

NeuralDiff: Segmenting 3D objects that move in egocentric videos

Vadim Tschernezki^{1,2} Diane Larlus² Andrea Vedaldi¹

¹Visual Geometry Group
University of Oxford

{vadim, vedaldi}@robots.ox.ac.uk

²NAVER LABS Europe

diane.larlus@naverlabs.com

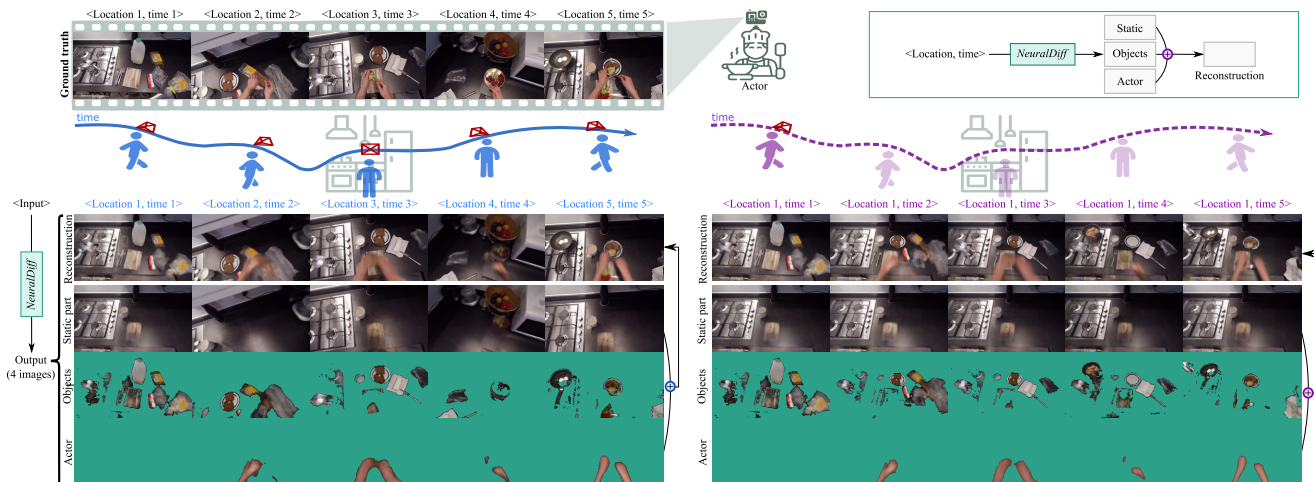


Figure 1: Given a video sequence captured from an egocentric viewpoint, we segment all the objects that the actor/observer interacts with. We achieve this by means of a neural architecture, NeuralDiff, that learns to decompose each frame into a static background and a dynamic foreground, comprising the manipulated objects, which seldomly move in the sequence, and the actor’s body, which moves continually and heavily occludes the scene. The neural network contains three streams that, via different inductive biases, reconstruct the background, the objects and the actor in 3D, and is thus able to render images and their segmentations also for viewpoints that do not exist in the original video sequence (see the left part of the figure, where the camera is assumed to remain at its initial position while the action unfolds).

Abstract

Given a raw video sequence taken from a freely-moving camera, we study the problem of decomposing the observed 3D scene into a static background and a dynamic foreground containing the objects that move within the scene. This task is reminiscent of the classic background subtraction problem, but is significantly harder because all parts of the scene, static and dynamic, generate a large apparent motion due to the camera large viewpoint change and parallax. In particular, we consider egocentric videos and further separate the dynamic component into objects and the actor that observes and moves them. We achieve this factorization by reconstructing the video via a triple-stream neural rendering network that explains the different motions based on corresponding inductive biases. We demonstrate that our method can successfully separate the differ-

ent types of motion, outperforming recent neural rendering baselines at this task, and can accurately segment the moving objects. We do so by assessing the method empirically on challenging videos from the EPIC-KITCHENS dataset which we augment with appropriate annotations to create a new benchmark for the task of dynamic object segmentation on unconstrained video sequences, for complex 3D environments.

1. Introduction

Given a video capturing a complex 3D scene, we consider the problem of segmenting the scene objects that move independently of the camera. Motion is a powerful cue for discovering and learning visual objects in an unsupervised manner. In fact, ‘detachability’, namely the possibility of moving a body independently of the rest of the scene, is

used by Gibson [8] to *define* objects. However, *measuring* detachment given only raw visual observations as input is not an easy task.

