


Neural Rendering: NeRF Takes a Walk in the Fresh Air

Originally published at <https://www.unite.ai/neural-rendering-nerf-takes-a-walk-in-the-fresh-air/>



A collaboration between Google Research and Harvard University has developed a new method to create 360-degree neural video of complete scenes using [Neural Radiance Fields](#) (NeRF). The novel approach takes NeRF a step closer to casual abstract use in any environment, without being restricted to [tabletop models](#) or [closed interior scenarios](#).



CLICK TO VIEW ANIMATED GIF

See end of article for full video. Source: <https://www.youtube.com/watch?v=YStDS2-Ln1s>

Mip-NeRF 360 can handle extended backgrounds and 'infinite' objects such as the sky, because, unlike most previous iterations, it sets limits on the way light rays are interpreted, and creates boundaries of attention that rationalize otherwise lengthy training times. See the new accompanying video embedded at the end of this article for more examples, and an extended insight into the process.

The [new paper](#) is titled *Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields*, and is led by Senior Staff Research Scientist at Google Research Jon Barron.

To understand the breakthrough, it's necessary to have a basic comprehension of how neural radiance field-based image synthesis functions.

What is NeRF?


It's problematic to describe a NeRF network in terms of a 'video', as it's nearer to a fully 3D-realized but *AI-based* virtual environment, where multiple viewpoints from single photos (including video frames) are used to stitch together a scene that technically exists only in the latent space of a machine learning algorithm – but from which an extraordinary number of viewpoints and videos can be extracted at will.



CLICK TO VIEW ANIMATED GIF

A depiction of the multiple camera capture points that provide the data which NeRF assembles into a neural scene (pictured right).

Information derived from the contributing photos is trained into a matrix that's similar to a traditional [voxel grid](#) in CGI workflows, in that every point in 3D space ends up with a value, making the scene navigable.



A traditional voxel matrix places pixel information (which normally exists in a 2D context, such as the pixel grid of a JPEG file) into a three-dimensional space. Source: ResearchGate

After calculating the interstitial space between photos (if necessary), the path of each possible pixel of each contributing photo is effectively 'ray-traced' and assigned a color value, including a transparency value (without which the neural matrix would be completely opaque, or completely empty).


Like voxel grids, and *unlike* CGI-based 3D coordinate space, the 'interior' of a 'closed' object has no existence in a NeRF matrix. You can split open a CGI drum kit and look inside, if you like; but as far as NeRF is concerned, the existence of the drum kit ends when the opacity value of its surface equals '1'.

A Wider View of a Pixel

Mip-NeRF 360 is an extension of [research from March 2021](#), which effectively introduced efficient anti-aliasing to NeRF without exhaustive supersampling.


NeRF traditionally calculates just one pixel path, which is inclined to produce the kind of '[jaggies](#)' that characterized early internet image formats, as well as [earlier games systems](#). These jagged edges were solved by various methods, usually involving sampling adjacent pixels and finding an average representation.

Because traditional NeRF only samples that one pixel path, Mip-NeRF introduced a 'conical' catchment area, like a wide-beam torch, that provides enough information about *adjacent* pixels to produce economical antialiasing with improved detail.



The conical cone catchment that Mip-NeRF uses is sliced up into conical frustums (lower image), which are further 'blurred' to create vague Gaussian spaces that can be used to calculate the accuracy and aliasing of a pixel. Source: <https://www.youtube.com/watch?v=EpH175PY1A0>

The improvement over a standard NeRF implementation was notable:




Mip-NeRF (right), released in March 2021, provides improved detail through a more comprehensive but economical aliasing pipeline, rather than just 'blurring' pixels to avoid jagged edges. Source: <https://jonbarron.info/mipnerf/>

NeRF Unbounded

The March paper left three problems unsolved with respect to using Mip-NeRF in unbounded environments that might include very distant objects, including skies. The new paper solves this by applying a [Kalman-style](#) warp to the Mip-NeRF Gaussians.

Secondly, larger scenes require greater processing power and extended training times, which Mip-NeRF 360 solves by 'distilling' scene geometry with a small 'proposal' [multi-layer perceptron](#) (MLP), which pre-bounds the geometry predicted by a large standard NeRF MLP. This speeds training up by a factor of three.


Finally, larger scenes tend to make discretization of the interpreted geometry ambiguous, resulting in the kind of artifacts gamers might be familiar with when game output 'tears'. The new paper addresses this by creating a new regularizer for Mip-NeRF ray intervals.



On the right, we see unwanted artifacts in Mip-NeRF due to the difficulty in bounding such a large scene. On the left, we see that the new regularizer has optimized the scene well enough to remove these disturbances.

To find out more about the new paper, check out the video below, and also the [March 2021 video introduction](#) to Mip-NeRF. You can also find out more about NeRF research by checking out [our coverage so far](#).

Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields



Originally published 25th November 2021

21st December 2021, 12:25pm – Replaced dead video. – MA

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.

Portfolio site: martinanderson.ai