

Denoiser



Unknown macro: 'html5-video'

RenderMan 25 features a completely new state-of-the-art Denoiser from Disney Research which uses machine learning to resolve partially converged images while maintaining image detail and temporal coherence, making the final result nearly indistinguishable from full convergence.

This advanced denoising technology is production-ready, having been developed by Disney Research in partnership with Pixar, Walt Disney Animation Studios (WDAS), and Industrial Light and Magic, and is based on novel research into how neural networks can denoise images.

This denoiser replaces the current offline denoiser in RenderMan developed by WDAS.

Highlights:

Machine Learning — Pixar's RenderMan Denoiser uses machine learning to remove noise from partially converged images with high accuracy and temporal coherence, resulting in a final image that is nearly indistinguishable from fully converged images. The denoiser has been trained on a broad range of production datasets from VFX at ILM to feature animation at Pixar and Disney.

Production Proven — This advanced denoising technology is reliable as it has been developed by Disney Research and used in production at Pixar, Walt Disney Animation Studios (WDAS), and Industrial Light and Magic, where it has been proven to be highly effective.

Highest Fidelity — The Denoiser stands out in preserving complex, intricate details that could otherwise be lost, making it ideal for demanding applications like feature animation and VFX. It significantly reduces rendering times while maintaining high fidelity, and has been occasionally called "magic" by industry veterans.

You can run the denoiser from either the [command line](#), the [UI](#), or from inside [Maya](#), [Katana](#), [Houdini](#), [Solaris](#), or [Blender](#)