Machine-learning Denoising in Feature Film Production

Henrik Dahlberg
Industrial Light & Magic

David Adler
Walt Disney Animation Studios

Jeremy Newlin
Pixar Animation Studios

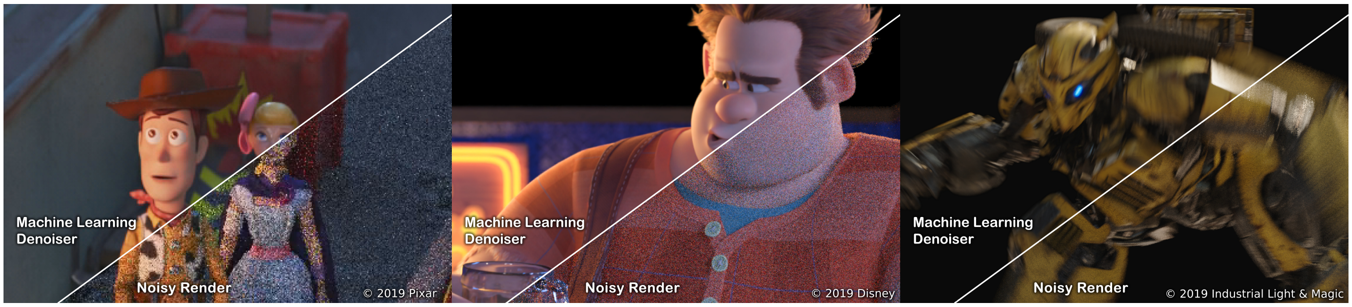

Figure 1: Production examples from *Toy Story 4* (left), *Ralph Breaks the Internet* (middle) and *Bumblebee* (right). In each case we show the clean image (upper-left) our denoiser produces from a noisy render (lower-right).

ABSTRACT

We present our experience deploying and using machine learning denoising of Monte Carlo renders in the production of animated feature films such as Pixar's *Toy Story 4*, Disney Animation's *Ralph Breaks the Internet* and Industrial Light & Magic's visual effects work on photo-realistic films such as *Aladdin (2019)*. We show what it took to move from an R&D implementation of "Denoising with Kernel Prediction and Asymmetric Loss Functions" [Vogels et al. 2018] to a practical tool in a production pipeline.

CCS CONCEPTS

• **Computing methodologies → Rendering**; *Image processing*; *Neural networks*; • **Applied computing** → *Media arts.*

KEYWORDS

Monte Carlo denoising, neural networks, production, rendering, collaboration

ACM Reference Format:
Henrik Dahlberg, David Adler, and Jeremy Newlin. 2019. Machine-learning Denoising in Feature Film Production. In *Proceedings of SIGGRAPH '19 Talks.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3306307.3328150

1 INTRODUCTION

Path-traced Monte Carlo rendering is now ubiquitous in feature animation and visual effects production [Keller et al. 2015]. Its random sampling of light paths causes visible noise in the final image unless the render is given a long time to converge. A common solution to this problem is to perform image-space denoising after the render completes.

Our three studios in collaboration with Disney Research Studios recently put machine learning denoising into production (Figure 1) and found it a significant improvement over past hand-designed linear regression denoisers. We discuss what we learned when implementing the approach of Vogels et al. [2018] as a practical production tool used on animated feature film such as Pixar's *Toy Story 4*, Disney Animation's *Ralph Breaks the Internet* and Industrial Light & Magic's visual effects work on photo-realistic films such as *Aladdin (2019)*. Our three production studios represent a variety of visual appearance, renderers, and production processes.

2 DENOISING APPROACH

Our denoiser follows the general approach of Vogels et al. [2018]. The neural network architecture has been modified further to speed up denoising.

Temporal Denoising: The denoiser reads a temporal region of seven frames, letting it denoise more effectively and improving temporal stability. At the start or end of a sequence of frames the denoiser supplies black imagery and features for the missing frames. During training some images were randomly replaced with black, so the network learned to tolerate missing images.

Inputs to the Denoiser: The input to training and to denoising is images rendered by Disney's Hyperion Renderer [Burley et al. 2018] or RenderMan. These include specular and diffuse decomposition of the rendered color, alpha (opacity), surface color and surface normal direction features, variance estimates of all of those, and motion vectors. Pixar has a layer of configurability on top of RenderMan that allows additional separation of the illumination. For example, some image content benefits from filtering subsurface in its own buffer.

Motion Vectors: The denoiser warps frames to match the center frame's pixel positions as guided by motion vectors. Pixar generates motion vectors using optical flow, which can track the motion

SIGGRAPH '19 Talks, July 28 - August 01, 2019, Los Angeles, CA, USA

ACM ISBN 978-1-4503-6317-4/19/07.
https://doi.org/10.1145/3306307.3328150

of volumetric data and can capture motion behind glass – advantages over rendered motion vectors. Pixar's networks were trained with optical flow motion vectors; those networks do not generalize well for denoising with motion vectors output from RenderMan - particularly on content that is already challenging, like hair.

