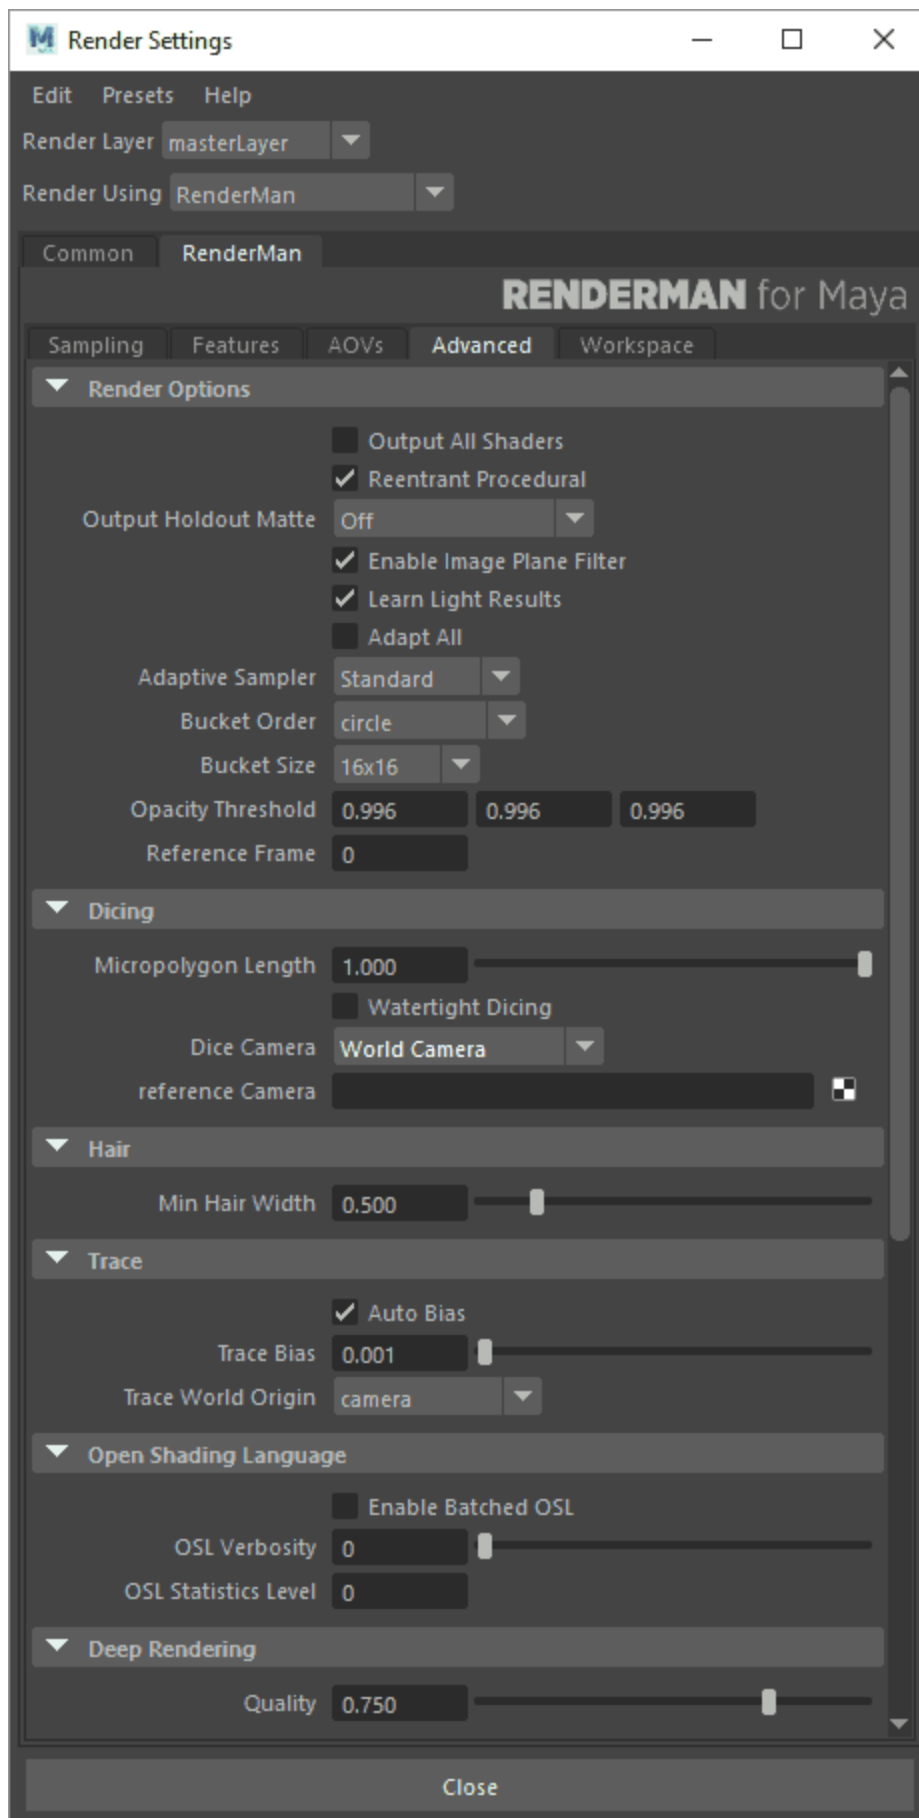


Advanced

Many of the settings in the **Advanced Tab** you will never need to edit. However, for the more advanced user who wants fine tuned control over RenderMan for Maya, they are here.



Render Options

Output All Shaders: Typically we export only material networks attached to objects, but if you're shading a procedural it may not be connected in a way that is exported. This option will export all networks regardless.

Procedural Reentrant: This allows procedurals to run simultaneously in multiple threads and greatly speed up renders with procedurals (hair). You most likely want to leave this on.

Output Holdout Matte: When using the [Holdout](#) workflow, output the Shadow AOV to this specified buffer.

Enable Image Plane Filter: Render the image plane using the Pixar Image Plane Filter (Maya uses an actual geometric plane)

Learn Light Results: This is a new option for scenes with many lights. The renderer will refine the light selection process to improve sampling. This is especially useful in scenes with many occluded lights to avoid wasting samples from shadowed surfaces. This version is not deterministic and small differences can be seen if not well converged on re-renders of the same scene.

Adapt All: Apply the sampling quality criteria (pixel variance) to all color/illumination AOVs. This can increase render time as some AOVs may have more noise than the Beauty and trigger more samples through all the AOVs.

Adaptive Sampler: Make use of the improved sampling algorithms instead of the default. You may need to slightly increase your Pixel Variance setting to avoid longer render times (but better resulting quality on complex images) Your min samples should be set to 0 which has special meaning that the square root of the max samples is used. If you manually input a low value, it will terminate early.

