

Renderer Options

- [Common](#)
- [Baking](#)
- [Shading](#)
- [Dicing](#)
- [Display](#)
- [Statistics](#)
- [Configuration](#)
- [RIB](#)
- [OSL](#)
- [Volumes](#)

Common

Name	Type	Default	Description
hider: adaptall	<i>int</i>	0	Enable adaptive sampling of all AOVs. When enabled, all channels (beauty pass + AOVs) are considered when deciding if a pixel is converging. The default (off) will ignore the AOVs and only the beauty pass is considered for pixel convergence. So turning off "adaptall" may speed up your render if some of your AOVs never reach the specified variance threshold. Value range: true/false.
hider: adaptive metric	<i>string</i>	variance	Specifies the metric used to drive adaptive sampling. "contrast" measures the contrast between samples in a pixel. "variance" measures the variance of the samples in a pixel. "halfbuffer" splits the samples into two classes and measures the difference between the average of each of the two sample classes. "contrast-v22" uses the relative contrast computation of RenderMan 22 and older. "variance-v22" uses the relative variance computation of RenderMan 22 and older. Supported values: contrast, variance, halfbuffer, contrast-v22, variance-v22.
hider: bluenoise	<i>int</i>	1	Setting to control whether the samples should have a visually pleasing blue noise distribution or not. Default is 1 (on). Set to 0 (off) if the samples in adjacent pixels need to be completely decorrelated. Value range: true/false.
hider: darkfalloff	<i>float</i>	0.025	De-prioritize low luminance samples in the adaptive sampler. Use this parameter to avoid oversampling dark areas while maintaining overall image quality. Higher values cull more samples, speeding up render time while increasing noise in darker areas. Value range: 0 to 1.
hider: decidither	<i>int</i>	0	Maximum allowed pixel decimation rate for interactive rendering. When making rapid edits to heavy scenes and this setting is non-zero, the renderer will try to reduce tearing by updating the whole screen progressively with a dissolve-like effect. Set this to zero to disable this and force it to fully update the buckets that it can get to. Value range: 0 to 6.
hider: exposurebracket	<i>float[2]</i>	-1.0 1.0	Intended exposure Bracket [min max] in post-production to help inform the adaptive sampling.
hider: extrememotiondof	<i>int</i>	0	Enable improved sampling accuracy. In some cases where a large amount of motion blur is used with depth of field, sampling artifacts could cause an undesirable result. This flag causes a slower but more accurate sampling method to be used to alleviate the artifacts. Due to the performance degradations of this sampling method, this flag should only be used when necessary. Value range: true/false.
hider: geomShadowTermBias	<i>int</i>	1	This setting will control when to compute and apply a bias to Poly Meshes to avoid the geometric shadow terminator artifacts that appear for shading points when they are visible from a light wrt their shading normal, but not visible wrt their geometric normal. The default value is 1 (on). Set to 0 (off) if you don't want this bias to be applied to the Poly Meshes in your scene. Value range: true/false.
hider: incremental	<i>int</i>	0	Enables the progressive display of results (rather than bucketed results). Value range: true/false.
hider: jitter	<i>int</i>	1	This option toggles the jitter (variation) in camera samples. By default, this hider option is enabled (set to 1), meaning that camera samples will be placed randomly in a pixel location. When it is turned off, camera rays are placed in the center of the pixel. When turned off, surface edges will appear aliased and jagged. Value range: true/false.
hider: maxsamples	<i>int</i>	0	An explicit limit for the number of samples. If maxsamples is equal to its default then the maximum number of samples will be set equal to 64. Note that minsamples defaults to the square root of the maxsamples setting. Value range: 0 or more.
hider: minextra samples	<i>int</i>	-1	This is the minimum number of extra samples that a pixel should continue shooting with after a sample fails to meet its convergence test with the adaptive sampler. Defaults to matching minsamples.
hider: minsamples	<i>int</i>	-1	Specifies the minimum number of samples for adaptive sampling. The default value is the square root of maxsamples.
hider: pixelfilter mode	<i>string</i>	importance	Specifies how pixel filtering of camera samples is performed. "Importance" (the default) warps the samples in screen space around the pixel center, and does not share sample results between pixels. "Importance" is needed by some postprocessing algorithms such as denoisers, and is necessary to get the visual benefit of blue-noise samples. It is also needed to produce variance AOVs. "Weighted" is the older method that blends together sample results from neighboring pixels. Supported values: weighted, importance.
hider: sample motion	<i>int</i>	1	This option toggles the sampling of motion blurred micropolygon in the hider, and is thus only meaningful if motion blur is being used in your scene. By default, this hider option is enabled (set to 1), meaning that micropolygons will be blurred and sampled in the hider. When disabled (set to 0), micropolygons will no longer be blurred before sampling; however, the dPdttime variable will still be computed for all objects. This variable can be output into a separate image using an arbitrary output variable (aov) and used to perform motion blur as a postprocess. Value range: true/false.