Denoising passes: The denoiser separately denoises 1. the render's surface color output; 2. diffuse divided by rendered surface color; and 3. specular. Then it multiplies denoised diffuse by denoised albedo and adds the result to specular to get the final denoised color. Alpha is optionally denoised in an additional pass.

Training: The neural network is trained on thousands of 7-frame sequences of frames from multiple renderers, productions, and studios. For each frame we give it a nearly noise-free reference render and a set of noisy renders with different numbers of samples per pixel, and consequently different noise levels. We find that the network generalizes well without retraining for a particular production. Disney used the denoiser extensively on *Ralph Breaks the Internet* without training on any of that show's imagery.

Alpha Denoising: The denoiser let us reduce the number of samples per pixel for color so much that certain content will have an unacceptably noisy alpha (since alpha comes from the same samples as color). For example, vegetation with nothing behind it in-render will composite poorly due to unresolved alpha. So we sometimes also denoise alpha. Neural networks trained only on color data have not generalized perfectly to alpha. For example, a color network may produce values of 0.9999 for a solid alpha input, which can create "pinhole" bright spots during compositing. Pixar solved this by retraining the denoiser specifically on alpha examples.

3 DEPLOYMENT AND PIPELINE INTEGRATION

Stand-alone versus Nuke Plug-in Denoising: Disney and ILM denoise in a stand-alone Python application. Pixar re-implemented denoising as a C++ Nuke plug-in, hoping that would give more compositing flexibility, though this increased its memory use.

CPU versus GPU denoising: The denoiser can run its neural networks on the CPU or GPU. The GPU version runs faster, but neither Disney's, Pixar's nor ILM's render farms have sufficient access to GPUs so we currently run on CPU.

Pipeline & Render Queue Tools: The studios have extended existing pipeline and queue tools to denoise a frame as soon as its 7-frame neighborhood of frames has completed rendering. The prior denoiser used a 3-frame neighborhood.

4 IMPACT ON PRODUCTION

Effectiveness: This denoiser produces few artifacts and reduces noise much better than Disney's previous denoiser [described in Burley et al. 2018, section 5.1.2], which was state-of-the-art before machine learning denoisers were developed. This not only saves computation time, but also allows artists to go through more creative iteration cycles. Disney quickly adopted the new denoiser, in the middle production of Disney's *Ralph Breaks the Internet*. Due to its later release date, *Toy Story 4* had more time to plan for the adoption of the new denoiser, which afforded them more flexibility in render farm projections and show rendering needs. At ILM, the denoiser was deployed in time to help facilitate the delivery of *Aquaman* and has been used to throughout the production of *Aladdin (2019)* and other shows.

Detail Preservation and Over-resolution: ILM's highly detailed photorealistic renders present a challenge for detail preservation. The new denoiser is the first denoiser to perform well enough to be adopted for production use at ILM. Over-smoothing can still occur when denoising renders with high frequency detail, however. ILM sometimes renders with significantly fewer samples per pixel but at a higher resolution, then denoises, and then down-scales the result, which often leads to less loss of detail.

Artist Control of Denoising Strength: Vogels et al. [2018] introduced denoiser training using "asymmetric loss functions". ILM uses this feature heavily for per-pixel control of the intensity of denoising. Artists can adjust the value of this parameter to avoid over-smoothing in cases where detail retention is important, e.g. for close-up character shots. We can also adjust the denoising intensity based on the amount of motion blur.

Render Time Savings: The new denoiser reduced render times at Disney and Pixar by a factor of 2 to 4 compared to their prior denoisers and even more compared to not denoising. Production generally took some of the improvement as reduced render time and some as improved visual quality. At ILM the denoiser reduced worst render times by a factor of 2 to 3 compared to not denoising, as well as reducing the burden on the digital compositors to clean up noise and fireflies manually.

REFERENCES

Brent Burley, David Adler, Matt Jen-Yuan Chiang, Hank Driskill, Ralf Habel, Patrick Kelly, Peter Kutz, Yining Karl Li, and Daniel Teece. 2018. The Design and Evolution of Disney's Hyperion Renderer. *ACM Transactions on Graphics (TOG)* 37, 3 (2018), 33. https://doi.org/10.1145/3182159

Alexander Keller, Luca Fascione, Marcos Fajardo, Iliyan Georgiev, Per H Christensen, Johannes Hanika, Christian Eisenacher, and Gregory Nichols. 2015. The path tracing revolution in the movie industry.. In *SIGGRAPH Courses*. 24–1. https://doi.org/10.1145/2776880.2792699

Thijs Vogels, Fabrice Rousselle, Brian McWilliams, Gerhard Röthlin, Alex Harvill, David Adler, Mark Meyer, and Jan Novák. 2018. Denoising with kernel prediction and asymmetric loss functions. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 124. https://doi.org/10.1145/3197517.3201388