Bucket Order: RenderMan subdivides the output image into small rectangular regions called buckets, and renders a single bucket at a time. This allows a small portion of the image to be in memory at any one time. The drop-down list allows users to choose the order in which buckets are displayed in "it" or the Render View. The default is Spiral. Note that only some bucket orders are supported for [checkpointing](#) and recovery.

Bucket Size: Specify the size of the buckets. Larger buckets will require more memory to render. Conversely, smaller buckets are less efficient for shading, but will use less memory during rendering. If your scene is using a lot of memory, you may want to try setting this field down to 8 by 8 or even 4 by 4 buckets, but most likely you will not change this.

Opacity Threshold: Controls the Opacity Threshold, if set too low, it may terminate rays early on semi-opaque objects. An object's opacity is controlled by the material Presence parameter and depth in supported integrators.

Reference Frame: Once per job passes and *static* objects are evaluated at the frame designated here.

Dicing

Micropolygon Length: Controls the tessellation rate for displacement. The default setting of 1 will dice displacements to 1 pixel's size. (As seen from the dicing camera, typically the render camera). Smaller numbers are finer geometric detail at the cost of creating more geometry.

Watertight Dicing: Closes pinholes in diced geometry or attempts to close gaps in displacement. Useful if you have holes in your solid object alpha.

Dicing Camera: Enables the use of a reference camera for dicing. This can be done globally, or restricted to specific objects (e.g. objects with displacements that "sparkle" as the camera moves) by also attaching this as a per-object attribute.