hider: sampleoffset	<i>int</i>	0	This allows several images to be rendered in parallel (with different sampleoffset values) and then combined. With non-adaptive sampling: Let's say you render four images with 256 samples each, with sampleoffsets 0, 256, 512, and 768. If you combine those four images, you'll get exactly the same image as if you had rendered a single image with 1024 samples. With adaptive sampling: Let's say you again render four images, each with "maxsamples" 256, with sampleoffsets 0, 256, 512, and 768. Let's say that due to adaptive sampling, some given pixel only gets 64 pixel samples in each of the four images. Then the combined image has been rendered with sample numbers 0-63, 256-319, 512-575, and 768-831. Due to the stratification of the samples, this is not quite as good as if you had rendered a single image with 256 consecutive samples. However, it is still better than rendering a single image with only 64 samples.
hider: samplestride	<i>int</i>	1	This facilitates rendering several images in parallel (with the same aamplestride but different sampleoffset) and then combine them incrementally during interactive rendering. Default value is 1. With non-adaptive sampling: Let's say you render four images with 256 samples each, with samplestride 4 and sampleoffsets 0, 1, 2, and 3. If you combine those four images, you'll get exactly the same image as if you had rendered a single image with 1024 samples, and during incremental/interactive rendering you'll get the benefit of best use of the stratified progressive sample sequences. If you can "assemble" images with samples 0-3, 4-7, 8-12, etc during interactive rendering, you've got the benefits of faster rendering without jumping to high samples early on. With adaptive sampling: Let's say you again render four images, each with "maxsamples" 256, with samplestride 4 and sampleoffsets 0, 1, 2, and 3. This time with adaptive sampling (PixelVariance > 0). Since each of the four images is more noisy than the combined image, the adaptive sampling will use more samples before deciding that the image error is acceptable than if it had been rendered as a single image. So the PixelVariance setting may need to be adjusted.
hider: type	<i>string</i>	raytrace	Render mode. Supported values: raytrace, bake.
Ri: Frame	<i>int</i>	0	Frame number.
Ri: PixelVariance	<i>float</i>	0.015	Controls adaptive sampling. This sets the upper bound on the acceptable estimated variance of the pixel values from the true pixel values. Value range: 0 to 1.
Ri: Shutter	<i>float[2]</i>	0 0	Scene motion shutter.
Shutter: offset	<i>float</i>	0	Offset shutter open/close.