If the video is taken from the viewpoint of a static camera, the problem of separating the static background from the moving foreground reduces to background subtraction. However, classic background subtraction techniques are inapplicable if the camera undergoes a motion that induces significant parallax. We may call this more challenging scenario *wide-baseline background subtraction*.

To understand this concept, consider for example an egocentric video of a person cooking. This actor intervenes in the scene by moving (and transforming) objects. However, egomotion is the dominant effect: by comparison, objects move only sporadically, and in a way that is hardly distinguishable from the much larger apparent motion induced by the viewpoint change. Extracting the moving objects automatically is thus very difficult, and essentially impossible for traditional background subtraction techniques.

One may use motion segmentation techniques to separate a scene in different motion components. However, these techniques generally require correspondences (*e.g.*, optical flow), they reason locally, across a handful of frames, and usually avoid explicit 3D reasoning. In short, they are of difficult applicability to video sequences such as the ones in Figure 1, comprising many small rigid objects that move only occasionally throughout a long sequence.

In this paper, we propose to leverage recent progress in neural rendering techniques [18] to develop a motion *analysis* tool to achieve the desired segmentation. We build on the ability of neural rendering to reconstruct accurately the appearance of a rigid 3D scene under a variable viewpoint, without requiring dense correspondences. Given the reconstruction of the background, it is then possible to measure the more subtle appearance ‘differences’ induced by the objects that move *independently* of the camera.

We further note that the 3D objects manipulated in the video also contain significant structure. Specifically, they move in ‘bursts’, changing their state as they are manipulated, but remaining otherwise rigidly attached to the background. We thus extend the neural renderer to also reconstruct the object appearance using a slowly-varying time encoding for them. In fact, we go one step further and introduce a third neural rendering stream that captures the actor observing and moving the objects. The intuition is that the actor moves continually, in a way that occludes the scene, with a motion linked to the camera and not to the scene (being the observer), leading to significantly different dynamics compared to the background and foreground objects.

Our technical contribution is thus a three-stream neural rendering architecture, where the streams model respectively i) the static background, ii) the dynamic foreground objects, and iii) the actor. Those are then composed to ex-

plain the video as a whole (see Figure 2). We design the streams differently, in order to incorporate inductive biases that match the statistics of each layer (background, foreground, actor). These inductive biases depart from previous neural rendering models because of the structure of the foreground, which is composed of several objects being manipulated at different intervals, and of the actor, a deformable body attached to the camera and not to the background.

The resulting analysis-via-synthesis method shows that neural rendering techniques are not only useful for synthesis, but also for analysis. In particular, we are the first to demonstrate the effectiveness of these techniques in interpreting challenging egocentric videos, providing cues for the extraction of detached objects in scenes with a complex 3D structure and dynamics.

We focus our empirical evaluation on egocentric videos because, with the emergence of AR, they are becoming increasingly popular and have the advantage of showing the interaction of actors with their environment. We expect this kind of videos to provide an enormous wealth of information for computer vision, particularly with the recent introduction of Ego4D [9]. They are also particularly challenging to process, providing an excellent test scenario for this class of algorithms.

For evaluation, we augment the EPIC-KITCHENS dataset [5] and manually segment all objects that move at some point during the scene, *i.e.* that are thus detached. With these annotations, we can assess neural renderers not only in terms of new view synthesis quality, but also, and more to the point, in terms of their ability to separate videos in the various dynamic components. With this, we also define a new *benchmark* for measuring progress in the challenging task of dynamic object segmentation in complex videos, thus inviting further research in the area. Using this data, which we call EPIC-Diff, we show that our model outperforms the direct application of existing neural rendering approaches such as NeRF [18] or NeRF-W [15] in the wide-baseline background subtraction problem.

2. Related work

Although unique, our research problem relates to several existing research directions that we mention below.

Background subtraction. Background subtraction techniques have been used for detecting moving objects in video sequences (see survey [2]). They typically assume that only the foreground moves, and many methods rely on a background initialization step which assumes that the first few frames give a good estimate of the background’s appearance. The background model can be updated during the sequence, but these methods fail when the background varies as much as in the egocentric videos that we consider [11].

