

NeRF: Training Drones in Neural Radiance Environments

Originally published at <https://www.unite.ai/nerf-training-drones-in-neural-radiance-environments/>



Researchers from Stanford University have devised a new way of training drones to navigate photorealistic and highly accurate environments, by leveraging the recent [avalanche of interest](#) in Neural Radiance Fields (NeRF).



CLICK TO VIEW ANIMATED GIF

Drones can be trained in virtual environments mapped directly from real-life locations, with no need for specialized 3D scene reconstruction. In this image from the project, wind disturbance has been added as a potential obstacle for the drone, and we can see the drone being momentarily diverted from its trajectory and compensating at the last moment to avoid a potential obstacle.

Source: <https://mikh3x4.github.io/nerf-navigation/>

The method offers the possibility for interactive training of drones (or other types of objects) in virtual scenarios that automatically include volume information (to calculate collision avoidance), texturing drawn directly from real-life photos (to help train drones' image recognition networks in a more realistic fashion), and real-world lighting (to ensure a variety of lighting scenarios get trained into the network, avoiding over-fitting or over-optimization to the original snapshot of the scene).



CLICK TO VIEW ANIMATED GIF

A couch-object navigates a complex virtual environment which would have been very difficult to map using geometry capture and retexturing in traditional AR/VR workflows, but which was recreated automatically in NeRF from a limited number of photos. Source:


<https://www.youtube.com/watch?v=5JjWpv9BaaE>

Typical NeRF implementations do not feature trajectory mechanisms, since most of the slew of NeRF projects in the last 18 months have concentrated on other challenges, such as [scene relighting](#), [reflection rendering](#), [compositing](#) and [disentanglement](#) of captured elements. Therefore the new paper's primary innovation is to implement a NeRF environment as a navigable space, without the extensive equipment and laborious procedures that would be necessary to model it as a 3D environment based on sensor capture and CGI reconstruction.

NeRF as VR/AR


The new [paper](#) is titled *Vision-Only Robot Navigation in a Neural Radiance World*, and is a collaboration between three Stanford departments: Aeronautics and Astronautics, Mechanical Engineering, and Computer Science.

The work proposes a navigation framework that provides a robot with a pre-trained NeRF environment, whose volume density delimits possible paths for the device. It also includes a filter to estimate where the robot is inside the virtual environment, based on image-recognition of the robot's on-board RGB camera. In this way, a drone or robot is able to 'hallucinate' more accurately what it can expect to see in a given environment.



The project's trajectory optimizer navigates through a NeRF model of Stonehenge that was generated through photogrammetry and image interpretation (in this case, of mesh models) into a Neural Radiance environment. The trajectory planner calculates a number of possible paths before establishing an optimal trajectory over the arch.

Because a NeRF environment features fully modeled occlusions, the drone can learn to calculate obstructions more easily, since the neural network behind the NeRF can map the relationship between occlusions and the way that the drone's onboard vision-based navigation systems perceive the environment. The automated NeRF generation pipeline offers a relatively trivial method of creating hyper-real training spaces with only a few photos.




CLICK TO VIEW ANIMATED GIF

The online replanning framework developed for the Stanford project facilitates a resilient and entirely vision-based navigation pipeline.

The Stanford initiative is among the first to consider the possibilities of exploring a NeRF space in the context of a navigable and immersive VR-style environment. Neural Radiance fields are an emerging technology, and currently subject to multiple academic efforts to optimize their high computing resource requirements, as well as to disentangle the captured elements.

NeRF Is Not (Really) CGI

Because a NeRF environment is a navigable 3D scene, it's become a misunderstood technology since its emergence in 2020, often widely-perceived as a method of automating the creation of meshes and textures, rather than replacing 3D environments familiar to viewers from Hollywood VFX departments and the fantastical scenes of Augmented Reality and Virtual Reality environments.



NeRF extracts geometry and texture information from a very limited number of image viewpoints, calculating the difference between images as volumetric information
<https://www.matthewtancik.com/nerf>

In fact, the NeRF environment is more like a 'live' render space, where an amalgamation of pixel and lighting information is retained and navigated in an active and running neural network.

The key to NeRF's potential is that it only requires a limited number of images in order to recreate environments, and that the generated environments contain all necessary information for a high-fidelity reconstruction, without the need for the services of modelers, texture artists, lighting specialists and the hordes of other contributors to 'traditional' CGI.

Semantic Segmentation

Even if NeRF effectively constitutes 'Computer-Generated Imagery' (CGI), it offers an entirely different methodology, and a highly-automated pipeline. Additionally, NeRF can isolate and 'encapsulate' moving parts of a scene, so that they can be added, removed, sped up, and generally operate as discrete facets in a virtual environment – a capability that is far beyond the current state-of-the-art in a 'Hollywood' interpretation of what CGI is.



CLICK TO VIEW ANIMATED GIF

A [collaboration](#) from Shanghai Tech University, released in summer of 2021, offers a method to individuate moving NeRF elements into 'pastable' facets for a scene. Source:
<https://www.youtube.com/watch?v=Wp4HfOwFGP4>

Negatively, NeRF's architecture is a bit of a 'black box'; it's not currently possible to extract an object from a NeRF environment and directly manipulate it with traditional mesh-based and image-based tools, though a number of research efforts are beginning to make breakthroughs in deconstructing the matrix behind NeRF's neural network live render environments.

Vision-Only Robot Navigation in a Neural Radiance World



□

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

Contact: martin@martinanderson.ai



Discover More