Baking

Name	Type	Default	Description
hider: bakebboxmax	<i>point</i>	1e30 1e30 1e30	Bounding box max for "integrator" baking of 3D point clouds (and for the integrator part of "all" baking. In world space. Ignored for "pattern" baking and for 2D baking.
hider: bakebboxmin	<i>point</i>	-1e30 -1e30 -1e30	Bounding box min for "integrator" baking of 3D point clouds (and for the integrator part of "all" baking. In world space. Ignored for "pattern" baking and for 2D baking.
hider: bakemode	<i>string</i>	pattern	Specifies whether the renderer should bake Pattern results or Integrator results or both. Supported values: pattern, integrator, all.
hider: bakeudimoffset	<i>int</i>	0	Offset when baking for a subset of udims. Used together with bakeudimstride. Default 0 means bake all udims when bakeudimstride is 1.
hider: bakeudimstride	<i>int</i>	1	Stride when baking for a subset of udims. Default 1 means bake all udims.
hider:invert	<i>int</i>	1	Specifies whether or not to invert the T texture coordinate when baking integrator results. Value range: true/false.
hider:primvar	<i>string[2]</i>	st,	Specifies dicing manifold used when baking integrator results.

Shading

Name	Type	Default	Description
lighting: minimumestimate	<i>float</i>	1e-6	Minimum estimate for light selection. Local lights whose estimate falls below this value will be ignored. Increasing this value may help speed up direct lighting at the expense of accuracy.
lighting: selection learning scheme	<i>int</i>	1	Employs a machine learning algorithm to more accurately importance sample lights, taking into account occlusion and light filters. Value range: true/false.
trace: bvhcompression	<i>float</i>	0	Provides a hint to the system to favor memory utilization versus speed when constructing the ray accelerator. Valid values are in the range [0, 1], where the higher the values the more compact the memory representation. The control behaves almost linearly with some discontinuities in the compression rate as different strategies are activated. The result obtained is likely to be scene dependent. As a rule of thumb, leave this control to 0 unless you need to reduce memory utilization in order to render a given scene or sequence. When tuning, set the control to 1 at first to see how much memory can be saved through compression and what time overhead is generated, then dial the number back to a lower value (if possible) and reach the desired balance between speed and memory usage. Value range: 0 to 1.
trace: worldoffset	<i>float[3]</i>	0 0 0	Used as the origin when the worldorigin option is set to "worldoffset".

trace: worldori gin	<i>string</i>	camera	Set a new world origin to give better floating point precision. Supported values: world, worldoffset, camera.
---------------------------	---------------	--------	---

Dicing

Name	Type	Default	Description
curve:minwidth	<i>float</i>	0	Minimum curve width in raster coordinates. Value range: 0 or more.
dice: offscreenmulti plier	<i>float</i>	1	Dicing rate multiplier for objects outside the viewing frustum. Value range: 0 or more.
dice:triangle	<i>string</i>	trim	The method used for dicing "triangles". The default method "trim" renders triangles by trimming a quad. The "blp" method renders a quad with two points being coincident. Supported values: trim, blp.
stitch: refwarning	<i>int</i>	1	Enable stitch warnings. Value range: true/false.

