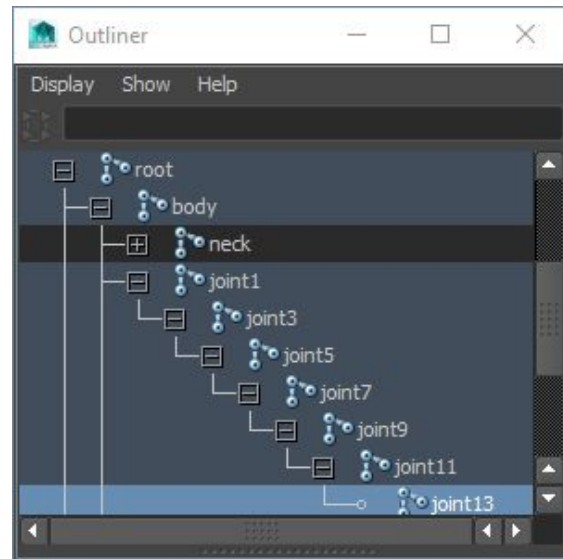




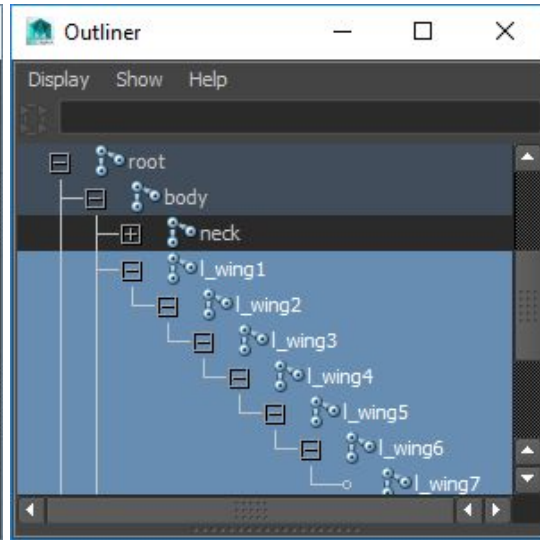
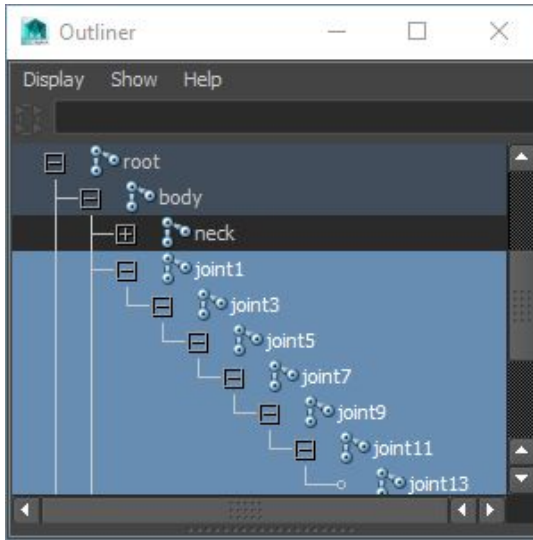

PRO-TIP: If you have the Outliner open, you can go into your viewport and select the leaf node.



Then, in the outliner, you can right-click Reveal Selected to expand.



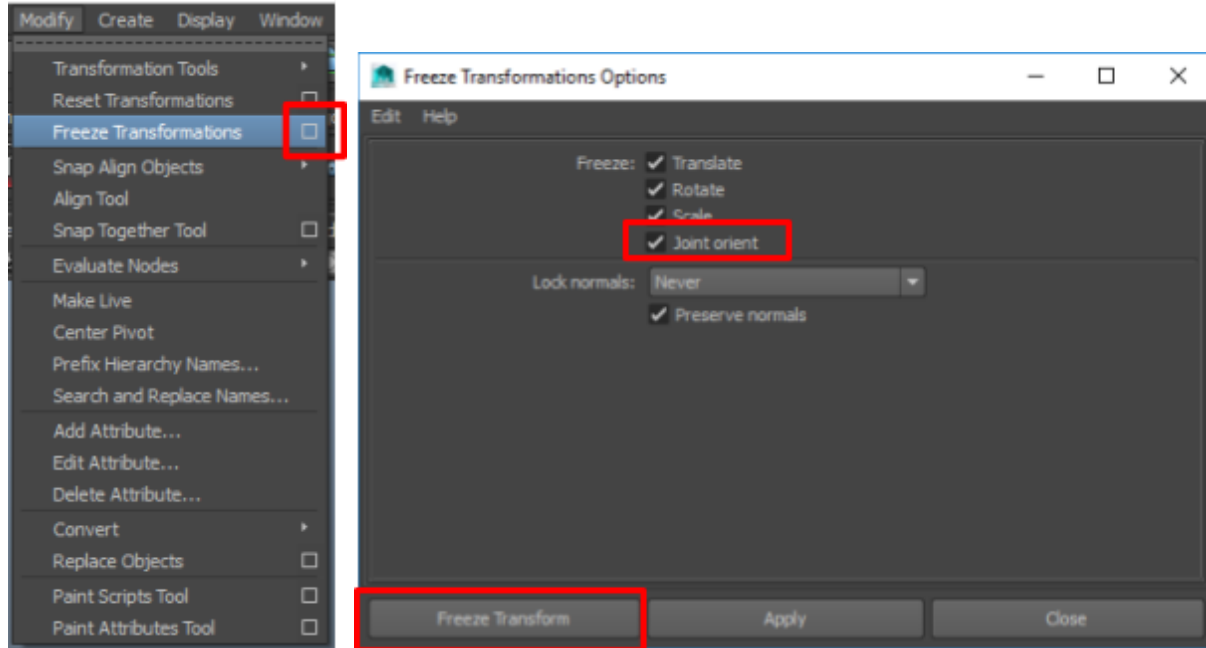
Click the top-most node and shift-click the bottom-most/leaf node. **MAKE SURE TO DO IT IN THIS ORDER.** You can rename at this point.



Finalizing Joints


Once you've finished moving around and rotating your joints such that they're in their default/base position, you need to freeze the transforms. To do this...

1. Select all your joints.
2. Go to Modify -> Freeze Transform and click the little box on the right to open the options.
3. Make sure "Joint Orientation" is checked and click "Freeze Transform".



Once you do this, you'll notice that the rotations zero out and any scale you set gets reverted back to 1.

HINT: Before doing anything with joints, make sure to turn on the joint xray feature in the viewports and turn off texturing (Press 4 to go into smooth shading mode only). Texturing sometimes makes it tough to see joints.

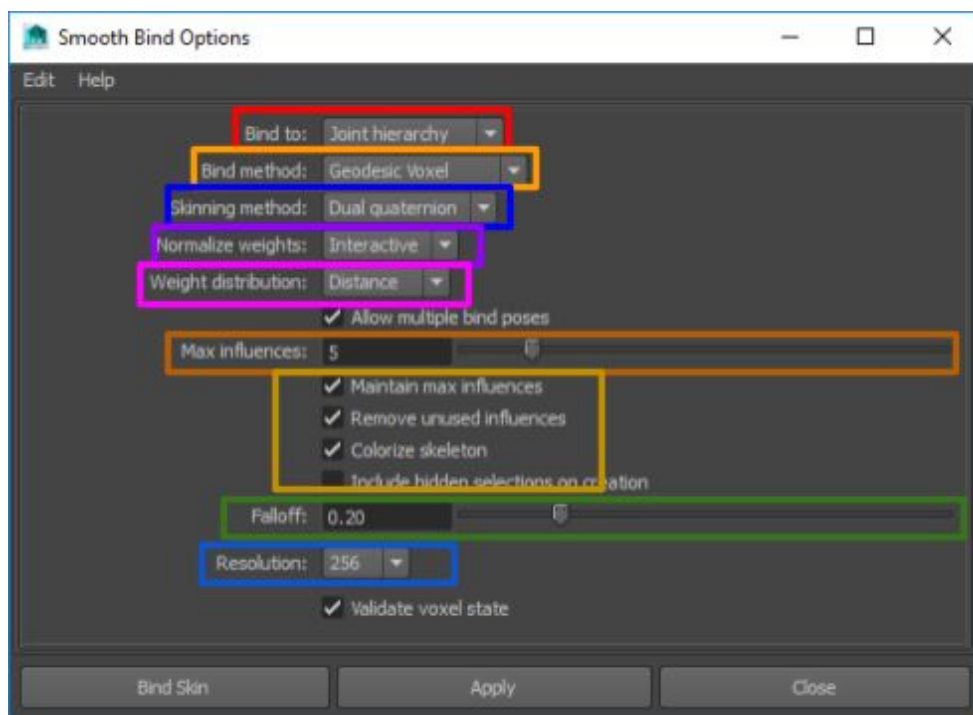



Creating Skinned Geometry

For joints to have any value, they need to be bound to geometry (also called skinning). Binding to geometry means that when you transform the joint those transformations also get applied to the geometry they're bound to (geometry in this case means a set of vertices). Since there may be multiple joints, a single vertex may be influenced by multiple joints.

Initially, we can assign the geometry of some object via the smooth bind tool.

1. Make sure you're in the Animation menu set.
2. Select both your joints and your object (multi-select), starting with the joints.
3. Go to Skin -> Smooth Bind and open the tool options (click the little box on the left).



NOTE: Many of these options are confusing. I've described them as best I could based on the content I'm learning from. If in doubt, use the options in the screenshots as the default.

Bind to

Options here include Joint hierarchy, Selected joints, and Object hierarchy. Typically you would use Joint hierarchy as you want your object to be bound to the entire skeleton.

Sometimes, we may want something different: we want to have joints that help control other joints, but we don't want these helper joints to be attached to any geometry. In that case, Selected joints may be more appropriate.

Bind method

For this, you almost always want to use Geodesic Voxel. This method automatically subdivides your geometry into voxels (voxel = cube / 3D pixel). You won't see this process taking place nor will you be able to view the voxel grid when it's done -- all of this is done internally.

Apparently Geodesic Voxel produces the most accurate binding. But, be advised that it does take a bit of time to apply (it doesn't happen instantly, you need to wait a few seconds/minutes for it to complete). If you can't deal with the wait, the next 2 best choices here are Closest in distance or Closest in hierarchy.

Skinning method

For this, you almost always want to use Dual Quaternion. This method tries to preserve the volume of your geometry, meaning that if you start twisting your joints you're less likely to end up with a sausage-linking effect that causes the volume of your object to go to 0.

Normalize weights

Interactive -- updates all of the weights every time you make a change to a vertex.

Post -- allows you to make changes to multiple vertices without having all of the weights update.

Post is usually a little easier to work with, but if you have a small number of vertices you can leave it as interactive. You can always change this option later on via the attribute editor / channel box. Once you finish skinning, select your object and choose the skinCluster node...



Weight distribution You almost always want to use Distance. This option looks at the distance each vertex is from each joint and weight it accordingly.

Max influences The max number of joints that will be attached to any single vertex.

Options Maintain max influences -- prevents going over the max influence setting (the option just above this one)

Remove unused influences -- removes influences with 0 weight.

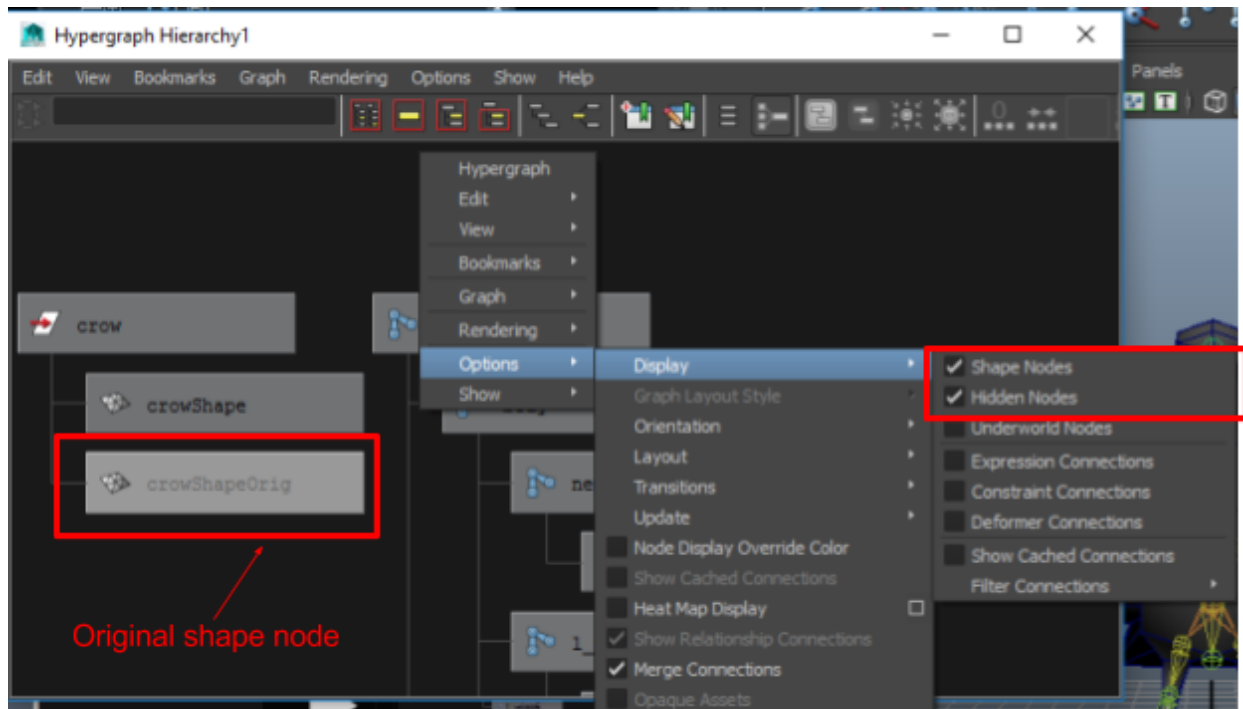
Colorize skeleton -- helps see which joints are bound to which vertices

Falloff No idea what this does. Seems to be related to geodesic voxel bind method.

Resolution Controls the number of voxels that'll be built for the geodesic voxel binding method. The higher the voxel count, the more detailed your weight distribution will be.

For low-poly/low-detail/low-resolution objects, you don't need a high number. For more details models (e.g. if you have a model with tons of little appendages like a centipede), you need to jack this number up.

No matter what binding method you choose, the end result will always be the same in the hypergraph. If you go into the hypergraph and turn on shape nodes + hidden nodes, you'll notice that the shape node for your geometry has been forked off to another node and a new shape node has taken its place.



The original shape node holds on to what the original object actually looked like before skinning. Transforming the joints causes deformations from that original object shape.

Updating Skinned Geometry

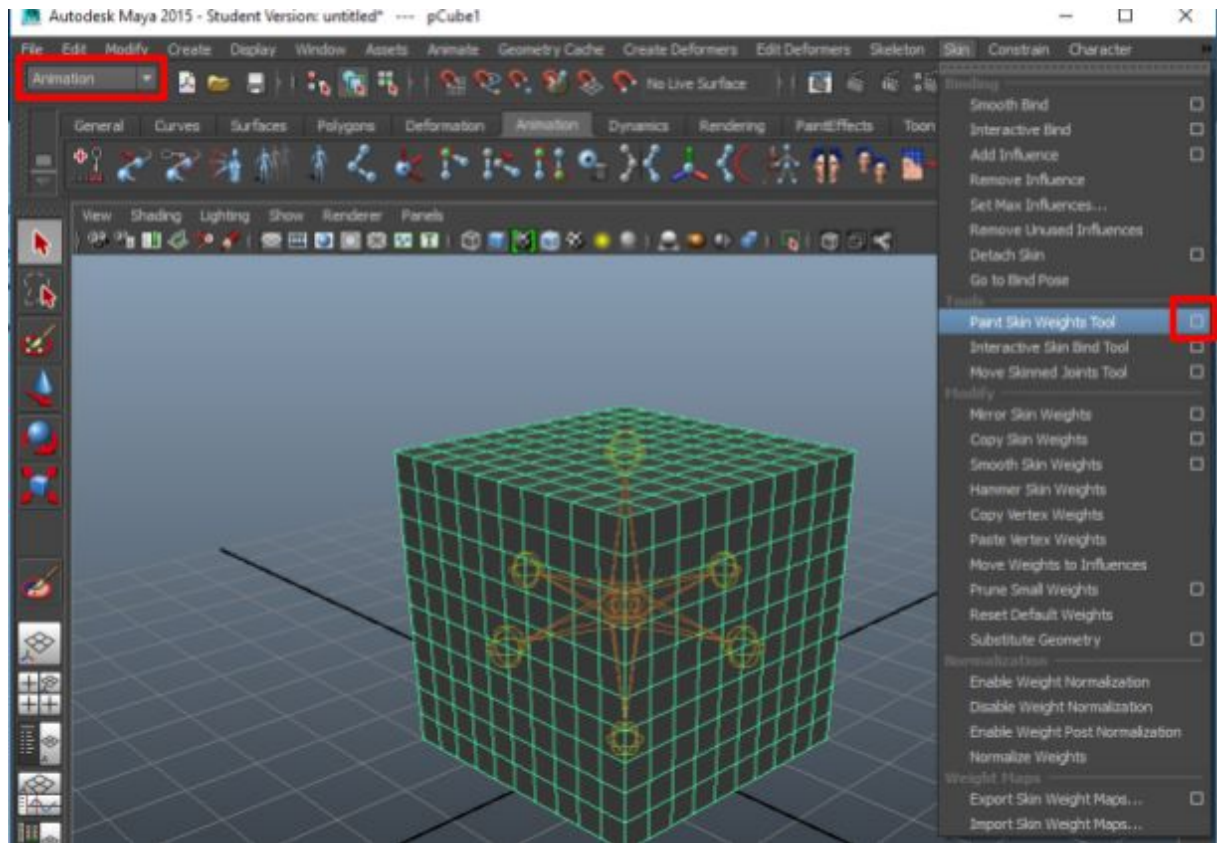
Once you've skinned your geometry, things may not come out exactly as you expect them to. Transforming certain joints may influence parts of your geometry that they shouldn't. For example, if you had a model of the hand and you curled in a finger, you don't want the joints in that finger to modify the geometry for the wrist.


The following sections describe some tools you can use to modify/correct the skin weights on vertices.

Painting Weights

You can paint the weights on your geometry using the Paint Skin Weights tool.

1. Select your object
2. Go into the Animation menu set.
3. Open the Paint Skin Weights tool, go to Skin -> Paint Skin Weights Tool (choose the little box on the right to open the tool options as well)...





Before you start painting, there are a few things you should do...

1. Select which joints your changes should **influence** and lock off all the other joints by clicking the lock button next to the name. Note that you can multi-select here and lock/unlock at the same time.
2. **Value** determines how much influence to assign to the selected joints....

A value of 0 will REMOVE.

A value of 1 will ADD.


If you don't want to fiddle with this value too much, you can always temporarily flip this while painting but holding down the Ctrl key (e.g. if value is set to 1, it'll be set to 0 while

Ctrl is held down).

NOTE: There are some intricacies with adding/removing weights that are described in the respective subsections.

3. Enable the **color ramp**. Which vertices are influenced are shown as grayscale by default. For lighter influence weights, it's hard to tell the difference. With Color Ramp selected it becomes much easier (small influences show up in black/blue and large influences show up in red/white).


In the example below, you can see a lot more definitively which vertices are influenced by the selected joint and by how much they're influenced.



Adding Weights

Remember that the total weight on a vertice has to always equal 1. That means that when you're increasing the influence/weight of a joint on some vertices, the influences/weights that other joints have on those vertices will need to be decreased accordingly.

You can use the locking mechanism to indicate which joints influences/weights should be taken from...




In the above example, I'm working on joint2. The only joints I have unlocked are joint2 and joint5. When I start painting vertices such that they're more influenced by joint2, it'll take those influences from joint5.


NOTE: Be careful when you're painting with locked joints as the colors showing up during the painting process are not representative of the total weights on the vertices. In the example above, you're setting things up such that you can only take weight away from joint5 and add it to joint2 -- for some vertex, you can't take away more influence than joint5 has on it.

Your paint job may look like it it's assigned ALL of the weights for some vertices to joint2, but that isn't the case. As soon as you click away from joint2 and back again you'll see the colors for the true weights.


Original weights for joint2...



Painting all of one side to joint2...



Clicking away from joint2 and clicking back on it again...



You've only taken weight/influence away from joint5 and assigned it to joint2. All the

other joints that have an influence/weights on those vertices still have the same influence/weight values! That's why the color ramp doesn't stay white.

If you do the same thing without joint locks, it'll work the way you expect it to.

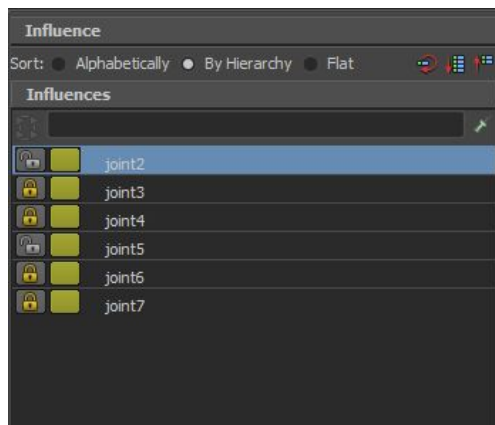
Removing Weights

NOTE: This is pretty much the same as the Adding Weights subsection, but we're removing here. The same process applies except we're removing instead of adding.

NOTE: Remember that to remove weights, you can set value to 0 in the Tool Settings of the paint tool, or you can just hold Ctrl while you paint.

Remember that the total weight on a vertice has to always equal 1. That means that when you're decreasing the influence/weight of a joint on some vertices, the influences that other joints have on those vertices will need to be increased accordingly.