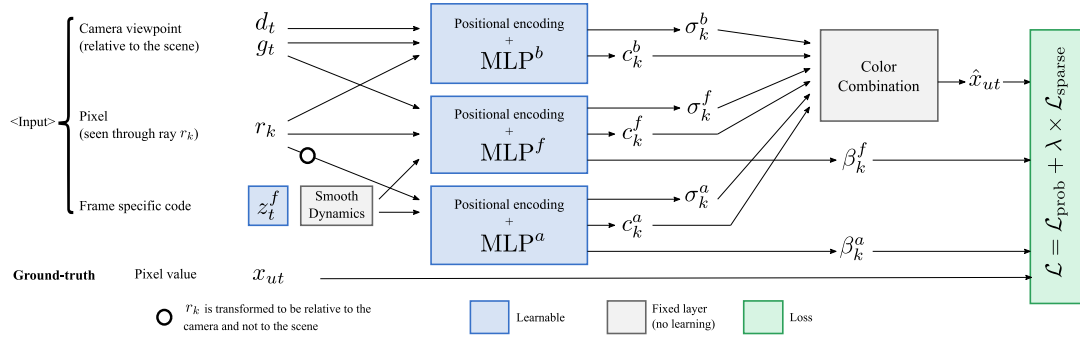


Figure 2: **Overview of NeuralDiff.** Given only the camera viewpoint g_t and a frame specific code z_t^f (learned latent variable), our three stream architecture learns to predict the color value of pixel x_{ut} by combining information coming from the static model of the background, and from two dynamic components, one for the foreground objects, and one for the actor. The parameters of this model are learned using a probabilistic loss \mathcal{L} .

Motion segmentation. Closely related to background subtraction, motion segmentation is the more general task of decomposing a video into individually moving objects [39, 16, 21]. These techniques generally rely on optical flow, which is subject to ambiguities [19], and reasons only locally without a 3D representation. Occlusions are one of the main challenges [37]. Many methods fail when dynamic objects temporarily remain static, even for a few frames [39, 16]. These issues prevent applying standard motion segmentation methods to egocentric videos (where objects rarely move and the actor heavily occludes the scene).

Discovering and segmenting objects in videos. This long-standing problem is related to background subtraction and motion segmentation [1, 22, 31, 10, 35]. For example, one can use a probabilistic model that acts upon optical flow to segment moving objects from the background [1]. In [29, 3], pixel trajectories and spectral clustering are combined to produce motion segments. The method of [17] uses image collections to reconstruct urban scenes and to discover their dynamic elements such as billboards or street art, by clustering 3D points in space and time. Recent work revisits classical motion segmentation techniques from a data-driven perspective [36, 33, 35], *e.g.* using physical motion cues to learn 3D representations [30] or learning a factored scene representation with neural rendering [38]. Our work departs from these approaches and learns a holistic representation capable of handling occlusions and sporadically moving objects, without the need for multi-view videos as in [38]. We approach the problem on complex real-world data as opposed to clean synthetic ones as in [30].

Neural rendering. Neural rendering is a way of synthesizing novel views using a neural architecture and classic volume rendering techniques. It was introduced as Neural Radiance Fields (NeRF) in [18] for static scenes. By default, it poorly handles dynamic scenes and occlusions.

It is extended in [15] as NeRF-W for the case of unconstrained photo collections to deal with transient objects occluding the static scene. Another research direction is studied in [34, 28] to model the 3D geometry of one or several static objects. Recent work has also focused on modeling dynamic scenes, mostly with monocular videos as input [23, 13, 25, 4, 7, 32]. Most approaches [23, 25, 4, 24] combine a canonical model of the object with a deformation network, or warp space [32], still starting from a canonical volume. Closer to our work, [13] and [7] combine a static NeRF model and a dynamic one. Yet, none of those methods explicitly tackle the segmentation of 3D objects in such long and challenging video sequences.

3. Method

Given a video sequence x , we wish to extract a corresponding mask m that separates the foreground objects from the background. We define as background the part of the scene that remains static throughout the entire video, generating an apparent motion only due to the camera viewpoint change. We define as foreground any object that moves independently of the camera in at least one frame.

Traditional background subtraction techniques solve a similar problem by predicting the appearance of each video frame *as if* the foreground objects were removed; given this prediction, the foreground objects can be segmented by taking the difference between the measured and predicted images. Predicting the appearance of the background under an occluding foreground object is relatively easy for a static camera, where correspondences with other video frames where the object is not present can be trivially established. However, the prediction is much more challenging if the viewpoint is also allowed to change.