Display

Name	Type	Default	Description
bucket: order	<i>string</i>	horizontal	Render bucket order. Supported values: horizontal, vertical, zigzag-x, zigzag-y, spacefill, random, spiral, circle.
bucket: orderorig in	<i>int[2]</i>	-1 -1	First bucket for "spiral" and "circle" order. Image center is [-1 -1].
checkpoi nt:asfinal	<i>int</i>	0	When enabled, this keeps the extra channels and the checkpoint tag when writing the final image for the render. The final image will be identical to any other checkpoint, rather than a slimmed down image. This means that once your image has reached the quality you've set and it completes, it can always be restarted by the user later. Value range: true/false.
checkpoi nt: command	<i>string</i>	<i>empty</i>	System command to be called after checkpoint. If system calls are enabled, then after a checkpoint has been written, the specified command will be called. This is synchronous; the rendering threads are quiescent while this runs and will not resume again until the process returns, avoiding possible race conditions if the command takes a token %i will be replaced with the current increment, zero-padded to 5 digits. The token %e will be replaced with the elapsed time in seconds, zero-padded to 6 digits. The token %r will be replaced with the reason for this update to the checkpoint files (either completely `finished`, `exiting` early due to exitat option, or a normal `checkpoint`). Literal % characters may be inserted with %%.
checkpoi nt:exitat	<i>string</i>	<i>empty</i>	How long render should run before writing a final checkpoint and terminating. Value can also be a string such as "10m" or "0.5h". Set to 0 or empty to disable.
checkpoi nt: interval	<i>string</i>	<i>empty</i>	Write framebuffers to disk every N progressive steps. If value is positive the units will be interpreted as seconds. If value is negative its units will be interpreted as increments. Value can also be a string such as "10m" or "0.5h". Setting option to 0 or empty will disable checkpoints.
checkpoi nt: keepfiles	<i>int</i>	0	If set to true then prman will write to a new, numbered image for each checkpoint. Note: This will disable the final write of a checkpoint file which means the render cannot be recovered using the 'recover' option. Value range: true/false.
deep: flagvolu mes	<i>int</i>	1	Auto is off for dtex, on otherwise. Supported values: Off (0), Auto (1), On (2).
Ri: CropWin dow	<i>float[4]</i>	0 1 0 1	Crop window [left right top bottom] in NDC space.
Ri: FormatP ixelAspe ctRatio	<i>float</i>	1	Output image pixel aspect ratio. Value range: 0 to 1.
Ri: FormatR esolution	<i>int[2]</i>	640 480	Output image resolution.
Ri: FrameA spectRat io	<i>float</i>	-1	Output image aspect ratio.
Ri: Screen Window	<i>float[4]</i>	0 0 0 0	Screen window [left right top bottom] in screen space.

Statistics

Name	Type	Default	Description
shade:chiangCompatibilityVersion	float	24.0	Setting to allow PxrChiang to work in compatibility mode for older versions of the Bxdf, so renders from different versions can match their results. The default value is 24.0.
shade:debug	int	0	Shading debug level. Value range: true/false.
shade:incorrectPointOpacityCalculation	int	0	Use old method of combining RiPoint falloff with presence or opacity. This has been fixed in R25. Setting this option will fall back to the old, incorrect, behavior for compatibility purposes. This option will be removed in R26 and the new, correct, method will always be used. Value range: true/false.
shade:roughnessmollification	float	1.0	Specular Roughness Mollification. This selectively increases material roughness on secondary bounces after diffuse or broad glossy reflections and on highly curved surfaces. This reduces noise and fireflies, but can also make caustics blurrier or disappear entirely. Set this value to 0.0 when you want to render ground-truth caustics. The default value is 1.0.
shade:shadowBumpTerminator	int	1	Setting to control whether bump shadow terminators are smoothed in non-hair bxdfs. Value range: true/false.
shade:subsurfaceTypeDefaultFromVersion24	int	0	Use old default value for PxrSurface 'surfaceType' parameter. Old default is 0 (Jensen dipole diffusion), new default is 5 (path-traced subsurface scattering). Value range: true/false.
statistics:displace_ratios	float[2]	0.1 1	Suppress reporting of displacements that, when divided by the max displacement, fall in the specified range.
statistics:filename	string	empty	File name for summary statistics reported in plain text.
statistics:level	int	0	Enable statistics reporting. Value range: true/false.
statistics:maxdispwarnings	int	100	Modify the maximum reported displacement issues. If the value is set to 0, then all displacements issues are reported. Value range: 0 or more.
statistics:shaderprofile	string	empty	
statistics:stylesheet	string	empty	Stylesheet for XML detailed statistics.
statistics:texturestatslevel	int	0	Enable texture statistics reporting. Value range: true/false.
statistics:xmlfilename	string	empty	File name for detailed statistics reported as XML.