You can use the locking mechanism to indicate which joints influences/weights should be taken from...



In the above example, I'm working on joint2. The only joints I have unlocked are joint2 and joint5. When I start painting vertices such that they're less influenced by joint2, it'll increase joint5's influence on those vertices accordingly.

NOTE: Be careful when you're painting with locked joints as the colors showing up during the painting process are not representative of the total weights on the vertices. In the example above, you're setting things up such that as you remove weights, they're only getting taken away from joint2 and being added to joint5 -- for some vertex, you can't take away more influence than joint2 has on it.

Your paint job may look like it's remove ALL of the weights for some vertices from

joint2, but that isn't the case. As soon as you click away from joint2 and back again you'll see the colors for the true weights (or click on the other joints and you'll see there is still color on those vertices from the other joints).

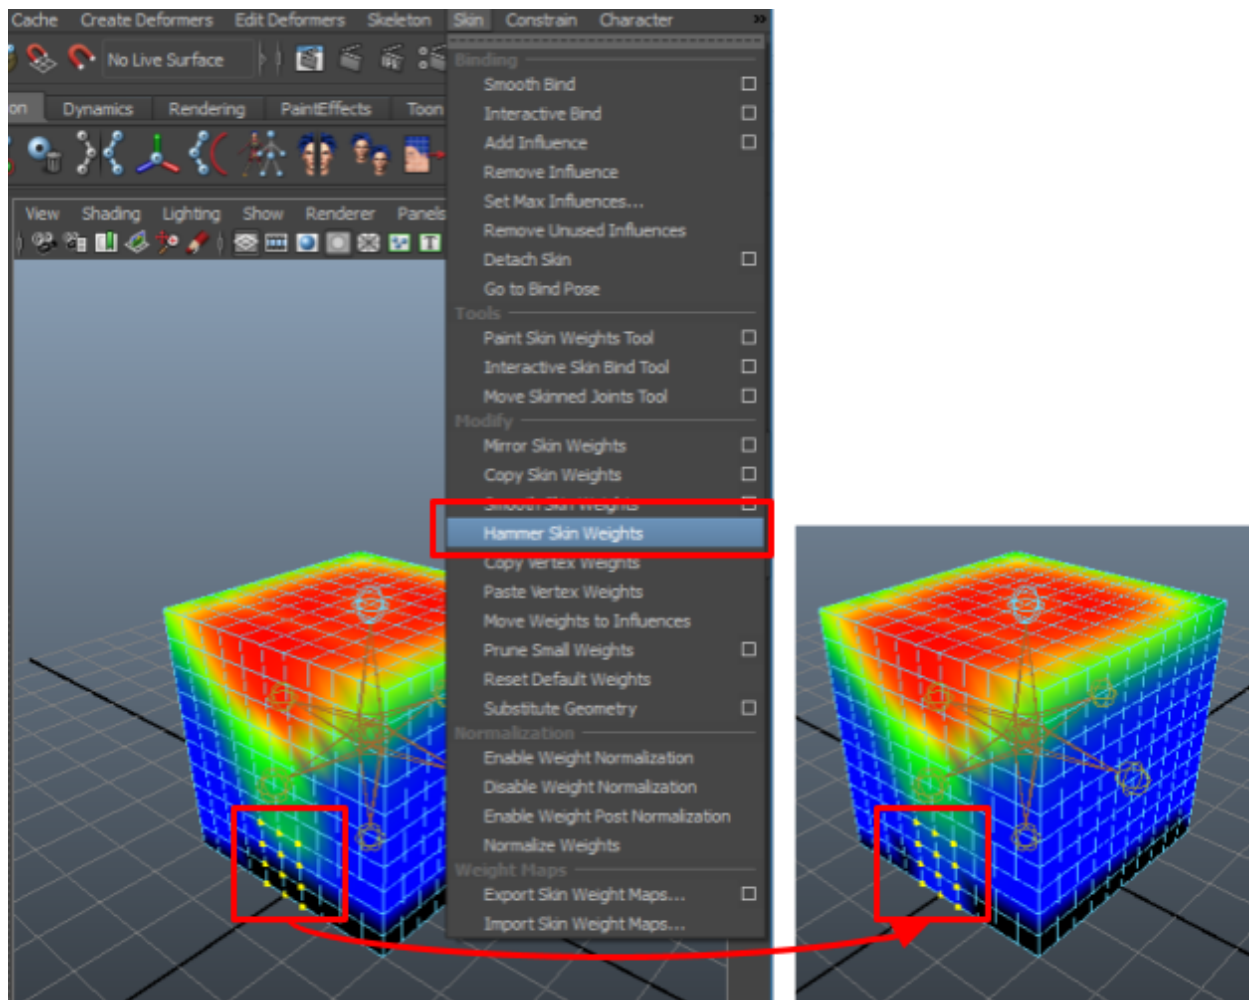
SEE THE EXAMPLE IN THE ADDING SUBSECTION. THIS IS JUST THE OPPOSITE. SAME THING APPLIES.

Hammering Weights

HINT: Use this instead of painting when you want to make small tweaks.

You can pick one or more vertices and "hammer" in the weights of the vertices that surround it/them by using the Hammer Skin Weights tool. The tool will assign weights by averaging the surrounding vertices.

To do this, go into vertex selection mode and select the vertices you want to "hammer", then go to Skin -> Hammer Skin Weights (make sure you're in the Animation menu set).



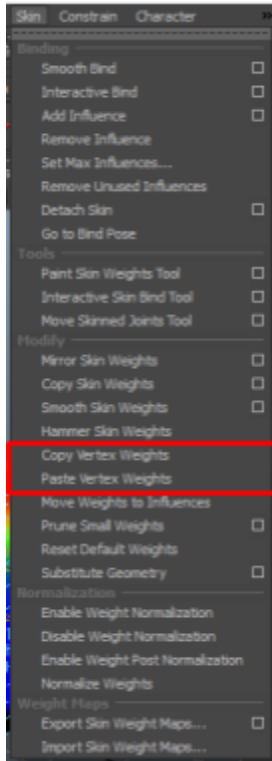
Copy/Paste Weights

HINT: Use this instead of painting when you want to make small tweaks.

You can copy the weights from one vertex to another using the Copy Vertex Weights / Paste Vertex Weights tool.

Make sure you're in the Animation menu set...

1. Go into vertex selection mode.
2. Select the vertex you want to copy, then go to Skin -> Copy Vertex Weights.
3. Select the vertex you want to paste to, then go to Skin -> Paste Vertex Weights.

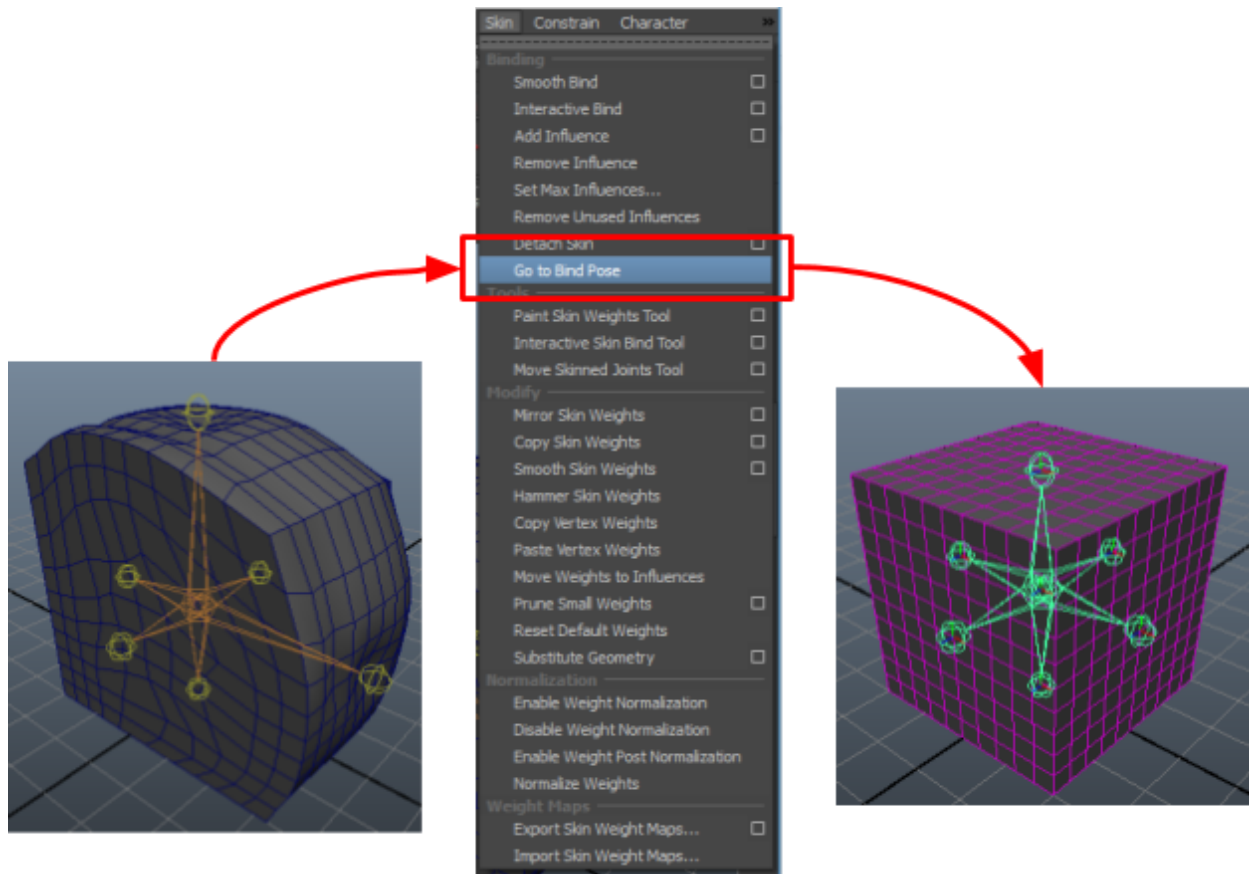


Mirroring Weights

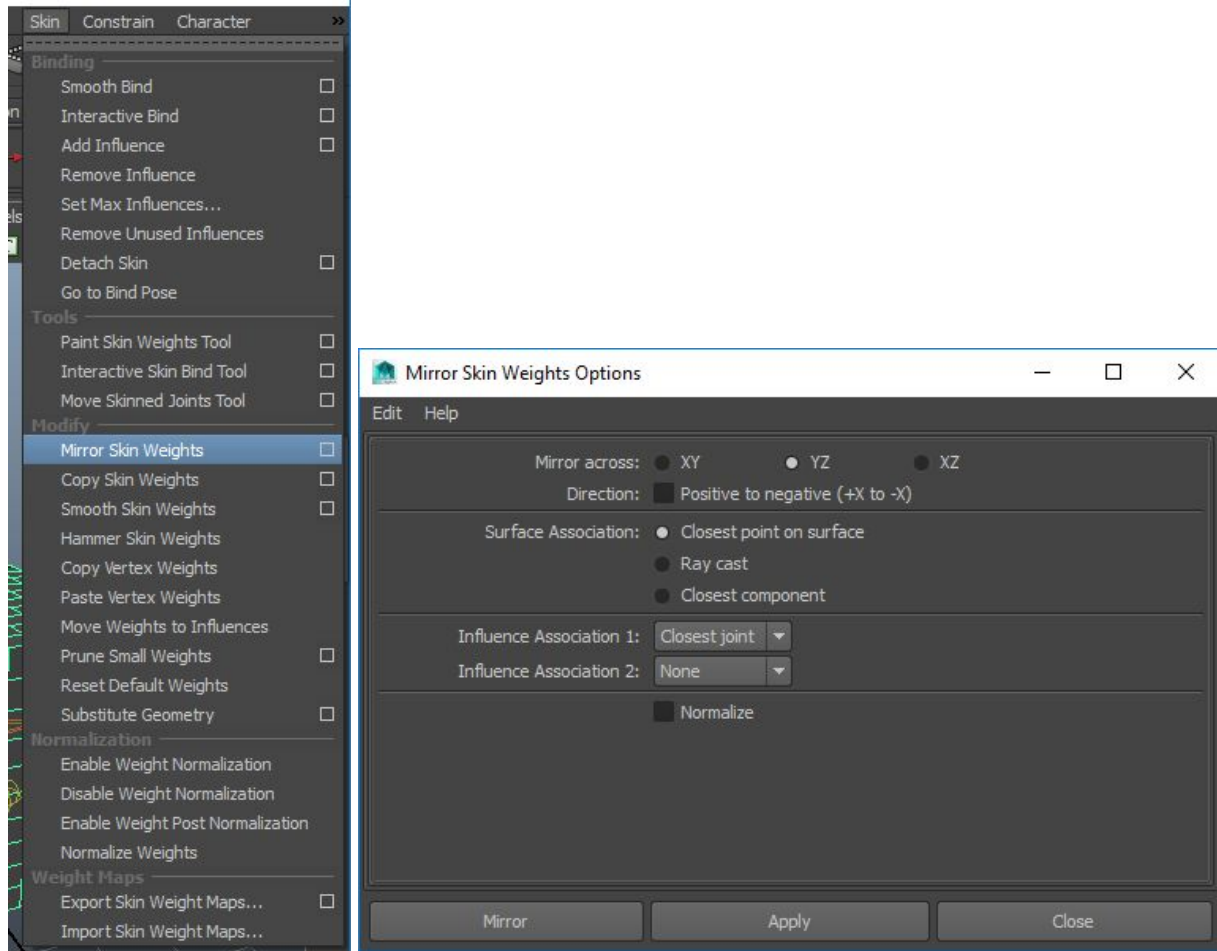
HINT: This is a big timesaver. You only have to properly weight half your model.

For mirroring weights to work, you must be in the default pose and the model/rig must be symmetrical.

To get back to your default pose, first select the root joint, then go to Skin -> Go to Bind Pose...




Once you've done that, you can mirror the weights by selecting the object (not the joints) and going to Skin -> Mirror Skin Weights (choose the little box on the right to open the options)...



The first 2 properties specify which plane you want to mirror across (which plane cuts your model symmetrically in half) + the direction to mirror from.

So If I choose YZ and I have Positive to negative checked, my model must be cut in half by the YZ plane, and I want weights copied from the side where X is positive to the side where X is negative...




For the rest of the options, pick the defaults in the screenshot. They're almost always what you want.

Updating Joints Post-skinning

If you've already created your skin but you realized that you made a mistake and you need to tweak the joint positions/orientations, you can do this using the Move Skin Joints Tool.

1. Go into the Animation menuset.
2. Open the Paint Skin Weights tool, go to Skin -> Paint Skin Weights Tool.
3. Select the joint you want to change and change it.

Once you select the joint, the associated geometry will go into wireframe mode and you can move around the joint without deforming the geometry. You can switch away to the select tool once you're done and your changes will be saved.



One thing to keep in mind here is that unlike normal joint selection, selecting a joint here does not automatically select its children. If you want to move around a joint chain / multiple joints, you need to multi-select them (Shift+LMB click each joint).

IK (Inverse Kinematics) Handles

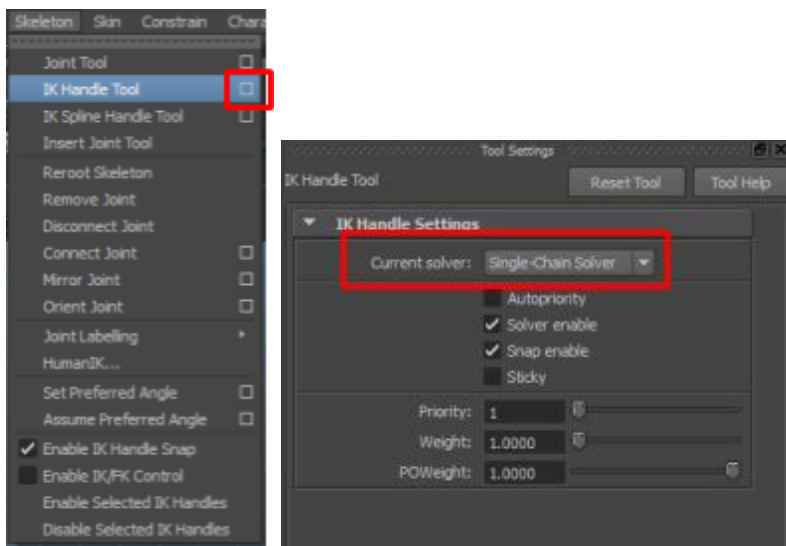
IK (Inverse Kinematics) is a system used to pose joint chains in a realistic manner using a single action. As you pose/animate joint chains, the IK handle figures out how to rotate and pose all the different joints in the chain.

IK handles are used for things like: limbs, fingers, tails, necks, spines, tentacles, etc...

Typically, a Maya IK handle requires you to specify an end joint (also called end effector). All joints between that end joint/effector and the root joint (can you choose a value that isn't root?) will get oriented and positioned automatically.


To create an IK handle on your skeleton...

Go to Skeleton -> IK Handle Tool and open the options (little box on the right).

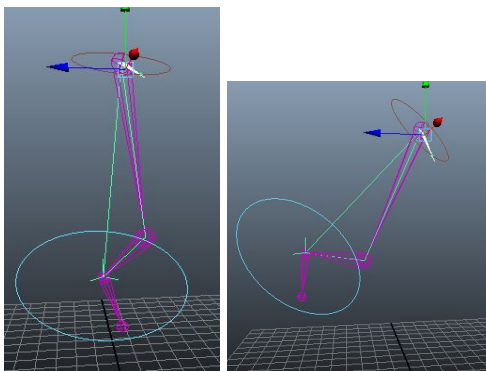


In the tool settings, you can choose which type of IK handle you want: single-chain solver or rotate-plane solver.

In the following example, I made a simple leg and used a single-chain solver and between the hip joint and the ankle joint. Note how there's a line running through the 2 joints that I selected. Now when I move the end joint, the entire chain (minus the foot because it isn't part of the IK handler) translates/rotates itself to look proper.



What's the difference between the single-chain (SC) and rotate-plane (RP) solvers? I don't know. The SC solver lets you rotate your leg skeleton from the end joint/end effector, but the RP solver has this circular hat on top of the start joint where the plane adjusts as the end effector moves around. RP rotations are done on whatever plane that hat is on + you have to rotate using the translate tool instead of the rotate tool....




NOTE: To access the translate tool on the rotational-plane top hat thing, select your end effector and hit T. To get back to being able to translate the end joint/effector, hit the normal translate tool hotkey (W)

NOTE: Why would you ever want a rotational-plane handler? Maya's website says why... It sounds like for things like rotating legs/shoulders? Constrain the rotational plane handler to something so that you can make the leg automatically turn in certain situations? For example, if you're leaning in, you want the leg to rotate inward? I still don't really know what this is for.

<https://knowledge.autodesk.com/support/maya/learn-explore/caas/CloudHelp/cloudhelp/2016/ENU/Maya/files/GUID-9942FFB5-65C2-46E2-B5A3-297667A9FB5D-htm.html>


The rotate plane solver is ideal for posing joint chains (such as arms and legs) that you want to stay in the same plane. For example, the shoulder, elbow, and wrist joints of an arm driven by a rotate plane IK handle all stay within the same plane as the elbow rotates. The plane itself can be rotated from the shoulder joint by the pole vector.



Rig Controls

“Rigging” is the art of adding controls to animate an object. Adding these controls makes animating these objects easier.

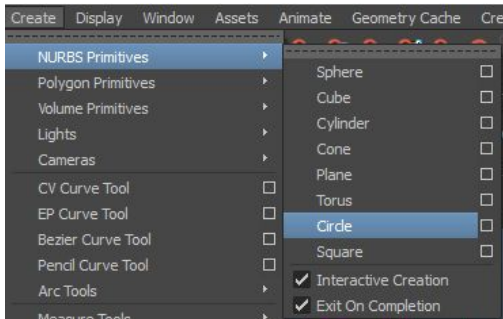
Rarely ever do you want to animate by manually manipulating the joints/skeleton. The typical workflow is to add control handles on top of the joints and move those instead.




In the screenshot above, the crow has a bunch of NURBS circles constrained to various joints on its skeleton. For example, if you wanted to orient the head, rather than manually moving the head joint you would move the NURBS circle that's around its head.

How would you go about creating a rig control? (make sure you're in the animation menu set)

1. Create your NURBS circle

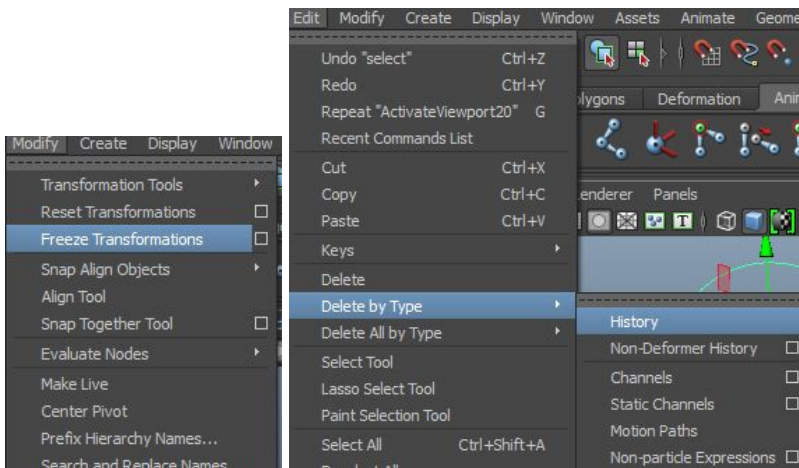


- Center and orient the NURBs circle around the joint you want to control

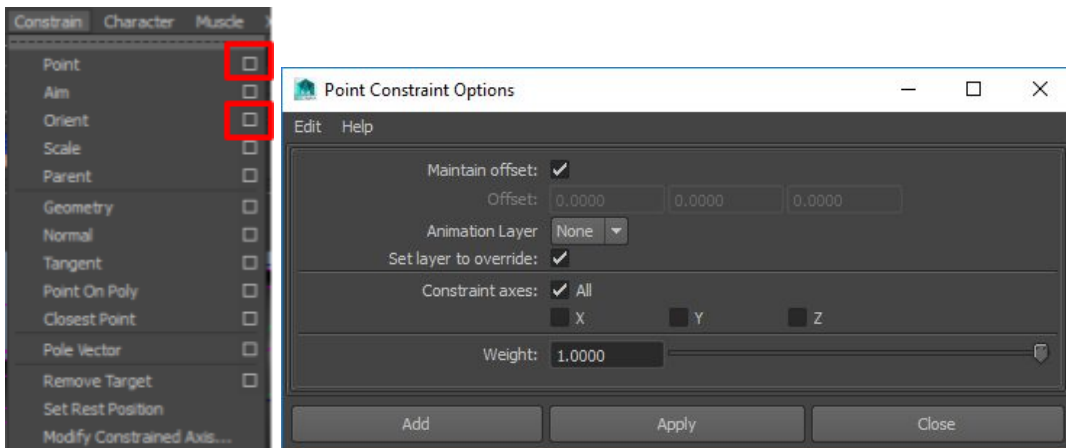


HINT: Use snap-to-point tool (hold V before moving) to position your circle exactly in the center of the joint.