Reference Camera: Specify the camera to be used as your dicing camera here. Click on the widget to create a new camera, or right click in the field to select an existing camera. Note that the Reference Camera needs to be different from the render camera.

Hair

Min Hair Width: Sets a minimum width for hairs/RiCurves. Narrower hairs will be clamped to this value and their opacity will be reduced to compensate for the difference in visual weight. This may improve anti-aliasing quality if not too small.

Trace

Trace Auto Bias and Trace Bias: RenderMan will automatically compute a trace bias for displacing ray origins off a hit surface. However, if you wish to turn off the auto bias, you can set the bias manually here. Per-object settings should have Auto Bias turned off. Increase the value slowly until the artifact (self-shadowing, dark lines, etc.) disappears.

Trace World Origin: Move the world origin for increased floating point precision for raytracing. Sometimes when rendering very large scenes, the camera is moved to a position quite far from the origin. This may introduce precision issues depending on your scene units, values exported to the renderer are very large in these cases. Possible values are:

- "world" (default) uses the true world origin
- "worldoffset" use user provided trace:worldoffset as the raytrace world origin (the control will expose an XYZ coordinate field)
- "camera" use the camera position as the raytrace world origin, you may wish to use this in large scenes with multiple camera positions as the default.

Open Shading Language

Enable Batched OSL: Enable performance improvements on batched shading points through vectorized instructions on supported hardware. (Available on machines with a minimum of support for Advanced Vector Extensions (AVX). Additionally, it can take advantage of the 512-bit wide datapaths on machines that support Advanced Vector Extensions 512 (AVX-512). This is *on* by default

OSL Verbosity: Valid in ranges of 0 to 5 with 5 having the most messages. See the [OSL documentation](#) for more information.

OSL Statistics Level: As above, valid in ranges of 0 to 5 with 5 having the most messages.

Deep Rendering

Quality: Controls the decimation rate of merging samples. A setting of 1.0 has no merging of samples and will generate huge files. The default of 0.75 maps to Deepshadow Error 0.01 and when set to 0 maps to Deepshadow Error 0.16. Note this is not a linear mapping.

Crop Window

Enable: Use the crop window, off (default) renders the complete frame based on camera settings

Top Left/Bottom Right: Specify the coordinates (x and Y) for a crop window to render. This can be useful when rendering many iterations for debugging or lookdev when only a specific region is under consideration. You may also use this to shift around the rendered region to tile together a very large render (possibly for print) and send the different regions to a different render node. Please note using crop windows with the PxrVCM renderer may generate different results than a full frame since we only generate photons for the region being rendered as an optimization.

IES Profiles

Scene units to meters: While scenes are exported to renderers "unitless", IES files specify a unit measure. Use this setting to specify the scene unit to meter conversion factor. For instance, if 1 unit = 1 meter, use 1. If 1 unit = 1cm use 100.

Ignore Watts: Ignore the internal watts measurement in the profile

Cache Sizes

Allows users to customize the cache sizes for various assets, including textures, ray-traced geometry (specifically the tessellation cache), and opacity (as a function of Presence in materials). The values are in megabytes, and the renderer will automatically unload data from the cache when the limits specified here are reached. These values may be conservative based on your machine's available resources. If your render stats show many cache misses (reloading and re-reading) and your overall usage is below your available memory, you may want to increase these values to improve performance, keeping them reasonable and within your memory budget.

Statistics

Level: Enable this to output [diagnostic information](#) from render jobs. Useful for timing, debugging, and seeing what RenderMan is doing. Higher numbers are more verbosity. Note there can be a performance impact that adds up by requiring the renderer to lookup these strings.

XML File: You can specify a target destination for both "Old School" and XML statistics.

LPE Lobe Mappings

This section maps the specified lobes to an output. You will notice in the [Light Path Expression](#) documentation that per-lobe outputs rely on these being correct so we recommend leaving them at their default. You may use User Lobes in custom C++ BxDFs.

Example: You will note in the defaults that Diffuse 3 Lobe ID is Subsurface. The corresponding LPE is:

```
color lpe:CD3[DS]*[<L.>0]
```

Note the D3 lobe specified. This matches the IDs assigned here.

There are no D1, S1, or U1 lobes as these are "fallthrough" when something is unspecified. User 2 Lobe are the collective Albedo results (texture data).