In order to solve this problem, we build on recent neural rendering techniques such as NeRF [18] that can pre-

dict effectively the appearance of a static object, in our case the rigid background, under a variable viewpoint. In fact, we suggest that the foreground objects, which change over time, can *also* be captured via a (distinct) neural rendering function, further promoting separation of background and foreground. Next, we first discuss neural rendering in general, and then introduce our model.

3.1. Neural rendering

We base our method on NeRF [18], which we summarize here. A video x is a collection $(x_t)_{t \in [0, \dots, T-1]}$ of T video frames, each of which is an RGB image $x_t \in \mathbb{R}^{3 \times H \times W}$. The video frames are a function $x_t = h(B, F_t, g_t)$ of the static background B , the variable foreground F_t and a moving camera $g_t \in SE(3)$, where $SE(3)$ is the group of Euclidean transformations. The motion g_t is assumed to be known, estimated using an off-the-shelf SfM algorithm such as COLMAP [26, 27]. The background and foreground components comprise the shape and reflectance of the 3D surfaces in the scene as well as the illumination.

Rather than attempting to invert h to recover B and F_t , which amounts to inverse rendering, neural rendering *learns* the mapping h directly, as a neural network f , $h(B, F_t, g_t) \approx f(g_t, t)$, providing time and viewpoint to reconstruct the corresponding video frame x_t . By a careful design of the function f , the learning process can induce a factorization of viewpoint g and time t , thus generalizing f to new (unobserved) viewpoints. NeRF additionally assumes that the scene is static, meaning that the variable foreground F_t is empty, so the function simplifies to $f(g_t)$.

The model $f(g_t)$ is further endowed with a specific structure, which is key to successful learning [40]. Specifically, the color $x_{ut} \in \mathbb{R}^3$ of pixel $u \in \Omega = \{0, \dots, H-1\} \times \{0, \dots, W-1\}$ is obtained by a volumetric sampling process that simulates ray casting. One ‘shoots’ a ray $r_k = \ell_k K^{-1}(u)$ along the viewing direction of pixel u , where $K : \mathbb{R}^2 \times \{1\} \rightarrow \Omega$ is the camera calibration function and $\ell_0 < \dots < \ell_{M+1} \in \mathbb{R}_+$ are the sampled depths.

The pixel color is obtained by averaging the color of the 3D points $g_t r_k$ along the ray, weighed by the probability that a photon emanates from the point and reaches the camera. A neural network, $(\sigma_k^b, c_k^b) = \text{MLP}^b(g_t r_k, d_t)$, estimates the density $\sigma_k^b \in \mathbb{R}_+$ and the color $c_k^b \in \mathbb{R}^3$ of each point $g_t r_k$, where the superscript b denotes the fact that the quantities refer to the background ‘material’ and d_t is the unit-norm viewing direction.

The probability that a photon is transmitted while traveling through the ray segment (r_k, r_{k+1}) is defined to be $T_k^b = e^{-\delta_k \sigma_k^b}$ where the quantity $\delta_k = |r_{k+1} - r_k|$ is the length of the segment. This definition is consistent with the fact that the probability of transmission across several segments is the product of the individual transmission proba-

bilities. We can thus write the color of pixel u as:

$$x_{ut} = f_u(g_t) = \sum_{k=0}^M v_k (1 - T_k^b) c_k^b, \quad v_k = \prod_{q=0}^{k-1} T_q^b. \quad (1)$$

The network is trained by minimizing the reconstruction error $\|x - f(g_t)\|$, thus fitting a single video at a time.

3.2. Dynamic components

The method discussed above assumes that the scene is rigid. In our case, reconstructing the scene is more complex due to the variable foreground F_t . In order to capture this dependency, on top of the background MLP (MLP^b), we introduce a foreground-specific MLP, $(\sigma_k^f, c_k^f, \beta_k^f) = \text{MLP}^f(g_t r_k, z_t^f)$. It produces a ‘foreground’ occupancy σ^f and color c^f . Additionally, it predicts an uncertainty score β_k^f whose role is clarified in Section 3.3. We also introduce a dependency on a frame-specific code $z_t^f \in \mathbb{R}^D$, capturing the properties of the foreground that change over time.