- Freeze the NURBs circle's transforms + delete it's history



- Constrain the point and orientation of your NURBs circle to the joint (or IK handle)



NOTE: Order is important here. Select your circle first, then select your joint, then add the constraint. The first thing you select is where the constraint values are coming from.


Typically you would want to have Maintain offset checked... this prevents the existing transform/rotation settings on the circle from screwing up the joint that the constraint is being placed on. We don't really need this because we've already freezed transforms, but it usually doesn't hurt to have this one.

Constraint axes defines what axes your constraint is for. For example, setting it to X only means that you can only translate/rotate on the X-axis.

NOTE: The constraints that you're setting here through the menus, you can set graphically via the node editor. See the section on node editor for more information.

HINT: It's typical to have one super control where all the actual rig controls are parented to this super control. You can use the super control to easily move around the entire rig in its current state. If you don't have this, you would have to multiselect all the controls individually and move them around together.

In the example below, the big circle on the bottom is the parent of all the other controls.



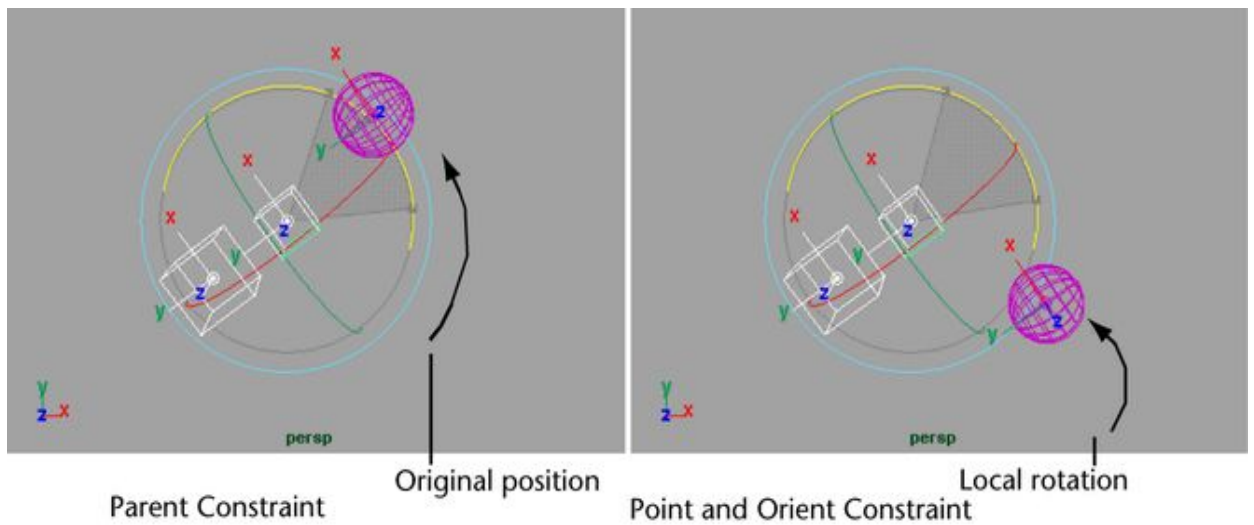
HINT: In addition to point/orient constraints, another useful constraint is the parent constraint. The following is from <http://help.autodesk.com/cloudhelp/2016/ENU/Maya/files/GUID-29785337-D109-48C5-AFC4-8A7A1D0C246F.htm...>

With a parent constraint, you can relate the position—translation and rotation—of one object to another, so that they behave as if part of a parent-child relationship that has multiple target parents. An object's movement can also be constrained by the average position of multiple objects.

When a parent constraint is applied to an object, the constrained object does not

become part of the constraining object's hierarchy or group, but remains independent and behaves as if it is the child of its targets. The constraining object is also known as the target object.

An object with a parent constraint does not behave the same as an object with a point and orient constraint. When a Parent constraint is used, rotating the target object(s) affects the constrained object's rotation along the world axis. When a Point and Orient constraint are used, rotating the target object(s) affects the constrained object's rotation along its local axis. This is shown in the following figure.



Maya 2015 Animation

Timeline

[Adding Keyframes](#)

[Deleting Keyframes](#)

[Copying/Pasting Keyframes](#)

[Moving/Scaling Keyframes](#)

Dope Sheet

[Moving Keyframes](#)

[Scaling Keyframes](#)

[Cut/Copy/Paste/Delete Keyframes](#)

Graph Editor

[Changing Keyframe Curves](#)

[Moving Keyframes](#)

[Scaling Keyframes](#)


[Cut/Copy/Paste Keyframes](#)

[Delete Keyframes](#)

Timeline


The timeline tool allows you to set keyframes on certain properties of your object (e.g. translation or scale). Once you have those keyframes set, the object will transition between keyframe values as you move between the keyframes.

You don't need to open up the tool, it's visible by default on the lower portion of Maya. This is the first place you'd go to for roughing out an animation of your model.



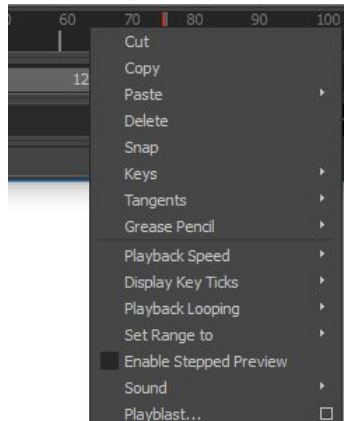
Timeline

The timeline allows you to scrub through your animation, set animation length / playback length, and see where keyframes are set on your selected object (where the red lines are in the timeline).




To scrub through your animation, click-and-hold LMB in the timeline and drag. The amount of frames exposed in the timeline is set by the playback range. As you move in the timeline, the current frame will update.

NOTE: If you right-click on the timeline, you get a context menu that'll give you some handy options such as playback speed and how to loop the animation.



Playback Controls

Playback controls allow you to play your animation as well as quickly navigate between your frames.




You can play the animation forward/backward, step through the animation by frame/keyframe, and easily jump to the start/end of your playback range.

One thing to remember is that your controls are limited to your playback range. If you have an animation of 200 frames but your playback range is only set to the first 60 frames, these controls will only take into account the first 60 (e.g. playing the animation will only play the first 60 frames).

Help Buttons

The first button enabled the “auto keyframe” option. This is described further in sections below.

The second button opens the animation preferences option. Many of these options are available by right-clicking the timeline, but some are not.



Adding Keyframes


NOTE: Keep in mind that the timeline only shows you the keyframes for the selected object. If you select away from the object, you won't see where your keyframes are anymore. The dope sheet will help with this (described later on).

NOTE: Make sure to delete history and freeze transforms on your object before doing any animation. History and existing transforms may cause instability with your animation. It'll definitely make things more confusing if you ever end up using the graph editor (graph editor described in a later section).

You add keyframes on select channels of an object (channels are the properties that show up in the channel box). For example, you can add keyframes for translation, or translation+rotation, or translation+scale, or some other combination of channels.

To add a keyframe...

1. Select your object
2. Select the frame you want to keyframe in the timeline
3. Select the channels you want to keyframe, right-click and select Key Selected




Once you do this, the channels you key'd will go red in the channel box + a red line will show in the timeline to indicate that a keyframe is there.

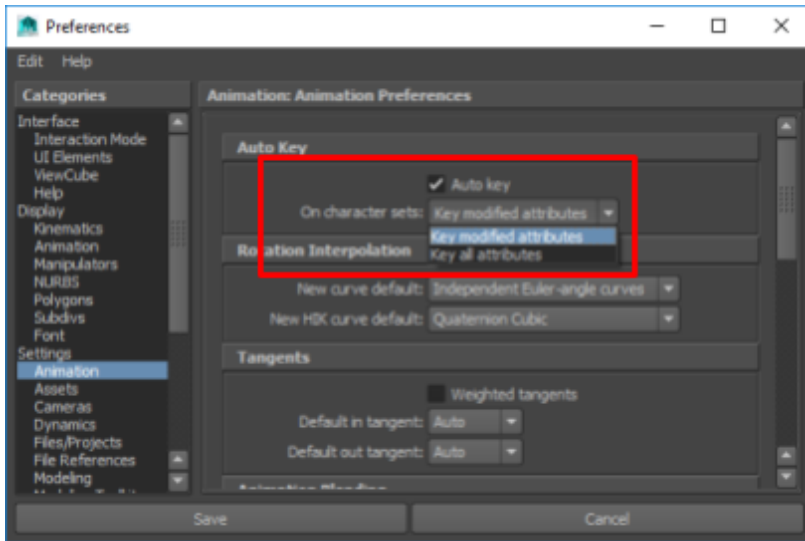


You can now navigate to another frame and repeat the exact same process to set another keyframe for the same object. When you go to a new frame and set the new channel value, you still have to right-click on the channel and select Key Selectable.

You may be tempted to think that you can just change the channel value and it'll automatically add a keyframe -- that won't happen unless you have the auto keyframe option toggled on.



NOTE: If you're having trouble with the auto keyframe option keying up EVERY channel instead of the specific channels you're animating on, go into the animation preferences and check the Auto key drop down.



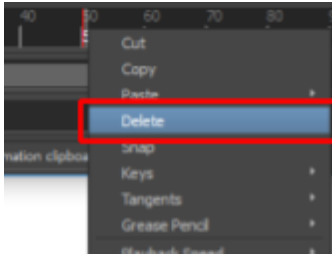
NOTE: The option underneath Key Selected is Key All Selectable. Rarely ever is this what you want. It will literally key every channel in the channel box, even boolean values like visibility. You should only key up on channels that you know you're going to be changing. You can always come back and add more channels if you need to.

If you ever do need it, the hotkey for Key All Selectable is S.

Deleting Keyframes

NOTE: Keep in mind that the timeline only shows you the keyframes for the selected object.

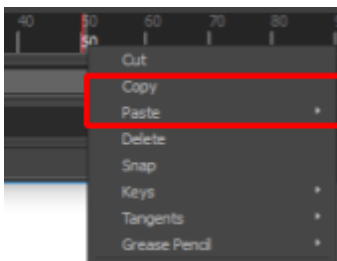
You can delete a keyframe by navigating to it in the timeline, right-clicking on it, and choosing Delete.




Copying/Pasting Keyframes

NOTE: Keep in mind that the timeline only shows you the keyframes for the selected object.

You can copy a keyframe by navigating to it in the timeline, right-clicking on it, and choosing Copy. Then, navigate to the frame you want to paste to and right-click and choose Paste.



Another way to copy over a frame to a keyframe is to go into the timeline and MMB click-and-drag the frame to copy to the frame to paste to. Once done, go into the channel box and key the selected frames just as you would do if you're adding a keyframe.



NOTE: If you use the MMB click-and-drag approach and have auto keyframe enabled, it still won't automatically pick up. For this specific case, you'll still need to go to the channel box to "key select" the relevant channels.

You can also copy and paste by ranges. See how to select ranges in the subsection on moving/scaling below.

Moving/Scaling Keyframes

NOTE: Keep in mind that the timeline only shows you the keyframes for the selected object.

You can move and scale keyframes in the timeline by Holding Shift + LMB click-dragging the area what you want to move/scale. The area you select will be highlighted red.

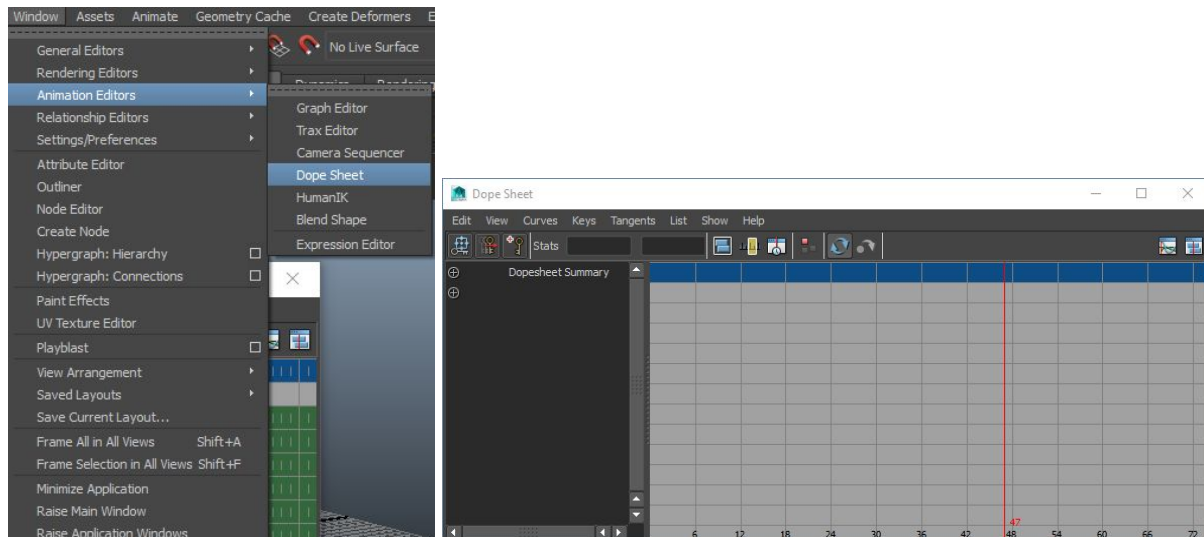
The outer arrows will allow you to scale, while the inner arrows will allow you to move.



Dope Sheet


Part of the problem with using the timeline is that you can't really view/modify the animations for multiple objects individually. The dope sheet fixes this by giving us a stacked view of the timelines for each of the objects we select.

You can open the dope sheet by going to Window -> Animation Editors -> Dope Sheet.

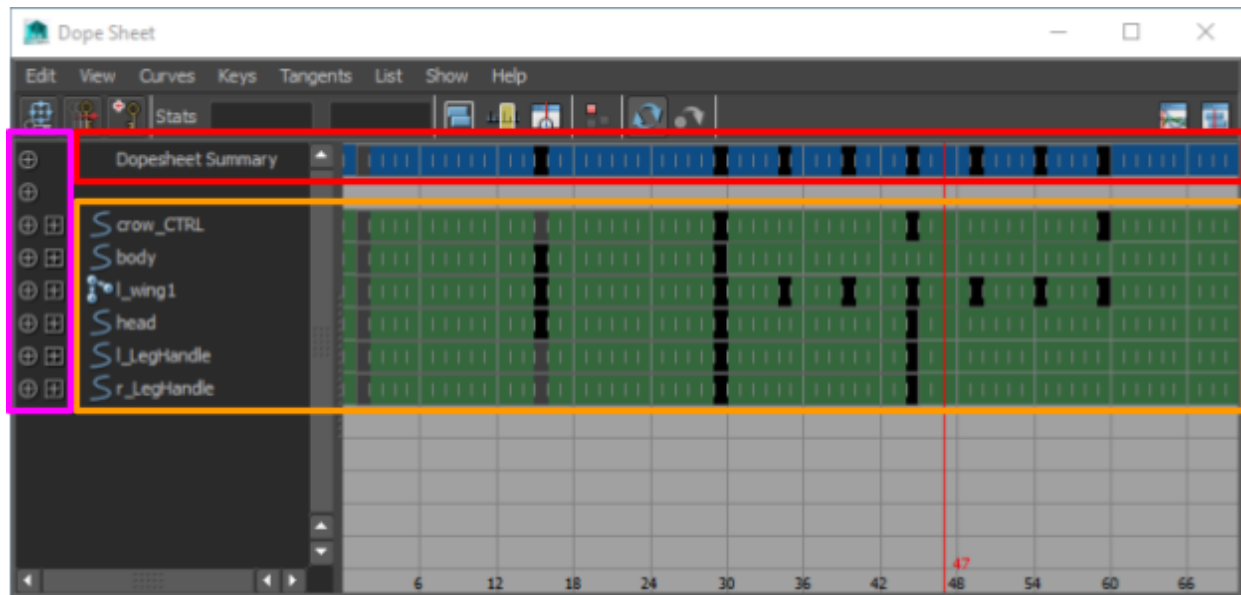


Before the dope sheet shows anything useful, you need to multi-select the various objects you've applied animation to. These could be: individual joints, rig controls, etc..

You can do this by directly selecting in the viewport, or you can go into the hypergraph and select the nodes which are slanted (slanted nodes = animated nodes).



As you select the nodes, they should show up in the dope sheet...



NOTE: Note the red line. It corresponds to where you are in the timeline. There doesn't seem to be a way to move this in the dopesheet, but if you move it in the timeline it'll sync to the dopesheet.

Selected


These are the items we've selected. The dopesheet shows frames they have keyframes on.

Summary


Summary is an aggregated view of all of our keyframes. As long as one of the objects we've selected has a keyframe on at some position, it'll show up on the summary as well (in that same position).

Expanders

The circular expander shows you which properties of the object are keyframed on.



The square expander allows you to drill down the hierarchy of the object.



NOTE: Standard viewport camera mechanics apply to the dopesheet....

- Alt+MMB to pan
- Alt+RMB (or scrollwheel) to zoom
- LMB click+drag to multi-select marquee (Shift/Ctrl modifiers apply here)


Moving Keyframes

You can select multiple keyframes and move them by first hitting W (same as viewport translate hotkey) and then MMB click+drag.

NOTE: Standard viewport selection mechanics apply here...

- LMB click+drag marquee selects a bunch of keyframes
- Shift + LMB click+drag adds to an existing selection
- Ctrl + LMB click+drag remove from an existing selection

In the following example, the highlighted keyframes were moved 3 frames...




Scaling Keyframes

You can select multiple keyframes and scale them by hitting the R key (same as viewport scaling hotkey). Once you're done scaling, you can hit W.

NOTE: Standard viewport selection mechanics apply here...

- LMB click+drag marquee selects a bunch of keyframes
- Shift + LMB click+drag adds to an existing selection
- Ctrl + LMB click+drag remove from an existing selection

In the following example, the highlighted keyframes were scaled down. Notice how the highlighted items have been framed up. You can scale by dragging one of the corners...

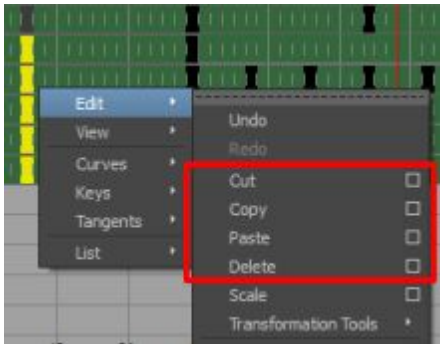


Cut/Copy/Paste/Delete Keyframes

You can select multiple keyframes and copy/paste/delete them by right-clicking and going to the Edit menu.

NOTE: Standard viewport selection mechanics apply here...

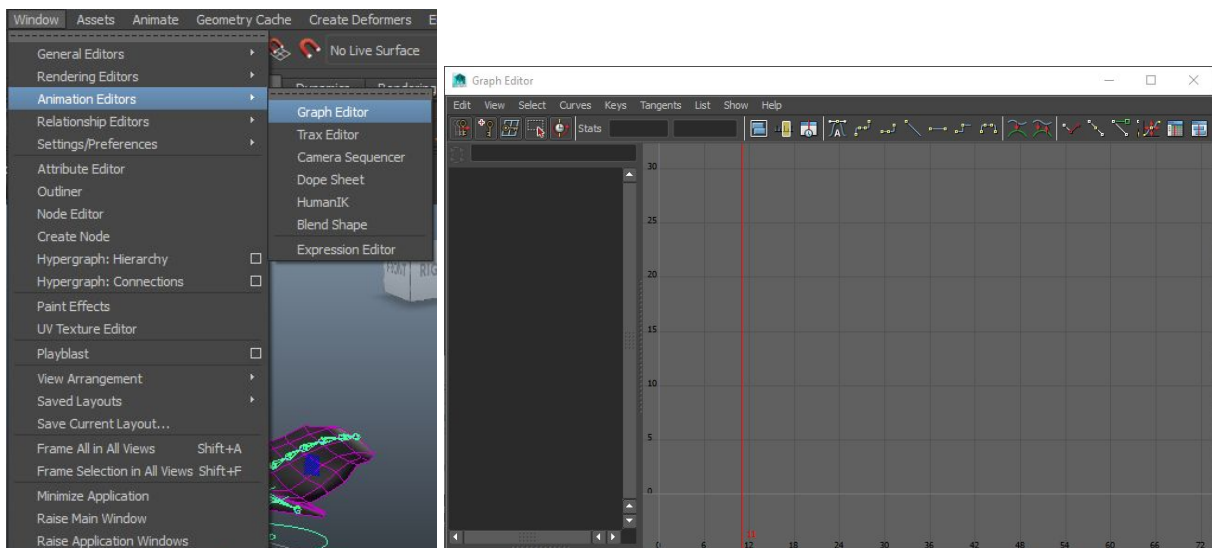
- LMB click+drag marquee selects a bunch of keyframes
- Shift + LMB click+drag adds to an existing selection
- Ctrl + LMB click+drag remove from an existing selection



Graph Editor

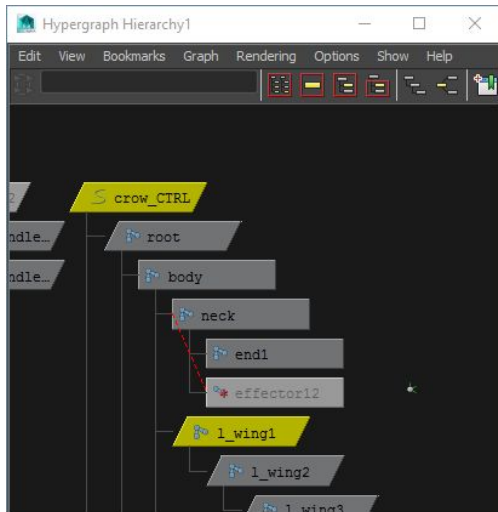
The graph editor is like the dope sheet, but instead of just showing you where the keyframes are, you get the ability to change how keyframes interpolate between each other.