Configuration

Name	Type	Default	Description
limits:brickmemory	int	2097152	Brickmap cache size in kB. Value range: 0 or more.
limits:bucketsize	int[2]	16 16	Size of render buckets in pixels.
limits:deepshadowerror	float	0.01	Control lossy compression scheme to reduce file size of deep output. Setting to a high value will result in lower numbers of samples stored for each pixel. Value range: 0 to 1.
limits:geocache memory	int	2097152	Geometry cache size in kB. Value range: 0 or more.
limits:gridsize	int	289	Value range: 1 or more.
limits:matrixcachememory	int	0	Matrix cache in kB. Value range: 0 or more.
limits:nurbcurvaturedeffactor	int	3	NURBs max curvature multiplier. Value range: 0 or more.
limits:octreememory	int	20480	None cache size in kB. Value range: 0 or more.
limits:opacitycachememory	int	1048576	Opacity cache in kB. Value range: 0 or more.

limits: othreshold	<i>float[3]</i>	0.99609375 0.99609375 0.99609375	Opacity Culling: When rendering scenes with a large number of semi-transparent layered objects (e.g. hair), the opacity culling threshold can be set for a significant time and memory savings. Essentially, a stack of visible points whose accumulated opacity is greater (in each channel) than the specified limit will be considered fully opaque by the hider, and objects behind the stack will be culled.
limits: pointmemory	<i>int</i>	20480	Point cache size in kB. Value range: 0 or more.
limits: procedural bakingclumps size	<i>int</i>	0	Value range: 0 or more.
limits: ptexturemax files	<i>int</i>	128	Value range: 1 or more.
limits: ptexturememory	<i>int</i>	32768	PTex cache size in kB. Value range: 0 or more.
limits: rendermemory	<i>float</i>	0	Memory limit as percentage of machine memory. Value range: 0 to 1.
limits: rendertime	<i>int</i>	0	Maximum render time in minutes. Value range: 0 or more.
limits: shadesize	<i>int</i>	289	The maximum number of vertices, not grid. Value range: 1 or more.
limits: texturememory	<i>int</i>	2097152	Texture cache size in kB. Value range: 0 or more.
limits: texturepert hreadmemory ratio	<i>float</i>	0.5	
limits: threads	<i>int</i>	0	Set the number of threads that the renderer uses. If a value of 0 is specified then all CPUs will be utilized for rendering. If a negative value '-N' is specified then prman will use all but N CPUs for rendering.
searchpath: h:archive	<i>string</i>	<i>empty</i>	Search path for Rib archives and DelayedReadArchive.
searchpath: h:dirmap	<i>string</i>	<i>empty</i>	Search path remapping.
searchpath: h:display	<i>string</i>	<i>empty</i>	Search path for display plug-ins.
searchpath: h: procedural	<i>string</i>	<i>empty</i>	Search path for DynamicLoad procedural plug-ins.
searchpath: h:rifilter	<i>string</i>	<i>empty</i>	Search path for Rif plug-ins.
searchpath: h:rixplugin	<i>string</i>	<i>empty</i>	Search path for shading plug-ins.
searchpath: h:shader	<i>string</i>	<i>empty</i>	Search path for OSL pattern plug-ins.
searchpath: h:texture	<i>string</i>	<i>empty</i>	Search path for texture files.

RIB

Name	Type	Default	Description
ribparse:varsubst	<i>string</i>	<i>empty</i>	RIB parse only.

OSL

Name	Type	Default	Description
osl: batched	<i>int</i>	1	Enable OSL optimizations that allow execution of multiple shading points via vectorized instructions on supported processors (AVX, AVX2, AVX512). Value range: true/false.

osl: statistics level	<i>int</i>	0	Set the level of OSL statistics to track. Level 2 and above introduces timing which can have a small effect on performance. Level 4 and above includes internal OSL statistics. Level 5 includes statistics per shader. Other levels reserved for future use.
osl: verbose	<i>int</i>	4	Sets the level of OSL messages to pass through.

Volumes

Name	Type	Default	Description
volume: aggregatespace	<i>string</i>	world	Space in which to compute aggregate volume metadata. Valid spaces are "world" and "camera". Generally the space should be chosen to minimize the range of motion of the volumes.