

Nerfies: Deformable Neural Radiance Fields

Keunhong Park¹, Utkarsh Sinha², Jonathan T. Barron², Sofien Bouaziz², Dan B Goldman², Steven M. Seitz^{1,2}, Ricardo Martin-Brualla²

¹University of Washington ²Google Research

Abstract

We present the first method capable of photorealistically reconstructing deformable scenes using photos/videos captured casually from mobile phones.

Abstract (long)

We present the first method capable of photorealistically reconstructing a non-rigidly deforming scene using photos/videos captured casually from mobile phones.

Our approach augments neural radiance fields (NeRF) by optimizing an additional continuous volumetric deformation field that warps each observed point into a canonical 5D NeRF. We observe that these NeRF-like deformation fields are prone to local minima, and propose a coarse-to-fine optimization method for coordinate-based models that allows for more robust optimization. By adapting principles from geometry processing and physical simulation to NeRF-like models, we propose an elastic regularization of the deformation field that further improves robustness.

We show that Nerfies can turn casually captured selfie photos/videos into deformable NeRF models that allow for photorealistic renderings of the subject from arbitrary viewpoints, which we dub “nerfies”. We evaluate our method by collecting data using a rig with two mobile phones that take time-synchronized photos, yielding train/validation images of the same pose at different viewpoints. We show that our method faithfully reconstructs non-rigidly deforming scenes and reproduces unseen views with high fidelity.

Video

Visual Effects

Using *nerfies* you can create fun visual effects. This Dolly zoom effect would be impossible without nerfies since it would require going through a wall.

Matting

As a byproduct of our method, we can also solve the matting problem by ignoring samples that fall outside of a bounding box during rendering.

Animation

Interpolating states

We can also animate the scene by interpolating the deformation latent codes of two input frames. Use the slider here to linearly interpolate between the left frame and the right frame.

Re-rendering the input video

Using *Nerfies*, you can re-render a video from a novel viewpoint such as a stabilized camera by playing back the training deformations.

Related Links

There's a lot of excellent work that was introduced around the same time as ours.

Progressive Encoding for Neural Optimization introduces an idea similar to our windowed position encoding for coarse-to-fine optimization.

D-NeRF and NR-NeRF both use deformation fields to model non-rigid scenes.

Some works model videos with a NeRF by directly modulating the density, such as Video-NeRF, NSFF, and DyNeRF.

Check out Frank Dellaert's survey on recent NeRF papers and Yen-Chen Lin's curated list of NeRF papers.