You can open the dope sheet by going to Window -> Animation Editors -> Graph Editor.

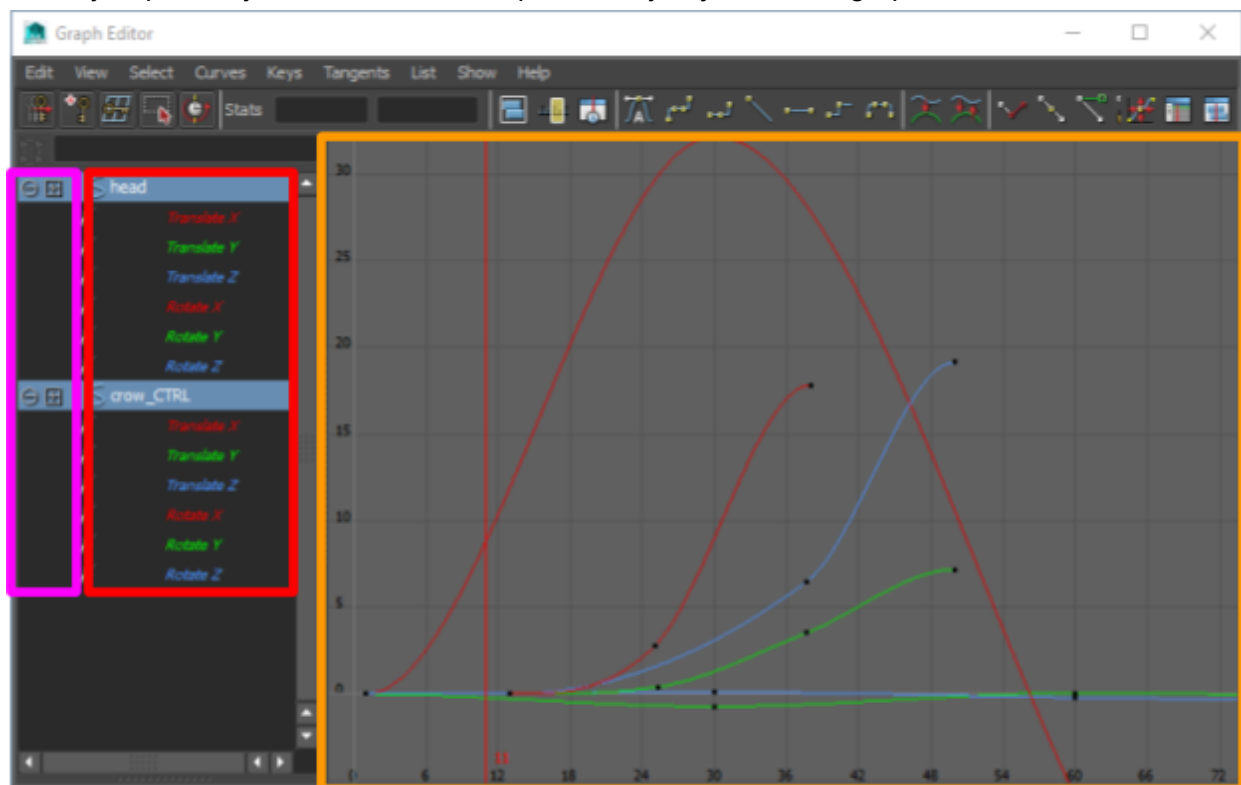


Before the graph editor shows anything useful, you need to select one or more objects that you've applied animation to. These could be: individual joints, rig controls, etc..

You can do this by directly selecting in the viewport, or you can go into the hypergraph and select the nodes which are slanted (slanted nodes = animated nodes).



As you select the nodes, they should show up in the dope sheet. However, unlike the dope sheet, you probably don't want to load up too many objects in the graph editor.



NOTE: Note the red line. It corresponds to where you are in the timeline. There doesn't seem to be a way to move this in the graph editor, but if you move it in the timeline it'll sync to the dopesheet.

Selected


These are the items we've selected. You can select multiple or select individual ones. The plots will show for what you have selected here.

Graph

Contains the plots for the stuff you have selected.

- * Black dots on the plots represent keyframes
- * X-axis is the frame number
- * Y-axis is the value of whatever channel it is that's being plotted

You can select one or more keyframes at the same time. When you do, the dot will turn from black to yellow and the plot will change color to white (these indicate selection).



As you can see above, when keyframes are selected, handles show up that let you manipulate the curve via the tangent (the handles represent the tangent). You can select these tangents just as you would anything else.



In the example above, the right-side tangent was selected.

More on tangents later.

Expanders

This is exactly like the dopesheet...

The circular expander shows you which properties of the object are keyframed on.

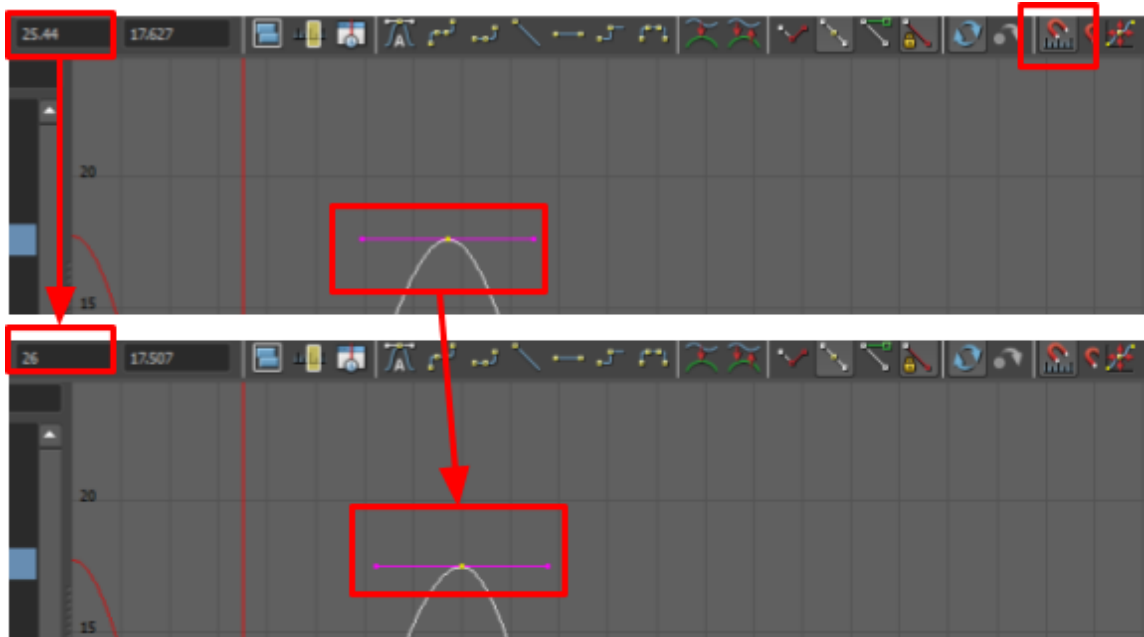
The square expander allows you to drill down the hierarchy of the object.

NOTE: Standard viewport camera mechanics apply to the graph editor....

- Alt+MMB to pan
- Alt+RMB (or scrollwheel) to zoom


- LMB click+drag to multi-select marquee (Shift/Ctrl modifiers apply here)

NOTE: One bizarre thing you'll notice in the editor is that, initially, your keyframes show up at non-whole numbers (e.g. a keyframe on frame 26 may show up at 25.44). There's no explanation as to why this happens. But, if you want you can use the snap to grid feature in the toolbar to move them to the correct spot (how to move described further in subsections below).



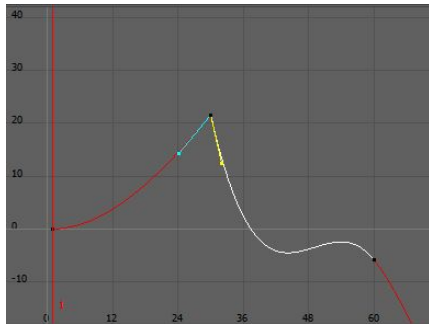
Changing Keyframe Curves

You can change how keyframes interpolate between one another by changing the curvature between them (e.g. transition slower/faster into a keyframe). You can do this by either moving the keyframe up or down, or modifying the tangent handles...



By default, tangent handles are unified. That means that if you move one handle, the other handle will move along with it. You can break this to have more control over how your curve transitions.


To do this, first select your keyframe, then go to the toolbar and click the break tangent button. Once you do this, you'll be able to move the handles independent of each other...



To get back, just click the button to the left of the break tangent button (called unify tangent). If you want you can also select both of your handles and click the auto-tangent button to get your 2 handles back to a straight line...




In addition to the auto-tangent button, there are several other buttons that will change the tangent to some pre-defined shape.




Of particular interest in the highlighted list above is the 3rd one and 5th one...

- 3rd = linear tangent -- breaks the tangent and makes the curve sharp



- 5th = step tangent -- breaks the tangent and turns the curve into a step





Moving Keyframes

You can select multiple keyframes and move them by first hitting W (same as viewport translate hotkey) and then MMB click+drag.

NOTE: Standard viewport selection mechanics apply here...

- LMB click+drag marquee selects a bunch of keyframes
- Shift + LMB click+drag adds to an existing selection
- Ctrl + LMB click+drag remove from an existing selection

In the following example, the highlighted keyframes were moved 3 frames...



NOTE: You can drag past neighbouring keyframes. If you do, it'll just switch positions.

Scaling Keyframes

You can select multiple keyframes and scale them by hitting the R key (same as viewport scaling hotkey) and then MMB click + drag.

NOTE: Standard viewport selection mechanics apply here...

- LMB click+drag marquee selects a bunch of keyframes
- Shift + LMB click+drag adds to an existing selection
- Ctrl + LMB click+drag remove from an existing selection



Some important things to remember with this...

1. WHERE YOU CHOOSE TO DRAG FROM IS WHERE YOU'LL BE SCALING FROM. A crosshair will show up when you start scaling that shows you where you're scaling from.

Keyframes before this point will go left/down.

Keyframes after this point will go right/up.


In the example below, the crosshair is at the second keyframe. That means the keyframe before it will go left and the one after it will go right.



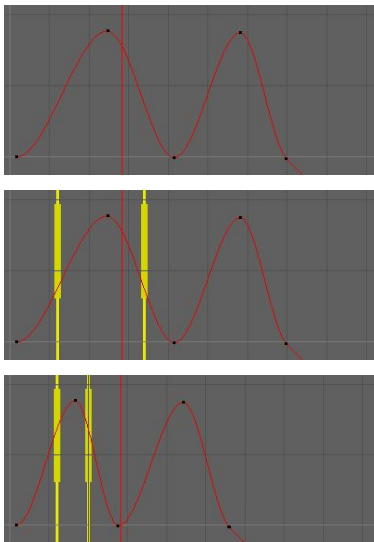
2. Once MMB is down, you can scale horizontally or vertically, not both. How you move the mouse after holding MMB defines which one happens.

Once determined, the mouse cursor will change to indicate which axis you're scaling about.

NOTE: If you want to scale just a portion of the plot without affecting the rest of it. Use the retime tool..



The retime tool lets you isolate a specific region (should have atleast 1 keyframe there) and just scale that region, without affecting anything on the outside of it. In the example below, note how we're isolating the first hump and shrinking it only. Everything before/after stays in the same place.




The tool will prompt you on what to do when you select it. If you need more information, search online.

Cut/Copy/Paste Keyframes

You can select multiple keyframes and cut/copy/paste them by going to the Edit menu. Note that when you paste, it'll PASTE WHERE THE TIMELINE CURRENTLY IS. Adjust the timeline to where you want to paste before pasting.

NOTE: Standard viewport selection mechanics apply here...

- LMB click+drag marquee selects a bunch of keyframes
- Shift + LMB click+drag adds to an existing selection
- Ctrl + LMB click+drag remove from an existing selection



NOTE: Unlike the dope sheet, this isn't exposed in the context menu. You have to go to the main menu for this.

In the following example... I copy my keyframes, then move the timeline to where I want to paste, then paste.




Delete Keyframes

You can select multiple keyframes and delete them by going to the Edit menu.

NOTE: Standard viewport selection mechanics apply here...

- LMB click+drag marquee selects a bunch of keyframes
- Shift + LMB click+drag adds to an existing selection
- Ctrl + LMB click+drag remove from an existing selection



NOTE: Unlike the dope sheet, this isn't exposed in the context menu. You have to go to the main menu for this.

Maya 2015 Rendering

[Interface](#)

[Settings](#)

[Maya Hardware 2.0 Options](#)

[mental ray Options](#)

[Image Quality](#)

[Image Based Lighting](#)


[Final Gather](#)

[Rendering](#)

[Batch Rendering](#)

Interface

You can access the rendering features of Maya by going to the rendering section of the toolbar.



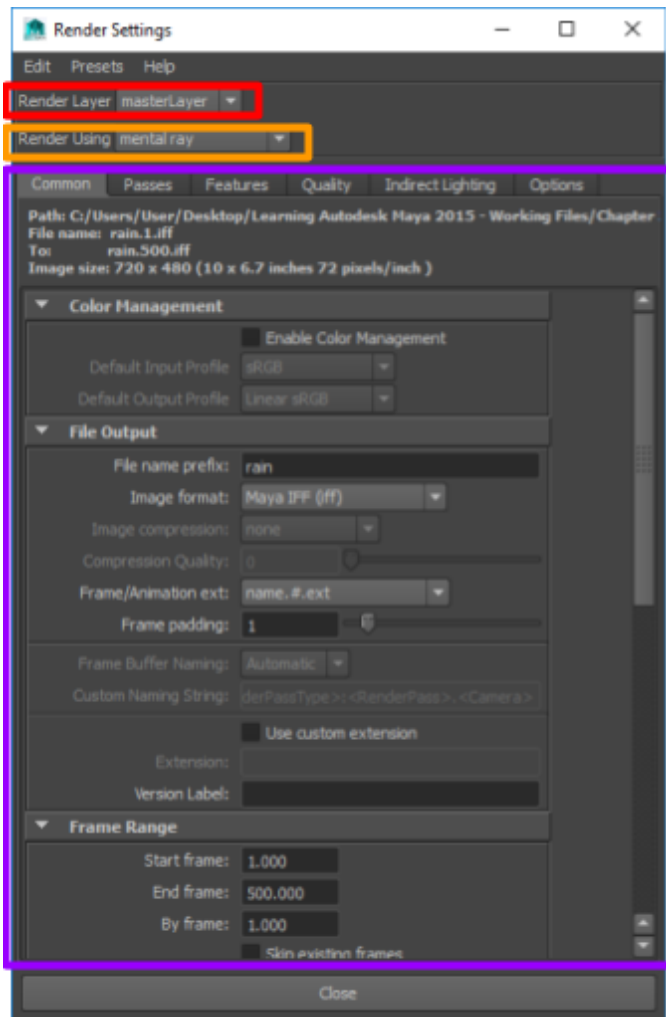
There 4 buttons take you to different areas...

- 1st button opens the rendering window.

If you had a scene rendered already or enabled IPR (explained below), you'd see it here.

- 2nd button renders the current scene.

This implicitly opens the rendering window and begins rendering the scene with whatever render settings were last saved.



Render layer

“You can use Maya’s render layers to separate the objects, lighting, and textures of a scene for easy compositing in other programs such as photoshop or aftereffects.”

<https://www.youtube.com/watch?v=xmRQY4N-J0E>

I don’t know enough about this to discuss it more, but it looks like by default there’s a “masterLayer” that renders everything. If you want, you can get to the render layers section by going to the layer editor in the Channel Box / Layer Editor tab and selecting the Render tab.

Render using

This is the renderer that’ll be used to output your image. The default renderer in maya 2015 is mentalray, which is a software renderer (non-GPU accelerated).

You can also use [Maya Hardware 2.0](#), which is the same renderer provided by the the Viewport 2.0 renderer.

I don't know what the other renderers do.

Options

This area contains the options for the renderer you've chosen. All renderers have a [Common](#) tab which contains basic options: where to output, pixel aspect ratio, which camera to use, etc...

Every tab after that is unique to the renderer chosen.

Note that if you've chosen the Maya Hardware 2.0 renderer, the settings you choose here won't affect your viewports. They'll only be used for the renders you explicitly do.

Maya Hardware 2.0 Options


I don't know enough about this to specifically talk about it. But, you might have to go into the settings and turn on things like screen-space ambient occlusion, anti-aliasing, etc.. to get a half-decent looking render.

mental ray Options

NOTE: On Windows 10 w/ Maya2015, mentalray has lots of problems. Batch rendering doesn't seem to work at all. Single renders sometimes cause access violation exceptions -- doesn't get fixed until Maya is restarted.

Image Quality

You can control quality by going to the [quality](#) tab. There are 2 options that you should look at here...



Quality

Overall image quality. The lower the quality, the more grainy the render will look (but the faster it'll complete).

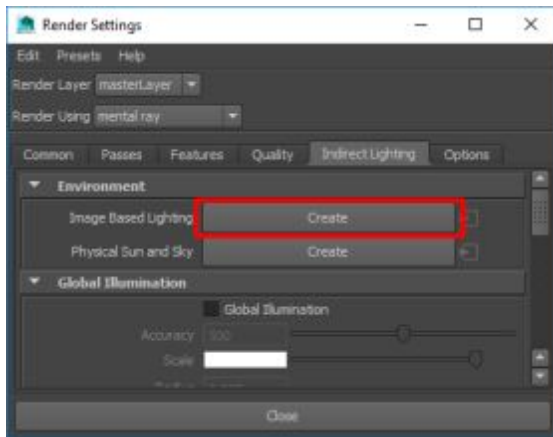
Progressive Mode

Means the same thing as progressive mode for images. Back in 56k modem days, when you downloaded an image off the internet it would first show in super blocky detail and would get higher and higher quality as more of the image was downloaded.


This is similar. Turning this on, you'll get a crappy render at first but as time goes on you'll get a better / more refined render. You can use this to quality spot problems with a render and stop it before it can finish (renders can take a super long time to finish).

Image Based Lighting

One cool feature of mentalray is that it can take an HDR image and use it as a light-source for your scene. To do this, make sure you have mentalray selected as your renderer, go to the indirect lighting tab, and choose to create an image based lighting (IBL) node ...


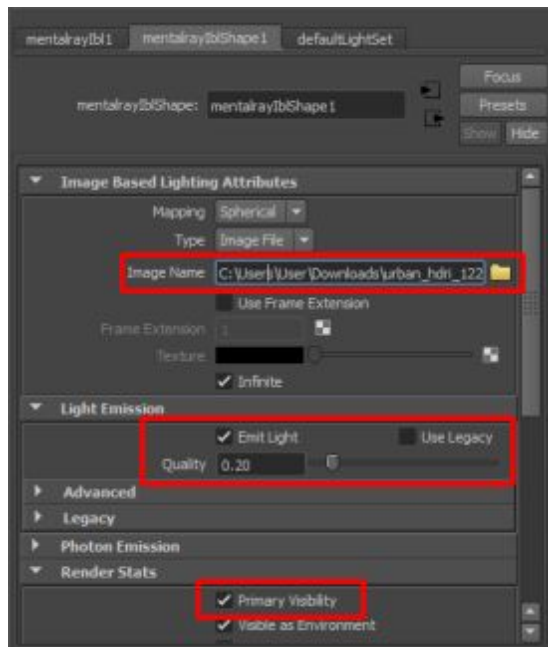


The IBL node is a giant sphere in your scene that encompasses all other objects.




You can select this node and assign an HDR image to it via the Image Name field. Mentalray will scatter light onto the other objects of your scene based on the HDR image you choose (this happens on render only, it isn't shown in the viewports).

Once you select your image, make sure you choose to Emit Light (adjust the quality as desired -- more quality = longer renders). If you don't want the actual image showing up in your render (you only want the light, not the image), turn off Primary Visibility.



NOTE: Note that IBL nodes aren't considered as a standard Maya lighting node. When Maya renders, if it doesn't find any standard Maya lighting nodes in the scene, it'll add in a default light. Having this default light active along with the IBL node's light will screw up the lighting in your scene.

You can turn off the default light in the common tab of the render settings.




Final Gather

Final gather is the raytrace-esque feature in mentalray. Here's the text explaining it from an article online...

Final Gather can be found in just about all of the most common render engines and is an essential step in creating realistic renders that closely mimic how light interacts with objects in the real-world. With that in mind, let's go over Final Gather, how it works and where it can be implemented into your next render. What is Final Gather? In the real world a lot of the light that you see is the direct result of indirect illumination (e.g., light that isn't coming directly from a source like a light bulb). This is the light that bounces off other surfaces to illuminate your surroundings. When creating realistic renders whether it be for an interior or visualization, indirect illumination is something you will want to consider using within your scene. Final Gather is a prominent feature in most rendering engines, including the render engine mental ray which is the main rendering engine for Maya, 3ds Max and Softimage.

To fiddle with the final gather settings, make sure you have mentalray selected as your renderer and go to the indirect lighting tab.



Accuracy

The higher this is, the more 'accurate' the ray tracing will be?

“Number of rays per final gather point cast into the hemisphere, to estimate incident global illumination. Higher values raise the accuracy of the result, which is expensive”

Point Density

Increases the “amounts of points” to be computed. More points = more points of light. The more amount of points we have, the more they can

get into individual nooks-and-crannies in the geometry of our scene.


“Controls the number of final gather points to be computed, performing the full and expensive final gather tracing.”

Point Interpolation Blends between the amount of rays in the scene.

“The number of final gather points to be considered for interpolation at a shading sample during rendering. Higher values smooth the final fathering result at little cost.”

NOTE: You probably need some type of light source for this to work. You can use an image based lighting node or set the camera background color for light. But, neither of these is an actual light source according to Maya. When Maya renders, if it doesn't find any standard Maya lighting nodes in the scene, it'll add in a default light. Having this default light active will screw up the lighting in your scene.

You can turn off the default light in the common tab of the render settings.




Rendering

You can render your scene using either the normal render button (first button) or the interactive photorealistic renderer button.

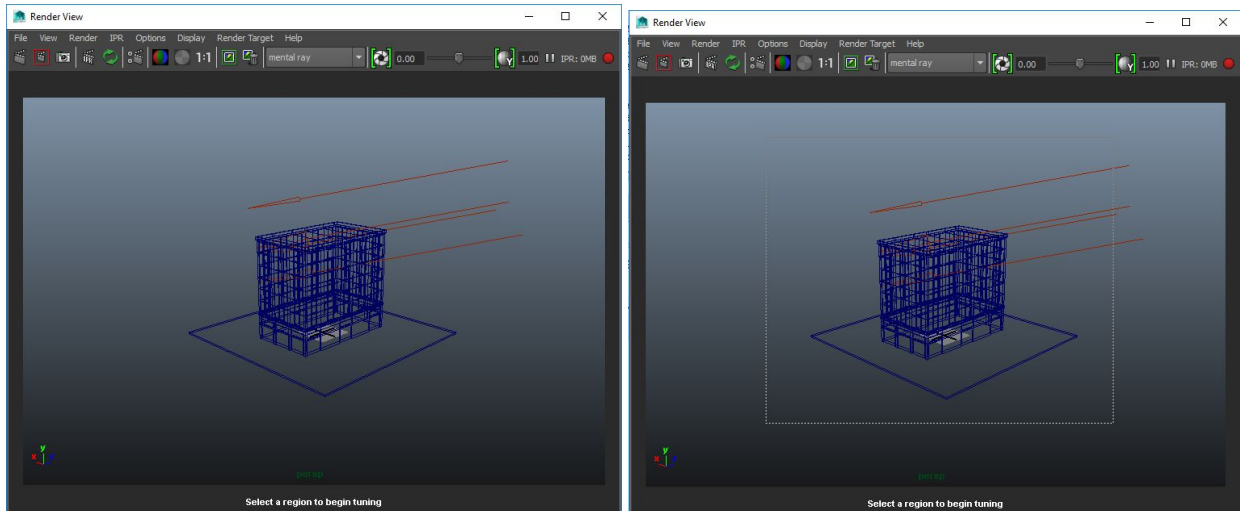


If you choose the first button (normal render), the render will begin immediately with whatever render settings you had set previously.

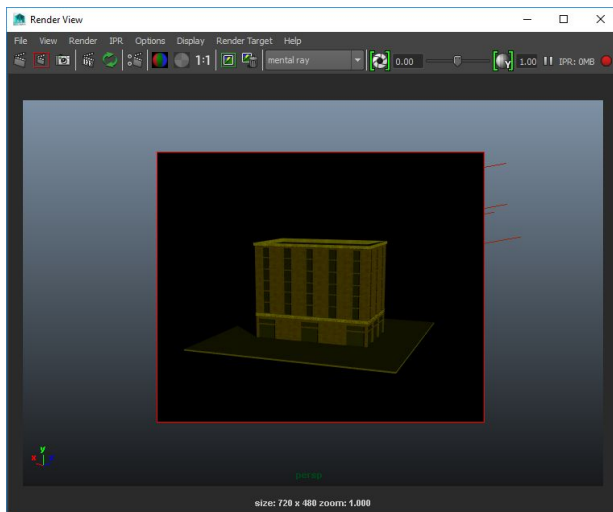


If you choose the second render button (IPR render), the render window will pop up with a wireframe.

You'll have to choose the region most important to by click-and-dragging across the area.



Once you do that, that region will start re-rendering as soon as something in the scene changes.

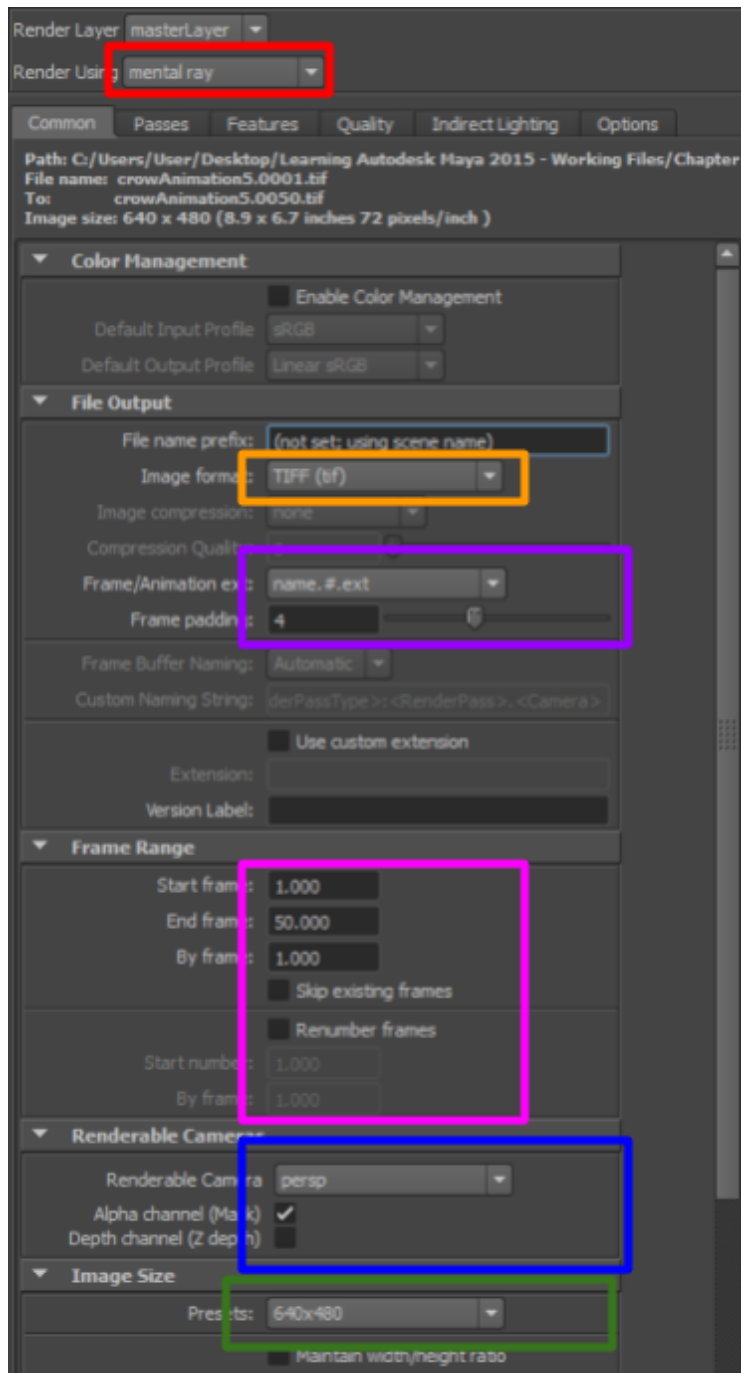


NOTE: You can kill the render at any time by hitting Esc.

HINT: It's a good idea to turn on progressive mode in the render settings. You can quickly spot problems in long running renders and cancel them.

Batch Rendering

Batch rendering is when you want to render multiple frames (when you have an animation). Here are the typical options you would set in the render settings for a batch render.



Render using

You typically want to use the production render rather than the Maya Software/Hardware renderers. The Maya renderers are what are used for Maya's viewports.

Image format

It sounds like the industry format is TIFF.

Naming convention It sounds like the industry convention is name.#.ext, where # is the frame number and padded up to four 0s.

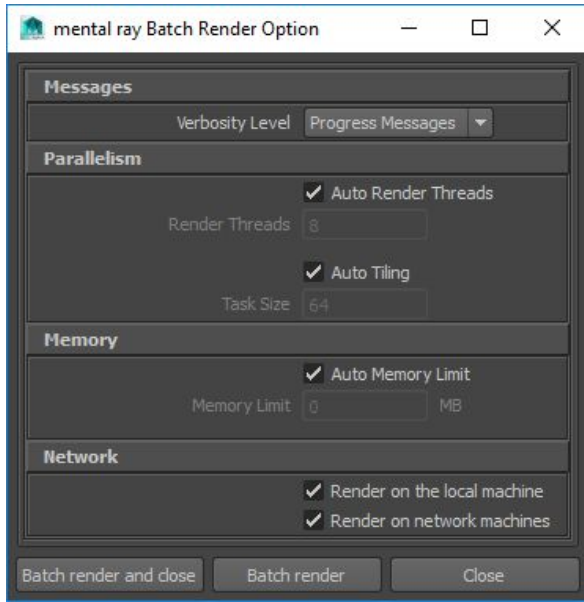
Frame Range You would typically set this to span your entire animation range. If you already have some frames rendered out, you can skip existing frames. If you want to render after already rendered frames, you can choose to renumber frames.

Renderable Camera Choose the camera you want to render with and make sure alpha channel is selected.

Image Size Lots of presets to choose from for image size, including typical print and movie formats.

Once you've setup your render via the render settings, you can switch to the rendering menu set and choose Render -> Batch Render (click the little box on the right to open the options).






When you're ready, start your render. It'll happen in the background. If you have Verbosity Level set to Progress Messages, you can check out the MEL script output to find out how rendering is going.


NOTE: You can kill the render at any time by going to Render -> Cancel Batch Render (make sure you're in the rendering menu set).

Images should get written to the images subdirectory of the current project directory.



NOTE: On Windows 10 w/ Maya2015, mentalray has lots of problems. Batch rendering doesn't seem to work at all.

Once rendering is complete, you can view the batch render using Maya's fcheck software. Open up fcheck (it should be in the start menu where Maya is located) and go to File -> Open Animation and choose any tiff from the render.



Maya 2017 Arnold Renderer

[Introduction](#)

[Setup](#)

[Renderer Setup](#)

[Arnold RenderView Setup](#)

[Arnold Light Manager Setup](#)

[Arnold RenderView](#)

[Arnold Light Attributes](#)

[Color and Intensity](#)

[Illumination Specifics](#)

[Shape, Spread, and Quality](#)

[Shadows and Volumetrics](#)

[Material/Shader Weights](#)

[Arnold Light Types](#)

[Area Light](#)

[Skydome Light](#)

[Mesh Light](#)

[Photometric Light](#)

[Arnold Object Smoothing/Subdivision](#)

[Adaptive Error](#)

[Arnold Standard Material \(aiStandard\)](#)

[Diffuse](#)

[Specular](#)

[Reflection](#)

[Refraction](#)

[Sub-surface scattering](#)

[Emission](#)

[Caustics](#)

[Arnold Skin Shader \(aiSkin\)](#)

[SSS Layers](#)

[Specular and Sheen](#)

[Arnold Environments](#)

[Background](#)

[Sky Shader](#)

[Physical Sky Shader](#)

[RaySwitch Shader](#)

[Atmosphere](#)

[Fog](#)

[Volume Scattering](#)

[Arnold Depth of Field](#)

[Arnold AOVs](#)

[Textures in Arnold](#)

Introduction

Arnold the default renderer in Maya 2017. It's a monte-carlo based raytracer.

The two main reasons to use Arnold are...

1. Industry adoption -- Arnold is used by most vfx/animation studios for rendering their scenes.
2. Easy of use -- Setting up arnold is much easier than other similar renderers (e.g. renderman/mentalray/etc...). This is especially true for artists. You don't have to do things like set up prepasses or cache files for things like shadows. Arnold calculates everything it needs on the fly.

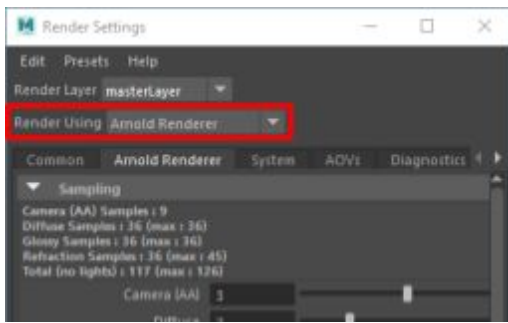
NOTE: What is is a pre-pass and what are cache files used for? Unsure, but a Google search seems to hint that this has something to do with a concept called deferred rendering/shading/lighting. You do multiple passes over the scene to perform calculations for things that will be used in the render, then do the actual render.

NOTE: This isn't a simple raytracer. There's some fancy stuff happening here. See <https://computergraphics.stackexchange.com/a/18> for a overview of the topic.

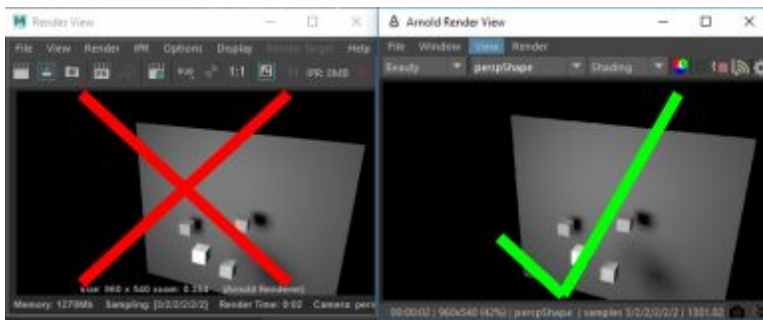
Setup

Before you do anything, you should setup your environment for working with Arnold.

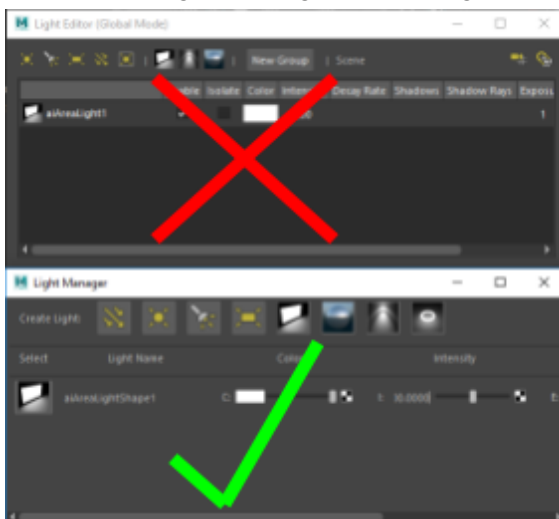
1. Make sure Arnold is set as the renderer (it should be by default).



2. Use Arnold RenderView to view renders (different from Maya's render view).



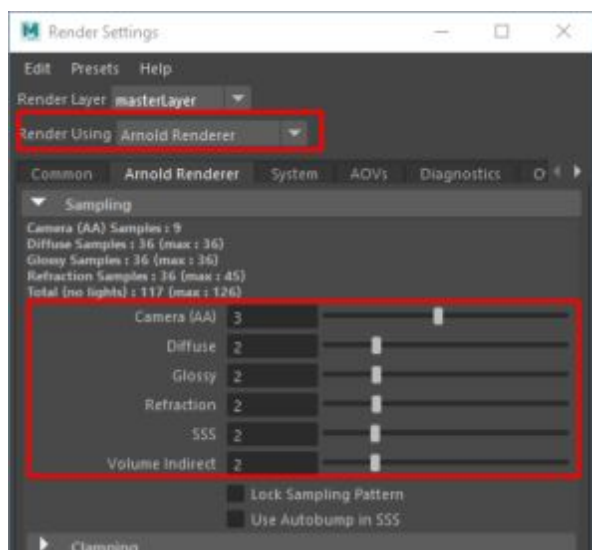
3. Use Arnold Light Manager to view lights (different from Maya's light editor).



Renderer Setup

You can make sure that Arnold is the renderer being used by going to render settings and choosing Arnold Renderer. This should already be the default.





The main quality control setting here is under the sampling section. The defaults here are good for draft renders. If you want a final render, you would bump these numbers up.

The main quality setting here is Camera (AA). A setting of 10 will generate a good quality render. But, keep in mind that the number of samples goes up exponentially as you increase this value. The more samples you have the longer your render will take.

NOTE: Camera (AA) slider goes up to 10, but it can be bumped even higher by typing in a value.

In some cases, you may notice noise/dithering with some parts of your render even though you have a high Camera (AA) value. You can find out which parts of your shader are having issues by inspecting AOVs (AOVs are explained later on).

For example, if you notice that subsurface scattering is having issues, you can bump up the the SSS value instead of bumping up the the main Camera (AA). This ends up being not as computationally expensive.

NOTE: One thing you'll notice about arnold vs other renderers is that there is no default light! If you have no lights in your scene, your render will show up pitch black.


Arnold RenderView Setup

Arnold RenderView is similar to Maya's render view, but streamlined for use with Arnold. You should always use the Arnold RenderView as opposed to the regular Maya Render View.

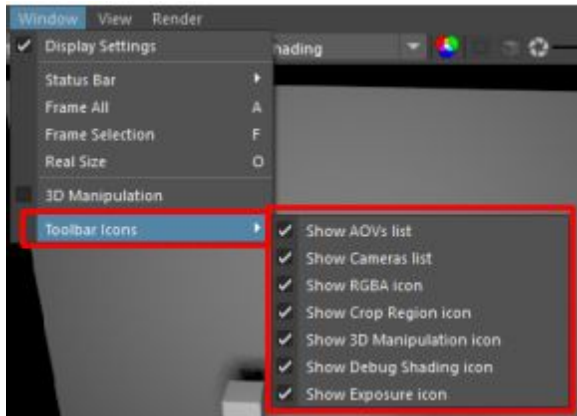
NOTE: Full list of features available here:

<https://support.solidangle.com/display/AFMUG/Arnold+RenderView+Window>


To get to the Arnold RenderView... go to Arnold -> Arnold RenderView (menuset doesn't matter)



Once you're in Arnold RenderView, enable ALL of the toolbar options. You can do so by going to Window -> Toolbar Icons and making sure everything is selected...



Here's a quick overview of the different things in Arnold RenderView. They'll be covered in more detail further on...



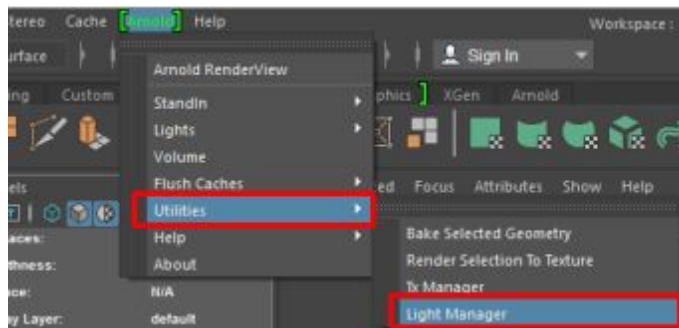
Arnold Light Manager Setup

Arnold Light Manager is similar to Maya 2017's Light Editor, but streamlined for use with Arnold. You should use Arnold's light tool as opposed to the regular Maya Render View.

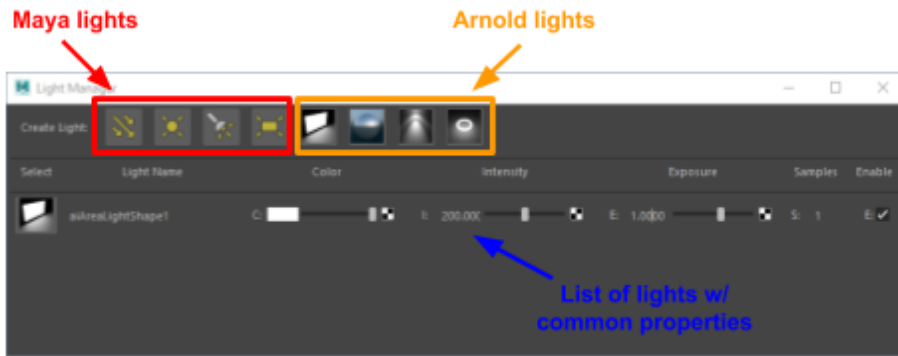
NOTE: Full list of features available here:

<https://support.solidangle.com/display/AFMUG/Light+Manager>

To get to the Arnold Light Manager... go to Arnold -> Utilities -> Light Manager (menuset doesn't matter)

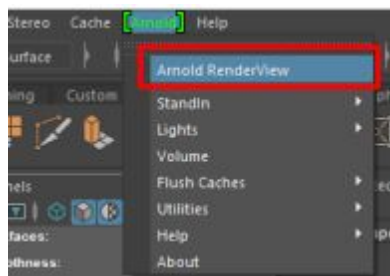


Here's a quick overview of the different things in the light manager...

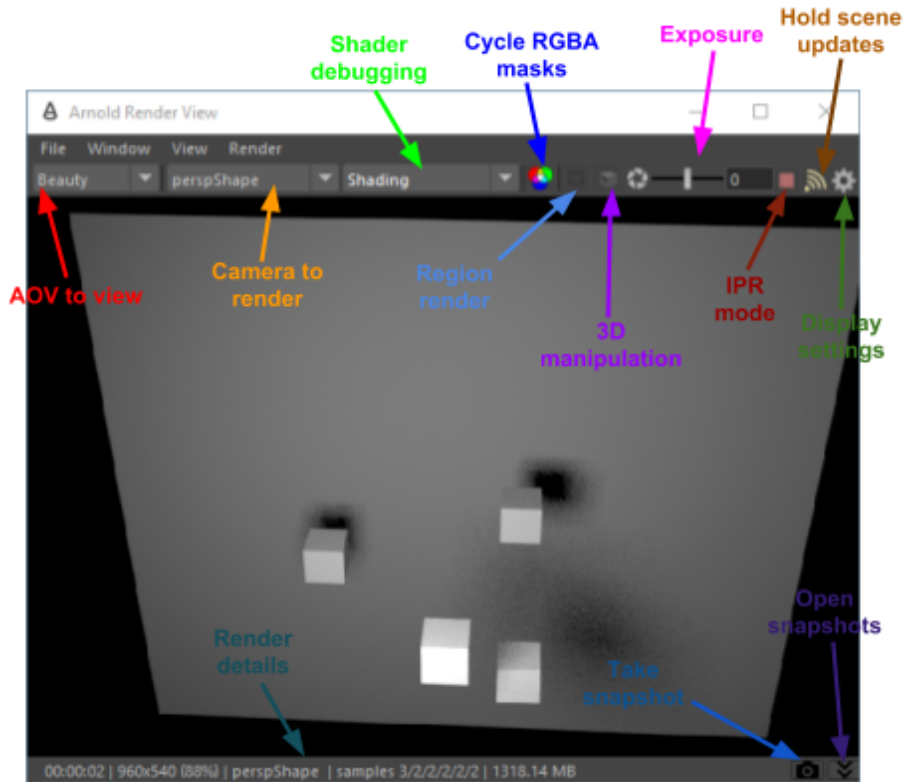


Arnold RenderView

To get to the Arnold RenderView... go to Arnold -> Arnold RenderView (menuset doesn't matter)



NOTE: Once you're in Arnold RenderView, enable ALL of the toolbar options. You can do so by going to Window -> Toolbar Icons and making sure everything is selected...



NOTE: One thing you'll notice about arnold vs other renderers is that there is no default light! If you have no lights in your scene, your render will show up pitch black. Keep this in mind if you ever do a render but it comes out pitch black.


NOTE: Slow renders? Set a lower test resolution by going to View -> Test Resolution

AOV to view

AOVs will be described later on, but if you have AOVs set you can select them in this drop down and view them here. For example, if you wanted to see how a specular pass would look just on its own, you can have that as an AOV and view it directly in this view.

The "Beauty" value is the one selected by default. This is the option for the full render.

HINT: Here's a description of AOVs from the arnold website... AOVs (Arbitrary Output Variables) provide a way to render any arbitrary shading network component into different images. For example, an artist might find it convenient to separate direct and indirect lighting contributions and later recombine them during compositing.



Camera to render


Select the camera you want to render with this dropdown.

Shader debugging

Select different ways of shading the scene. Use this to help with debugging your scene. For example, if you select normal your render will show the normals for the object in your scene, and you can easily spot things like polygons facing the wrong way.

Region render

Allows you to select a region on the render that will continuously get rendered instead of the entire scene. Use this when renders are really taking a long time and you only want to focus on one specific area.



If you mouse into the box, you'll see crosshairs in the center. Click-and-drag the crosshairs to move the box around, or click-and-drag the corners to expand/shrink the box.

Cycle RGBA masks

Pick individual channels in the render: red, green, blue, or alpha.

3D manipulation

If this is selected, you can pan, rotate, and zoom via the Arnold Render View just like you normally do when you're in a viewport (Alt+MMB/LMB/RMB). You can also select objects through the Arnold Render View by clicking on it (but you can't do fancy selections like marquee or multi-select).

Exposure

Expose the display image up or down, in f-stops.

IPR mode

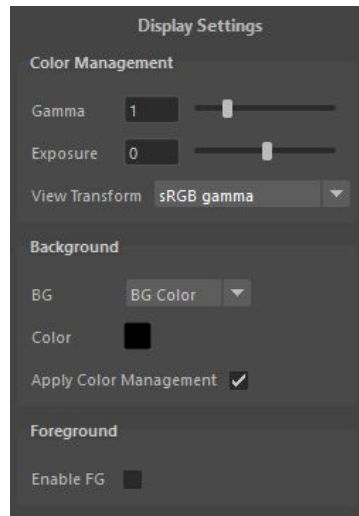
This enables/disables automatic updates whenever the scene is changed. The user will then manually trigger a new render by calling a Render.

NOTE: A lot of times Arnold won't pick up new things when it does scene updates like this. You'll need to go to Render -> Update Full Scene (Ctrl+U shortcut).

Hold scene updates If IPR mode is enabled, enabling this will stop renders from happening on a change.

If IPR mode is disabled, this turns into a button that you can click to render the scene.

Display settings Opens up a panel on the side that lets you change various settings with your display...



Gamma is self-explanatory.

I don't know what View Transform does.

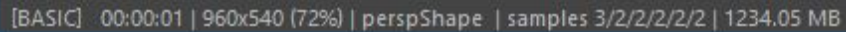
Exposure mirrors the exposure option in the toolbar -- expose the display image up or down, in f-stops.

Background and Foreground are self-explanatory.


NOTE: Color management is not on by default as it is when you use the normal Maya Renderer. You have to manually set things like gamma here in the display settings.

Render details

This is from the Arnold website. It does a better job of explaining what's here than I can...



Just below the render buffer, a Status Bar provides information about the render in progress. When a previously stored snapshot is displayed, information about this previous render is displayed.

- If Debug shading mode is used, the debug mode is shown (OCCLUSION, UV, NORMAL, etc).
- If Isolate Selected mode is used, the name of the selected node is shown. A 'lock selection' option  is available for isolate selected, which ignores selection changes.
- "Rendering..." appears when a render is in progress.
- Render time appears when a render is finished. Note that this render time corresponds to the last AA step, which corresponds approximately to the render time it would take to render this in batch mode.
- Resolution: width x height.
- Zoom factor: 50%, 1:1, 200%, etc.
- Camera name: displays the name of the camera that is being rendered.
- Sampling info: AA samples / GI Diffuse Samples / GI Glossy Samples / GI Refraction Samples / GI SSS samples / GI Volume Samples.
- Memory usage: displays the memory used for the current render.

Take snapshot

Takes a snapshot of the current render and saves it in memory.

NOTE: When you view the snapshot again, you'll get the rendering details in the status bar.

Open snapshots

Open the snapshots panel on the bottom.



Click any snapshot to view. When you're viewing a snapshot, the take snapshot button turns into an eye instead of a camera. Click the eye to go back to viewing the render.

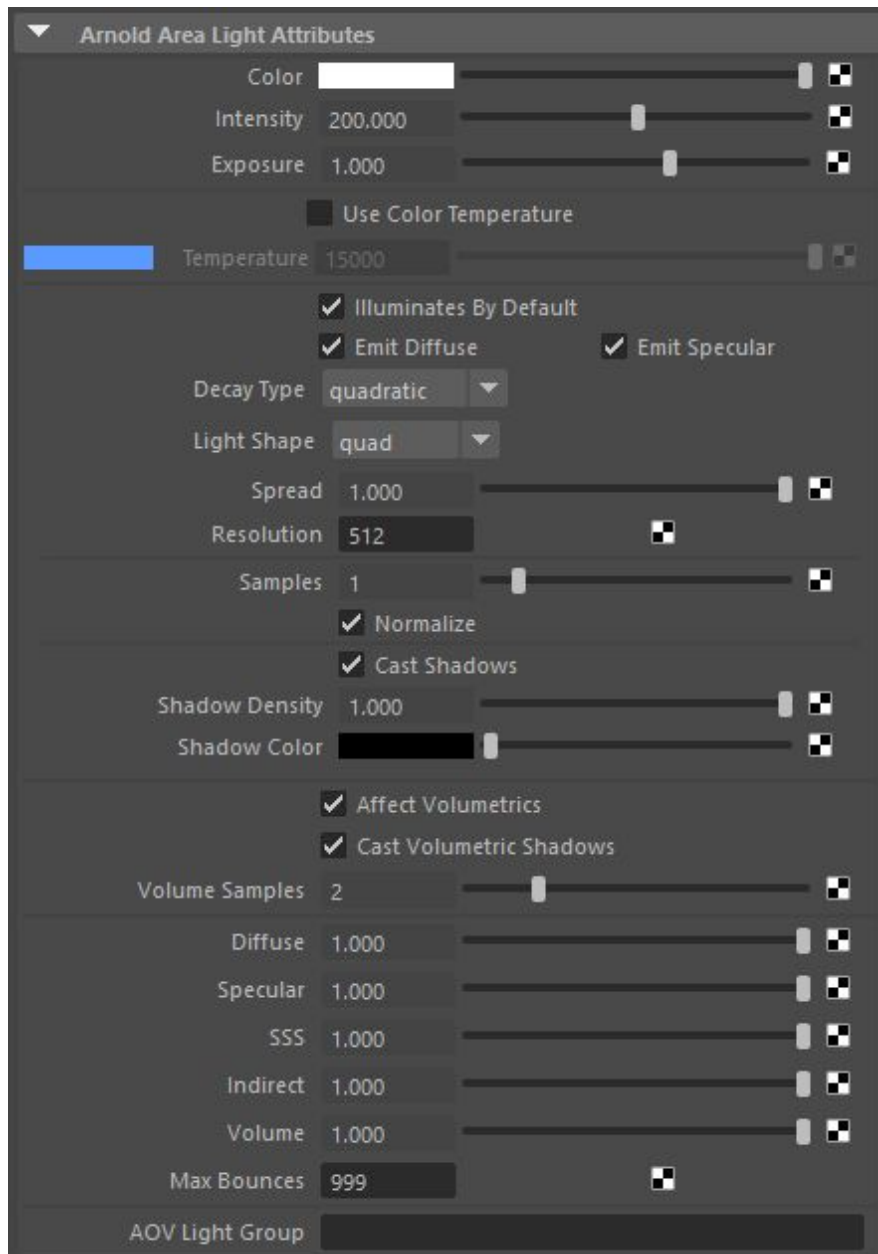


NOTE: When you view a snapshot, you'll get the rendering details for that snapshot in the status bar.


Arnold Light Attributes

These settings apply to all types of Arnold lights (more or less). Some of these settings can also be set on Maya lights -- you'll need to go to the Arnold section within the attribute editor to find them.

The following snapshot is for an Arnold area light. But, most of these settings can be found in pretty much every other light Maya light.



Color and Intensity



Color Color of your light.


Intensity Intensity of your light. This works the same way as a regular Maya light. The higher your intensity is, the more light will be given off.

Exposure Intensity of your light. This is different from the above setting in that it's measured in f-stops (just like real photography). It's designed to help you interact with real cinematographers. For example, the cinematographer may come and ask the artist to increase something by half a stop -- you can do that directly here in this setting.

This is exponential... each time you go up by 1, it doubles the intensity of your light.

NOTE: This is linked to the intensity in that the exposure uses the current intensity value as the basis for its setup.


NOTE: Unlike intensity, exposure won't show up in your viewport. It may render properly but if you have lighting turned on in the viewport, the object will not actually get lit in the viewport.



Temperature Changes the color based on Kelvin (e.g. a cool color or a warm color).

NOTE: If you set this, it'll override the color attribute.

Illumination Specifics




Illuminated by Def. Disable this if you don't want your light to illuminate everything / if you want your light to illuminate specific things. You'll need to manually link your lights if you do this. Look up Maya light linking for more information.

Emit Diffuse Allow the light to affect the material's diffuse component.

Emit Specular Allow the light to affect the material's specular component.

NOTE: It looks like this only applies if the rays coming from a light DIRECTLY hit the material. If the rays are INDIRECT (e.g. they bounce off something else before hitting the material in question), it will apply to the material's specular component.

In the example below, the only light in the scene has emit specular turned off. Notice how the specular highlight isn't showing when emit specular is turned off, but the reflections of the boxes (specular highlight of boxes) are still showing.



Decay Type The falloff for a light. Quadratic is the default and the most realistic (uses the inverse square law -- look it up if you want).


Recommended you don't change this from quadratic.

Shape, Spread, and Quality



Light Shape The shape of the light. Can be quad, disk or cylinder.


NOTE: This applies to area lights only? The shape of the light shows up in specular highlights.



Spread

Emits light focused in the direction along the normal. The default spread value of 1 gives diffuse emission, while lower values focus the light more, until it becomes a laser like beam at value 0. Low spread values can be noisier than the default high spread, so be careful when using them.

The following example is spread of 1.0 vs 0.5 vs 0.25 vs 0.1...



NOTE: This wasn't explained during the course. The description above is from the Arnold documentation.

Resolution

This has to do with the topic of importance sampling. This isn't fully described, but from what it sounds like importance sampling has something to do with the images you apply to your lighting (e.g. if you loaded up a texture for the color of your light). It splits the image up and looks for the brightest part and sample those more vs the dark parts where it won't sample them as much.

This is the resolution you use to control that.

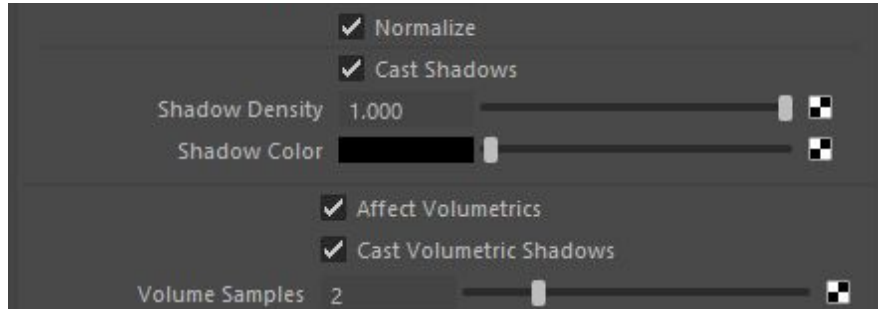
Samples

The quality of the light / the quality of the shadows. This is an override for the global samples setting.

The higher this value is, the longer things will take to render.

NOTE: Also check out the Volumetric Samples in the volume section.


Shadows and Volumetrics



Normalize

If this is turned on, the scale of the light controls how soft the shadows are rather than the intensity of the light (Applies to area lights only? Scaling does nothing with photometric lights).

That means that if this is on and you scale your light up, you will get softer shadows, NOT more light. The example below shows a scaled up light. The first render is with normalize on while the second is with normalize off.



Cast Shadows

If turned off, the light won't cast any shadows.

Shadow Density

The shadow strength.

It looks like if this is < 1.0 , some light will be let through.

Shadow Color

The color of the shadows being cast.

Affect Volumetrics

Lights up volumes in your scene (e.g. haze/fog/god rays/etc..).

A lot of the time you'll have filler lights to light up parts of your scene that are too dark, but you don't want those lights to effect the volumes. This setting helps with that -- volumetric lighting is expensive in terms of

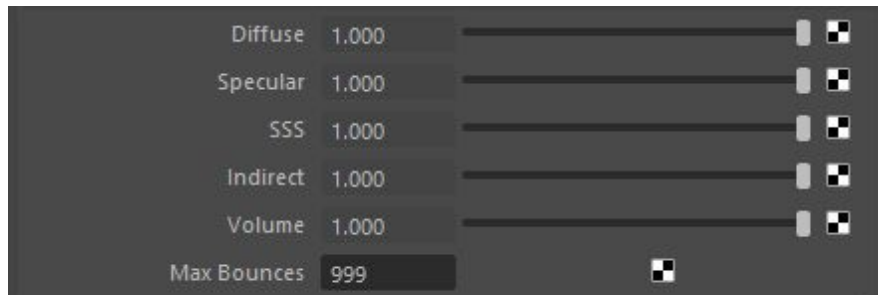
render time.

NOTE: Can't you just do this with light linking? Maybe this is here because light linking is a Maya-specific thing?

Volumetric Shadows Whether or not the light cast shadows in the volumetric light.

Volume Samples Similar to samples, but specific to volumes.

Material/Shader Weights



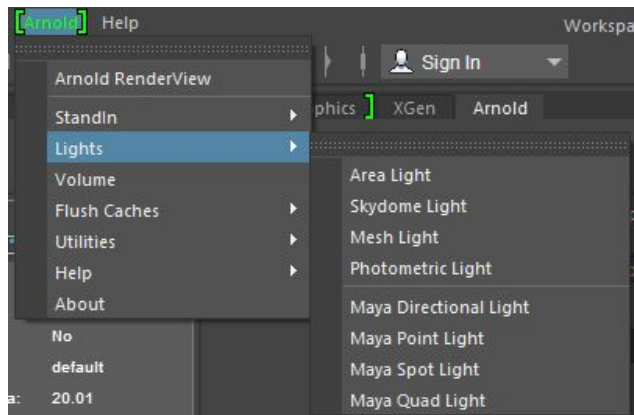
Specifies how much of the light effects different parts of a shader. For example, you can turn down the weight this light has on the specular of a material -- turning it all the way to 0 would be like unchecking the Emit Specular option.

Max Bounces controls the maximum number of times a light can bounce off a surface. For example if Max Bounces was set to 0, none of the rays would be bouncing off the polygon they hit.

Arnold Light Types

Arnold comes with 4 light types: Area lights, Skydome lights, Mesh lights, and photometric lights.



To add a light, you can use the Arnold shelf or you can go to Arnold -> Lights -> (pick light here).



Area Light

An area light is a basic light. It can be shaped as a quad, disk, or cylinder.

To add, go to Arnold -> Lights -> Area Light.




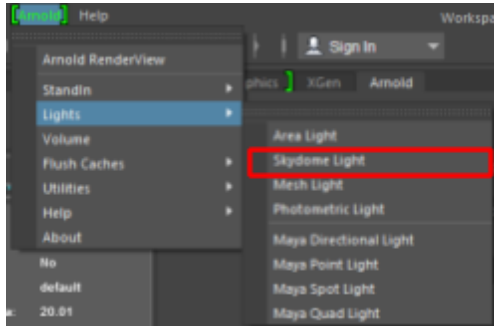
Note that, if normalize is on, the scale of the light controls how soft the shadows are. Make the light smaller for harder shadows, or larger for softer shadows.

If normalize is off, the scale controls the intensity.

Skydome Light

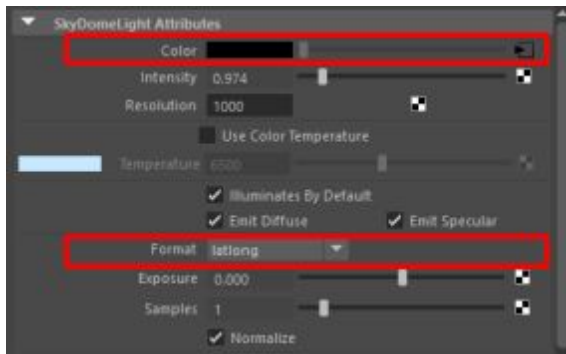
A skydome light is a giant light that goes all around your scene in a giant sphere/dome shape. Typically used for ambient/fill lighting?

To add, go to Arnold -> Lights -> Skydome Light.



NOTE: This seems very similar to a mentalray IBL node (image-based lighting node).

Typically, what you want to do with a skydome light is to apply an image texture to the color attribute (instead of assigning a solid color). The type of image you should be using is either a latlong, angular, or mirrored_ball.



NOTE: Resolution has to do with importance sampling. This increases the accuracy of light coming from the image. See the resolution description of the lighting attributes section for more information. You typically want to bump this up for interior scenes (or so the course video says).

NOTE: These image formats are discussed further on <http://www.hdrlabs.com/tutorials/index.html>. The text is replicated below.

What are the differences between a light probe, spherical image, LatLong, vertical cross, cubic (etc) versions of HDRs?

Depending on what application you are using, the names of the shapes of the HDR will change.

Most 3D programs want/need an image that looks like an unwrapped world map. This file can be called "Equirectangular", "Spherical", "LatLong", and also a "Latitude Longitude" file. These are called different names, but they are all exactly the same!

Vertical Cross, Horizontal Cross, and Cubic HDRs... these are all just six 90-degree views of a scene or panorama. They just happened to be placed together in a certain shape (vertical or horizontal cross) or saved out as six separate cube faces as six different files. Retouching is pretty easy with these as long as you do not try to retouch any of the edge pixels. If you need to do some cloning or color correction across the seams, then convert the image to a spherical image and work on that. Then convert back to a cubic format if needed.

"Light Probes" are full 360x180-degree spherical (or LatLong) HDR files that have been converted into a ball shape. Visually, it appears that they are chrome balls, but in reality they have a mathematically much simpler distortion and they may or may not have been created and captured that way. A few 3D programs directly support Light Probes (ball-looking) images, but almost all support LatLong shaped HDR images. When you make an HDR from a series of pictures of a chrome ball, this too, is considered a "light probe". In fact, any fully immersive HDR panorama can be considered a light probe, no matter what unwrapping format it is in. They are also sometimes called "Light Maps" as well.



Light Probe
Angular



Spherical
Equirectangular
LatLong
Latitude / Longitude



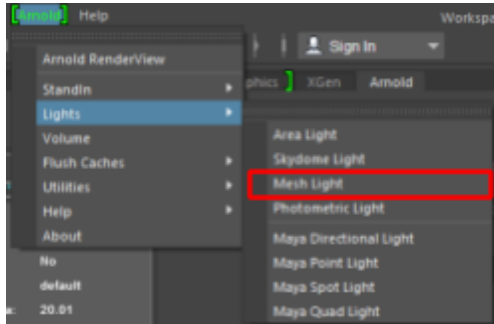
Vertical Cross
Cubic*

*6 faces can be attached or separate


Mesh Light

A mesh light turns your mesh into a light.

To add, first select your object then go to Arnold -> Lights -> Skydome Light.




In the following example, the cube with the red wireframe became a mesh light. The plane in the background has a high smooth specular. Notice how the mesh light causes a specular highlight but the actual cube giving off the light isn't visible anymore.



NOTE: You can make the object visible as well by going to the lighting properties (discussed below) and enabling Light Visible.

You can control the lighting properties by going to the shape node in the attribute editor and navigating to the Arnold section. If you want to change it back to a normal object, you can select


polymesh in the Arnold Translator dropdown (instead of mesh_light).



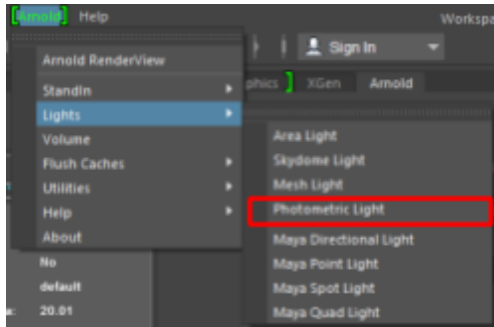
NOTE: No light will be given off in the viewport even if lighting is enabled.

Photometric Light

Photometric lights use data measured from real-world lights, often directly from bulb and enclosure manufacturers themselves. You can import IES profiles from companies such as [Erco](#), [Lamp](#), [Osram](#) and [Philips](#); their IES files provides accurate intensity and spread data for a given light model.



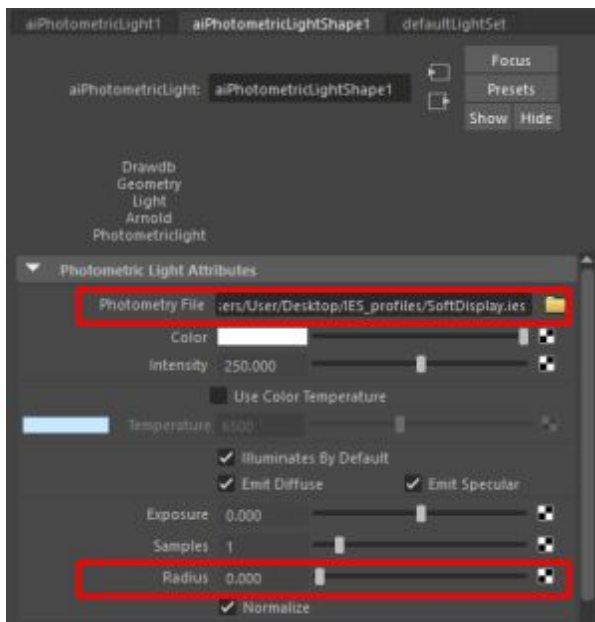
To add, first select your object then go to Arnold -> Lights -> Photometric Light.




Once you do that, go into the light attributes and select an IES profile to use (you can find a bunch for free on the Internet).

Also, unlike an area light, changing the scale does not soften the shadows. Instead, you can set how soft you want your shadows to be by using the Radius attribute.

NOTE: Scale does nothing? Even if normalize is on/off.





Arnold Object Smoothing/Subdivision

Typically, you don't want to use the standard Maya way of smoothing an object because it'll make your viewport super slow for large high-poly scenes. Instead, you can tell Arnold to automatically apply subdivision to an object when it renders.


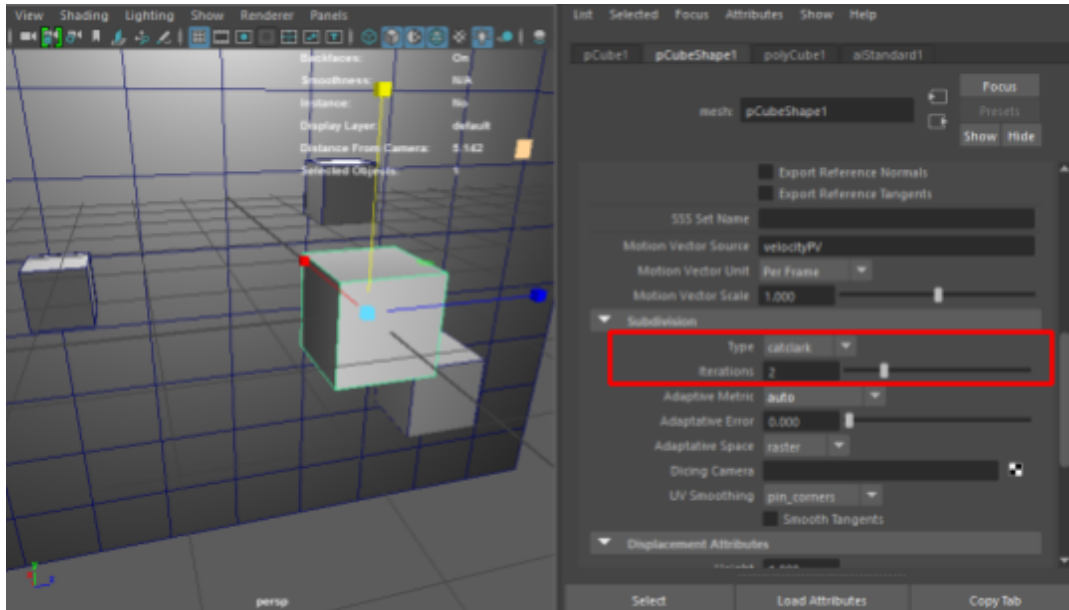
The standard Maya way of smoothing is to select an object and press 3 to apply catmull-clark subdivision (or 1 to turn it off). This applies it immediately, meaning that the new smoothed version is visible in your viewport.

The Arnold way is to go into the Arnold section of the shape node, select the subdivision type + number of iterations. The smoothing/subdivision will only be applied in the render, meaning your viewport will be fast because your objects will be low-poly.

The types of subdivision provided are...

- catmull-clark
- linear (doesn't smooth faces but does smooth normals)
- none


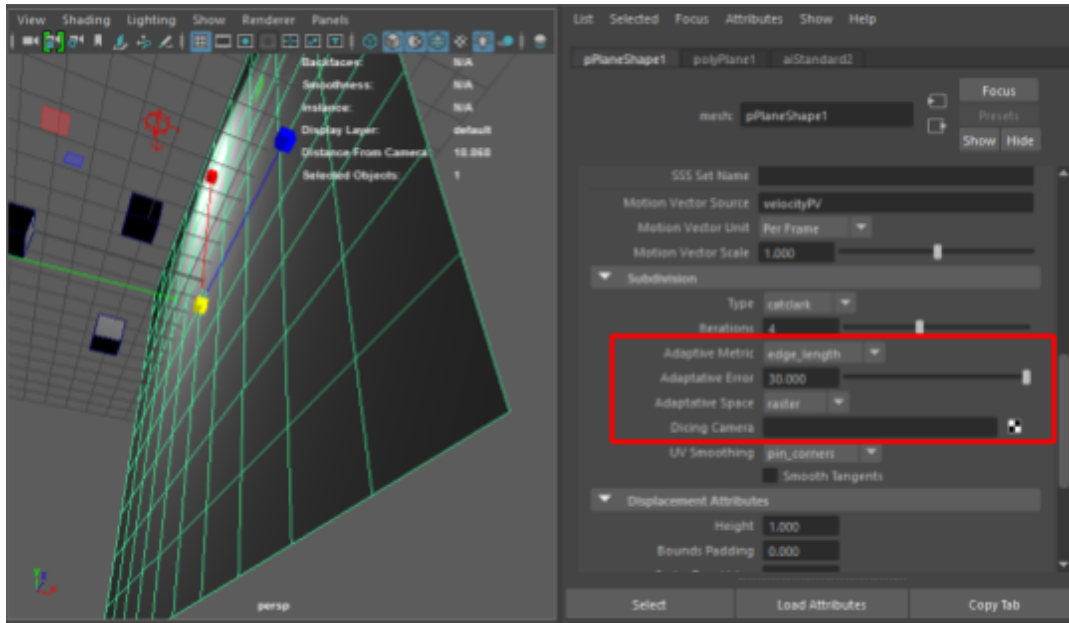
Here's an example of a catmull-clark subdivision with an iteration count of 2. Note that the viewport object is still a cube, while the renders have been rounded. The first render is a full render while the second render is a wireframe render showing the added/smoothed polys.



Adaptive Error

One cool feature of the subdivision function is that you can apply level-of-detail-esque subdivision. That is, polygons that are closer to your camera will get subdivided more than polygons that are further away.

To do this, you can use the adaptive error slider.



Adaptive Metric

Adaptive metric is how things are calculated. You can use edge_length when you're doing displacements (measured in pixels -- this is what's used most of the time). Don't use auto or the other one -- they don't seem to work very well.

UV Smoothing

UV smoothing I'm not sure about, but the options here are the same as those in regular Maya smoothing. Look those up if you need more information?

Smooth Tangents

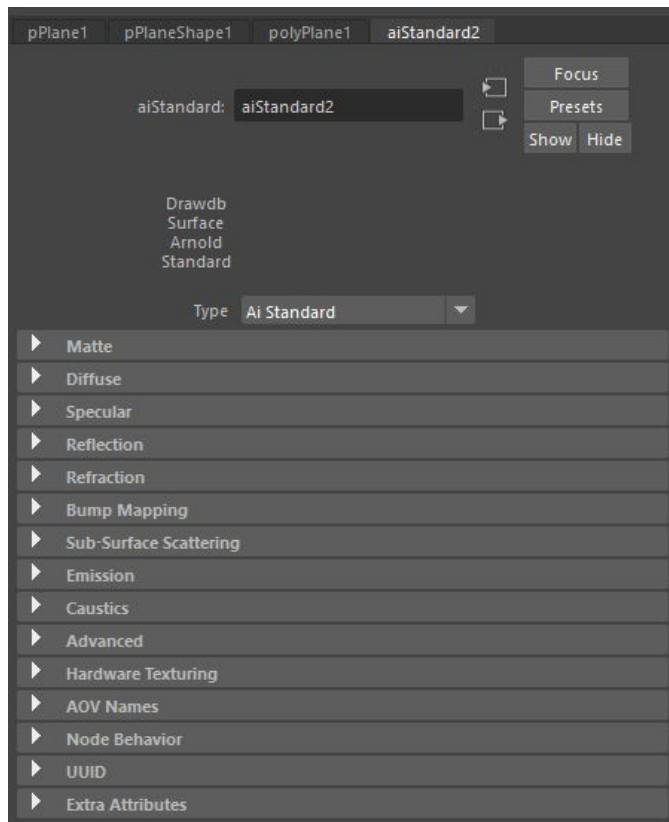
Smooth tangents is something you want to use if you have anisotropic specular highlights. The highlights can sometimes appear steppy because the polygons are no longer uniform.

NOTE: I tried to get steppy highlights to show up with an anisotropic specular that was heavily skewed -- couldn't do it. Maybe because my plane is flat so it doesn't really matter??? Need heavy curves in my object maybe???

One problem with this feature is that if you're animating your camera/object, faces will pop in and out via the tessellation that's happening. You can fix this by setting the Dicing Camera attribute to a camera that doesn't move. That camera will be used to make decisions on which polys to smooth.


Arnold Standard Material (aiStandard)

The aiStandard material shader is Arnold's basic material. You can use it to do pretty much anything you would do with a normal material: matte, diffuse, specular, specular highlights, reflection, bump/normal maps, etc...

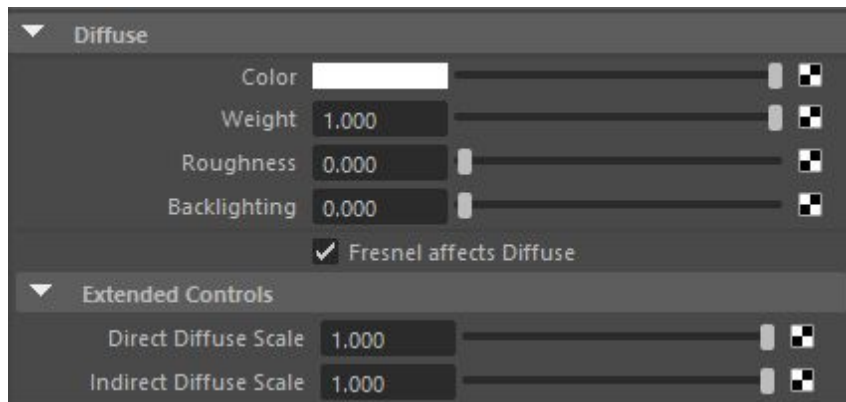


Diffuse

Diffuse reflection is the reflection of light from a surface such that an incident ray is reflected at many angles rather than at just one angle as in the case of specular reflection.




https://en.wikipedia.org/wiki/Diffuse_reflection



Color The diffuse color sets how bright the surface is when lit directly with a white light source (intensity at 100%).


Weight How much diffuse-ness to apply to this material.

Below is a weight scale of 0 to 1.




Roughness The diffuse component follows an Oren-Nayar reflection model with surface roughness. A value of 0.0 is comparable to a Lambert reflection. Higher values will result in a rougher surface look more suitable for materials like concrete, plaster or sand.

Below is a roughness scale of 0 to 1.



Backlighting Backlight provides the effect of a translucent object being lit from behind (the shading point is 'lit' by the specified fraction of the light hitting the reverse of the object at that point). It is recommended that this only be used with thin objects (single sided geometry) as objects with thickness may render incorrectly.


Below is a backlighting of 0 vs 1.



Fresnel affects

Specify whether Fresnel affects the diffuse component.

Disabled vs enabled.



Direct/Indirect scale


The amount of diffuse light received from direct/indirect sources only.

Remember that direct = directly from a light source

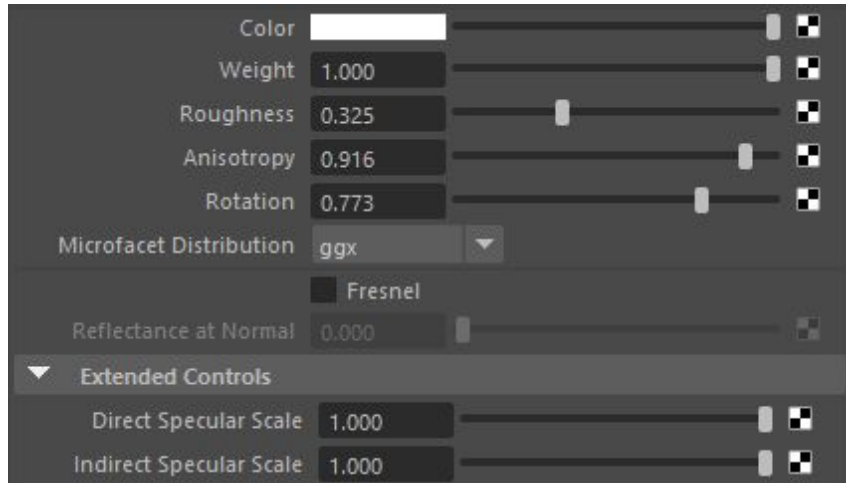
Remember that indirect = light that has bounced off something

Specular

Specular reflection, also known as regular reflection, is the mirror-like reflection of waves, such as light, from a surface.




https://en.wikipedia.org/wiki/Specular_reflection



Color The color the specular reflection will be modulated with. Use this color to 'tint' the specular highlight.

Weight How much specularness-ness to apply to this material.

Below is a weight scale of 0 to 1.




Roughness Controls the blurriness of the specular highlight. The lower the value, the sharper the reflection. A value of 0 will give you a perfectly sharp mirror reflection.

Below is a roughness scale of 0 to 1.



Anisotropy Anisotropy controls the 'skew' of the specular highlight when viewed from different angles. The default value for Anisotropy is 0.5, which means 'isotropic'. As you move this control towards 0.0, the surface is made more anisotropic in the U axis, and as you move the control towards 1.0 the surface is made more anisotropic in the the V axis.


Below is a anisotropy scale of 0 to 1.



Rotation

The rotation value changes the orientation of the anisotropic reflectance in UV space. At 0.0, there is no rotation, while at 1.0 the effect is rotated by 180 degrees.

Below is a rotation scale of 0 to 1.



NOTE: This has no effect when anisotropy is at 0.5 -- 0.5 is basically isotropic.

Microfacet Dist

Choose between GGX microfacet distribution or Beckmann (default). GGX is a microfacet distribution. It has a sharper peak and a larger tail than Beckmann. GGX is suitable for modeling light reflection from surfaces more realistically.

Fresnel

Specify whether Fresnel affects the specular component.

Direct/Indirect scale

The amount of specular light received from direct/indirect sources only.

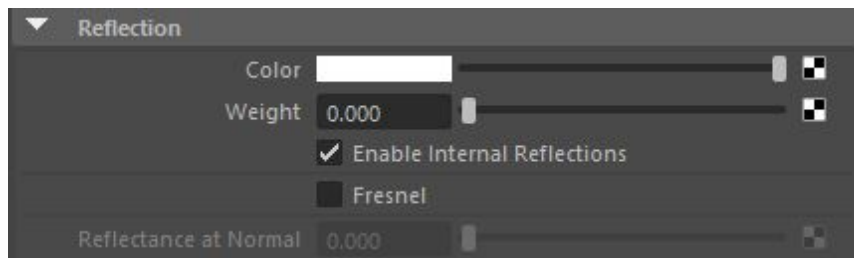
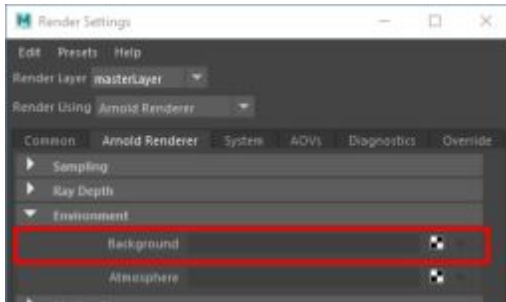
Remember that direct = directly from a light source

Remember that indirect = light that has bounced off something

Reflection

This is pretty much the same thing as specular but without all the options. The reflection from this will always be sharp -- you cannot use roughness to blur it.


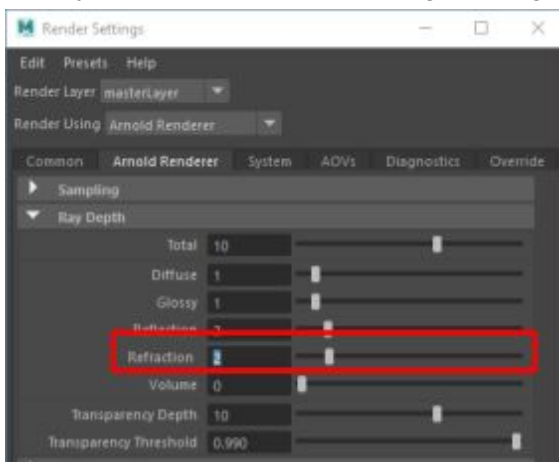
NOTE: This won't reflect skydome lights. If you want to reflect some type of environment in your reflections, you can go to render settings and set an environment background.



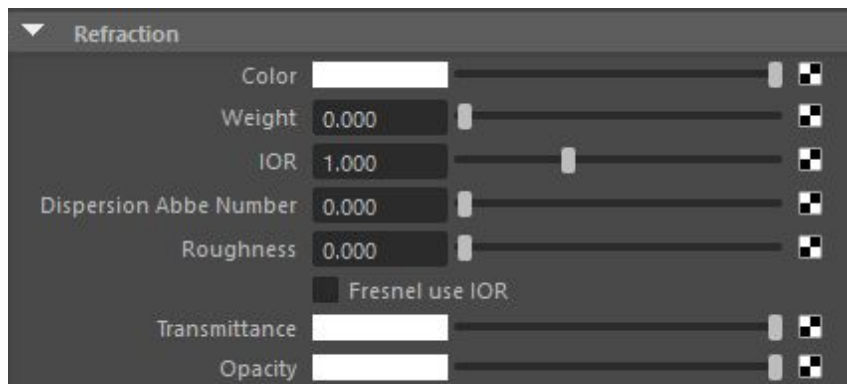
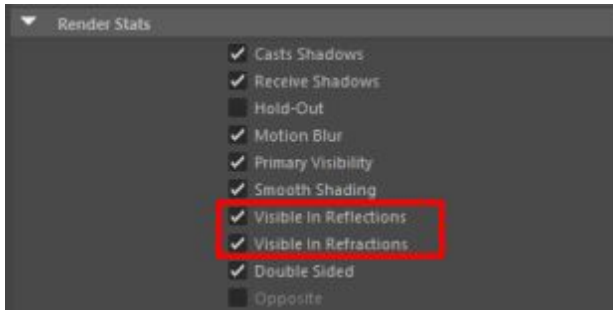
Refraction

This controls the see-through-ness of the material. Use this for things like glass or water.

NOTE: If you want this to work properly, it's super important to go to the Render settings and put in the right ray-depth values for refractions. This controls how many surfaces rays can go through. So for example, if you have a model of an empty wine glass in it, the ray would have to be able to go through AT LEAST 4 polygons/faces....



NOTE: Make sure your objects are visible in reflections/refractions by going to the Attribute Editor and navigating down to the Render Stats section. Some stuff (e.g. image planes) will have the visibility disabled by default.



Color

Transparency color multiplies the refracted result by a color. For tinted glass it is best to control the tint color via the Transmittance color since it actually filters the refraction according to the distance traveled by the refracted ray.


Weight

Weight allows light to pass through the material.

Below is a weight scale from 0 to 0.5 to 1.




Note that you must disable 'Opaque' for the mesh that has been assigned the Ai Standard shader



IOR

The index of refraction used. The default value of 1.0 is the refractive index of a vacuum, i.e., an object with IOR of 1.0 in empty space will not refract any rays. In simple terms, 1.0 means 'no refraction'.

Below is an IOR scale of 1 to 1.5.




NOTE: You can easily look up the IOR for things like water/glass/etc.. On the Internet.

Transmittance Color Like color, but tints the object more as the ray goes through the object. Thin parts of the object get less tint while thicker parts get more tint.


This is a very subtle effect, and a lot of times you won't notice a difference. Objects need to be scaled up.

Normal color vs Transmittance color (notice how transmittance gets darker the more it goes down in the water)...



Sub-surface scattering

Sub-Surface Scattering (SSS) simulates the effect of light entering an object and scattering beneath its surface. Not all light reflects from a surface. Some of it will penetrate below the surface of an illuminated object. There it will be absorbed by the material and scattered internally. Some of this scattered light will make its way back out of the surface and become visible to the camera. This is known as 'sub-surface scattering' or 'SSS.' SSS is necessary for the realistic rendering of materials such as marble, skin, leaves, wax, and milk.




Color The color used to determine the subsurface scattering effect.

Weight The amount of subsurface scattering.

Radius The radius of the area each sample affects. Higher values will smooth the appearance of the subsurface scattering. Results will vary depending on the scale of the object in your scene. See docs for more details.


Below is a radius scale from 0.25 to 0.5 to 1.



Emission


Makes your material give off light. If you go to the color and turn up to V in HSV to a super high value, this'll act like a mesh light.

NOTE: You shouldn't use this over a mesh light. Efficiency reasons or something like that.



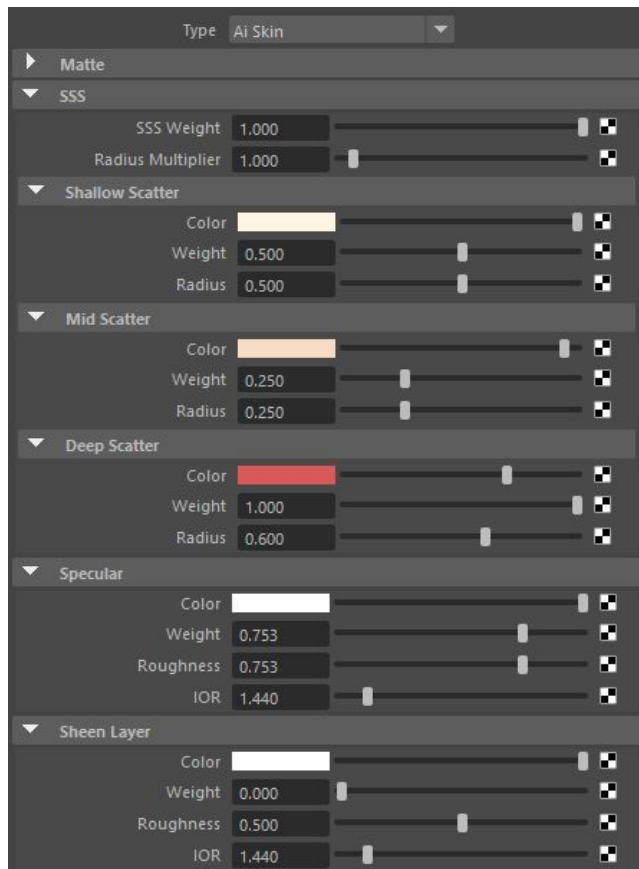

Caustics

“The light patterns at the bottom of a swimming pool” -- you can enable and disable them.



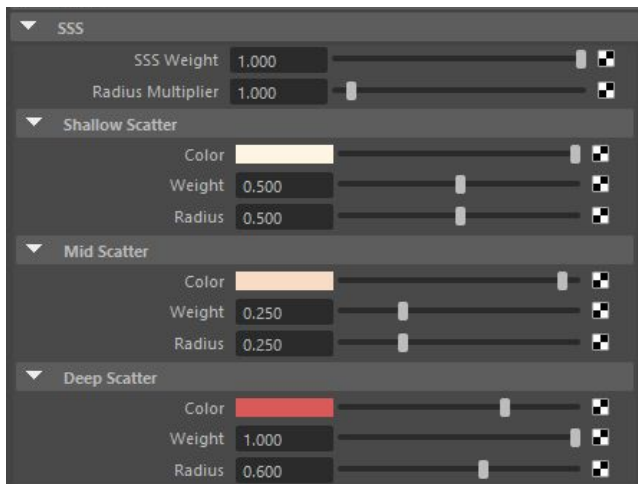

Arnold Skin Shader (aiSkin)

The aiSkin material shader is Arnold's skin material. You can use it for lots of different types of materials where rays go through the materials and scatter out: human/animal skin, vegetation, plastics, etc.. It provides you with 3 different layers of sub-surface scattering.



SSS Layers

Sub-Surface Scattering (SSS) simulates the effect of light entering an object and scattering beneath its surface. Not all light reflects from a surface. Some of it will penetrate below the surface of an illuminated object. There it will be absorbed by the material and scattered internally. Some of this scattered light will make its way back out of the surface and become visible to the camera. This is known as 'sub-surface scattering' or 'SSS.' SSS is necessary for the realistic rendering of materials such as marble, skin, leaves, wax, and milk.




SSS Weight

The overall contribution of the subsurface scattering effect.

Radius Multiplier

Overall multiplier of the SSS radius of the three layers.

Radius multiplier of 0 to 100...




NOTE: From the docs... SSS is very scale dependent. You will need to adjust the radius multiplier depending on the size of your model. If you were to render using the default Skin settings, you may get something

that does not look correct. Alternatively, adjusting the scene scale can have similar results.

- Color** The color used to determine the subsurface scattering effect.
- Weight** The amount of subsurface scattering.
- Radius** The radius of the area each sample affects. Higher values will smooth the appearance of the subsurface scattering. Results will vary depending on the scale of the object in your scene. See docs for more details.

Below is a radius scale from 0.25 to 0.5 to 1.



Specular and Sheen

The 'Specular' and 'Sheen Layer' settings are the same. It is recommended that you use 'Sheen Layer' for tight specular reflections and 'Specular' for broader specular reflections.



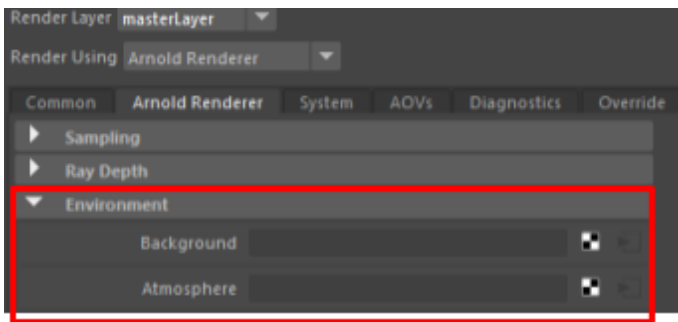
- Color** Specular highlight color.
- Weight** The amount of specular.
- Roughness** Roughness of the specular highlight.
- IOR** Index of refraction.
The amount of reflection will be dependent on the viewing angle of the

surface following the Fresnel equations (which depends on the IOR value).

I have no idea what this is used for... “turning it up has less of a fresnel effect -- that means that the faces facing away from us are reflective but also the faces facing towards us”

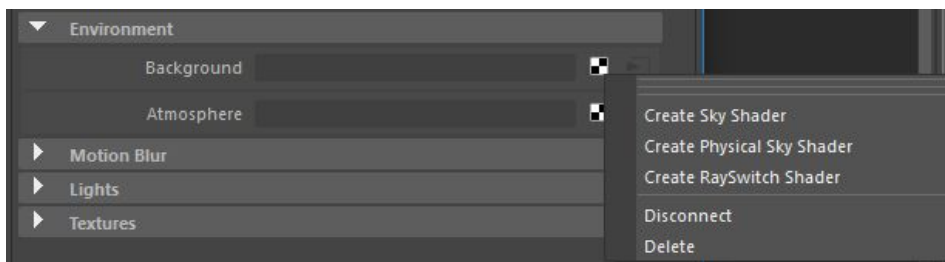
Arnold Environments

You can fiddle with your “environment” by going to Render Settings. Under the Arnold Renderer tab, there should be a Environment section.



Background

Background selects a “background” shader.



Sky Shader

This looks very similar to a skydome light, but it renders as well as gives off light? Here's what the website says...

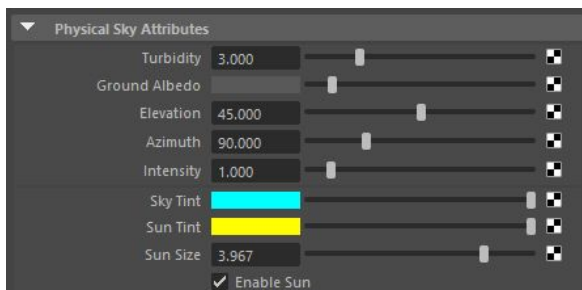
This shader is intended to be used as a global environment shader (or a background shader in Arnold terms), thus affecting all objects in the scene. By default it has a simple white color, and will illuminate the whole scene evenly using this color. For image based lighting, use the Skydome light node instead.

NOTE: This doesn't show in reflections/refractions by default -- check out the render stats section of the attributes and enable if needed.

Physical Sky Shader

Simulates a sky? Apparently there's something like this with pretty much every renderer. I don't know what all the options mean, but you can look them up individually if needed. Here's what the website says...

This shader implements a variation of the [Hosek-Wilkie](#) sky radiance model, including the direct solar radiance function. You can plug it into the environment, or to the color input of a Skydome light (which benefits from efficient importance sampling, sending more rays to bright areas around the sun) or add it directly as an environment shader. Currently this shader is invisible to GI diffuse and glossy rays. To use it as a light source you must attach it to a [skydome_light](#) with sufficient resolution to capture the small solar disk.



NOTE: Setting this under environment in the render settings did nothing. I had to link it to a skydome light to get what I wanted.

RaySwitch Shader

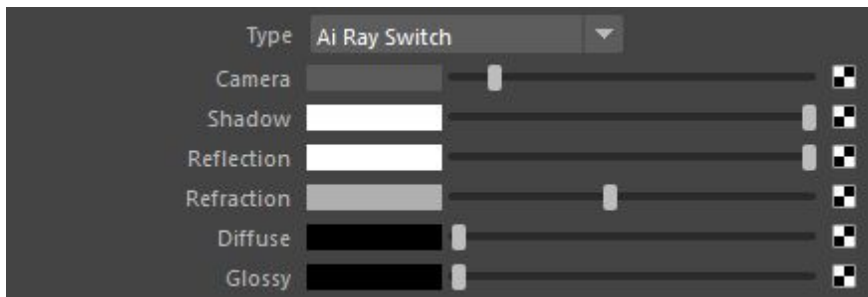
Unclear on what this is actually used for. Apparently it allows you to set what type of shader gets applied to what type of ray. For example, you can make reflections red (or link to some shader).

This is useful for "optimizing a scene". For example, if you have a complex shader but you don't want that shader being calculated for thing like reflections, you can link reflections in here to a simple version of that shader.

Here's what the website says...


This shader makes it possible to evaluate different shader trees per ray. This decreases the shading complexity of a scene and thus the render times, and increases artistic control. It can be used to remove unnecessary secondary rays (specular/sss), make speculars even more

glossy in specular rays, control the color of opacity in shadow rays to fake light scattering through a tissue or add a second specular lobe in camera rays only.



Atmosphere

There are two types of Atmosphere in Arnold, Fog and Volume Scattering. Fog simulates the effect of light scattering, which causes more distant objects to appear lower in contrast, especially in outdoor environments. Volumetric Scattering simulates light scattered by a thin, uniform atmosphere. It produces shafts of light and volumetric shadows cast from geometric objects.



Fog used for underwater effect vs volume scattering used to simulate rays of sunlight...



Fog

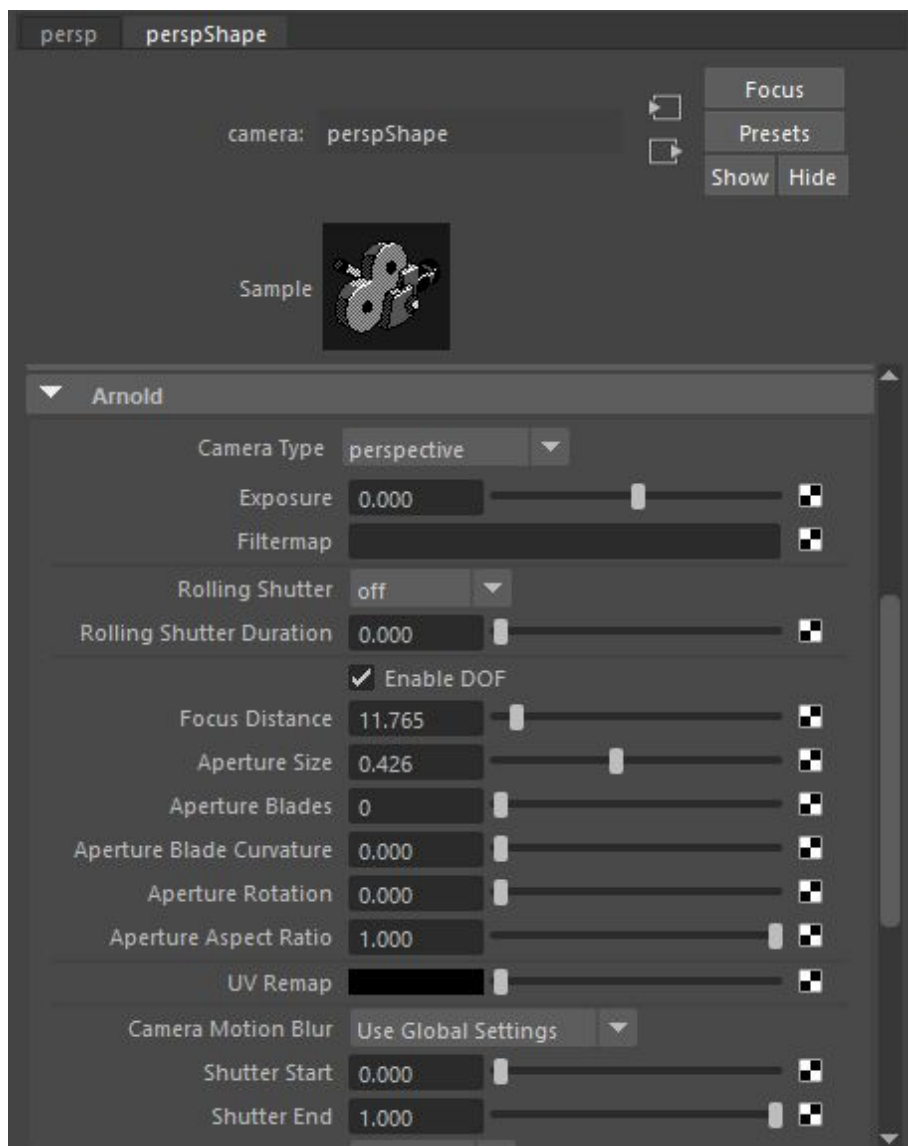
Don't know enough about fog to write here. Look up as needed.

Volume Scattering

Don't know enough about volume scattering to write here. Look up as needed.

Arnold Depth of Field



You can create a depth of field effect for your render through the camera you're rendering from. Select the camera and go down to the Arnold section of the attribute settings.



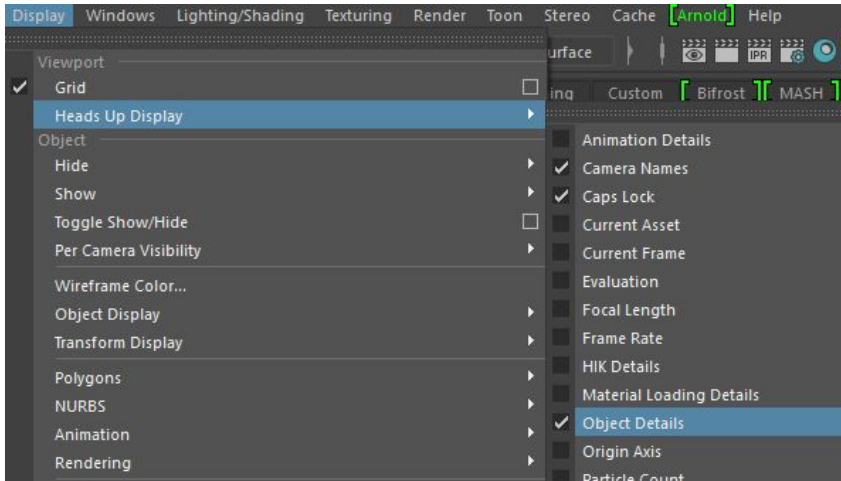
Click Enable DOF.

- Focus Distance controls the distance in which things will remain in focus.
- Aperture Size controls how blurry things will be for stuff that's out of the focus distance.
- Aperture Blades is bokeh related -- not sure what this has to do with blur? But setting this to 5 will give you a pentagon-ish lensflare bokeh type thing?
- Aperture Curves is bokeh related -- not sure what this has to do with blur? But setting this higher will make straight edges of the bokeh more curved?
- Aperture Rotation is bokeh related -- not sure what this has to do with blur? But sets the rotation of the bokeh lensflare pattern artifacts that show up?


The first 2 options are the main ones. Not sure about the rest.



HINT: Want to know the distance between your camera and an object? Use the heads up display by going to Display -> Heads Up Display -> Object Details. When you select an object in your viewport, you should now see how far it is from the camera. You can put this value in for Focus Distance to focus in on that object.



NOTE: You'll notice a lot of grain in the blurry parts. There's no way to fix that other than increasing the main Camera (AA) samples in the render settings.




Arnold AOVs

AOVs stand for Arbitrary Output Variables. You can use them to segment off specific parts of the render (e.g. the specular portion of a material) and save that as its own image. AOVs don't add much to the render time because Arnold needs to calculate these details anyways for doing the full render.


AOVs can be used for debugging your renders (e.g. checking out how the specular highlights look on their own) or for compositing (e.g. splitting out the indirect and direct lighting and combining in nuke).

To set up an AOV, go to the AOV tab in the render settings.



1. Select the material type in the 'AOV Groups'
2. Select what part of material to segment out into its own render in 'Available AOVs'
3. Click the >> button to move them over to 'Active AOVs'

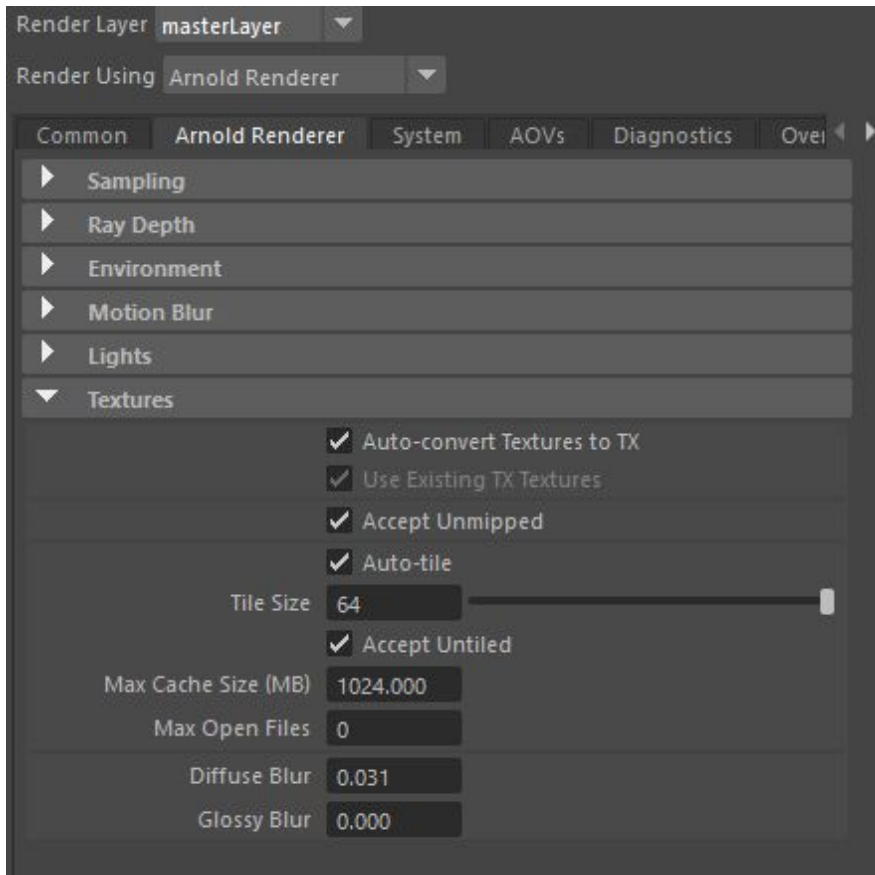
Now, in the Arnold Render View, you can view these AOVs by selecting them in the AOV dropdown.





Textures in Arnold

To get to the texture settings, go to the Arnold render setting under the Textures section.



Generally you shouldn't touch these settings. They're mainly knobs for controlling performance. Check out <https://support.solidangle.com/display/AFMUG/Textures+-+settings> for more information on what each of these parameters are for.

One important thing to note is that textures are automatically mipmapped. If you want to load a less-detailed version of the texture from the mipmap, you can change the mipmap bias attribute of your file node that loads up the texture (why would you ever want to do this?)...