The color x_{ut} of a pixel u is obtained by composition of multiple materials \mathcal{S} (e.g. the background and the foreground, so $\mathcal{S} = \{b, f\}$):

$$x_{ut} = f_u(g_t, z_t) = \sum_{k=0}^M v_k \left(\sum_{p \in \mathcal{S}} w^p(T_k) c_k^p \right),$$

$$\text{where } v_k = \prod_{q=0}^{k-1} \prod_{p \in \mathcal{S}} T_q^p \quad (2)$$

The factor v_k requires a photon to be transmitted from the camera to point r_k through the different materials (hence the transmission probabilities are multiplied). The weights $w^p(T_k)$ mix the colors of the materials proportionally to their density. Following NeRF-W [15], we can simply set

$$w^p(T_k) = 1 - T_k^p \in [0, 1]. \quad (3)$$

Smooth dynamics. The model (MLP) weights and the frame-specific parameters z_t can be optimized by minimizing the loss across all input frames

$$\min_{f, z_1, \dots, z_T} \frac{1}{T|\Omega|} \sum_{t=1}^T \sum_{u \in \Omega} \|x_t - f(g_t, z_t)\|^2$$

in an auto-decoder fashion. However, the foreground, while dynamic, does not change arbitrarily from frame to frame. In particular, most foreground objects are in most frames rigidly attached to the background. Because of this, the dependency on independent frame-specific codes z_t makes little sense; we replace it with a low-rank expansion of the trajectory of states, setting $z_t = B(t)\Gamma$ where $B(t) \in \mathbb{R}^P$ is a

simple handcrafted (fixed) basis and the motion $\Gamma \in \mathbb{R}^{P \times D}$ are coefficients such that $P \ll T$. We take in particular $B(t) = [1, t, \sin 2\pi t, \cos 2\pi t, \sin 4\pi t, \cos 4\pi t, \dots]$ to be a deterministic harmonic coding of time (meaning that z_t varies slowly over time.)

The method described above with a static part (as defined in Section 3.1) and a dynamic part describing the objects ($\mathcal{S} = \{b, f\}$) is the basic version of our proposed approach. We refer to it as `NeuralDiff` in what follows.

Improved geometry: capturing the actor. In egocentric videos, we further distinguish the foreground objects manipulated by the actor/observer, which moves sporadically, from the actor’s body, which moves continually. To model the latter, we consider a third MLP tasked with capturing parts of the actor’s body that appear in the frames. Formally, the actor MLP is similar to the foreground MLP: $(\sigma_k^a, c_k^a, \beta_k^a) = \text{MLP}^a(r_k, z_t^a)$, with the key difference that the 3D point r_k is expressed relative to the camera (v.s. $g_t r_k$ which is expressed relative to the world). This is due to the fact that the camera is anchored to the actor’s body, which therefore shows a reduced variability in the reference frame of the camera. By contrast, the background is invariant if expressed in the reference frame of the world; the same is true for the foreground objects when they are not manipulated, which is true most of the time. This inductive bias helps factoring the different materials. Here $\mathcal{S} = \{b, f, a\}$.

We refer to this new flavor as `NeuralDiff+A`.

Improved color mixing. Eq.(3), used in prior work to mix colors from different model components, cannot be justified probabilistically as it amounts to summing non-exclusive probabilities (nothing in the model prevents two or more materials to have non-zero density at a given point).

A principled mixing model is obtained by decomposing the segment δ_k in Pn sub-segments, alternating between the P different materials ($P = |\mathcal{S}|$, e.g. $P = 3$ if the background, foreground and actor are considered). In the limit, we can show that the probability that the photon is absorbed in a subsegment of material p is given by:

$$w^p(T_k) = \frac{\sigma_k^p}{\sum_{q=1}^P \sigma_k^q} \left(1 - \prod_{q=1}^P T_k^q \right). \quad (4)$$

The second factor in parenthesis is, evidently, the probability that the photon is absorbed by any of the materials. The first factor, which involves the densities rather than the probabilities, is the probability that a given material p is responsible for the absorption. Note that, differently from eq. (3), this definition yields $\sum_{p=1}^P w^p(T_k) = 1 - \prod_{q=1}^P T_k^q = 1 - v_k/v_{k-1}$ which is consistent with the definition of transmission probability v_k . A proof of eq. (4) is available in the supplementary material.

The `NeuralDiff` model can thus be improved by taking into account this more principled way of producing pixel

colors in the reconstruction. We refer to this improved version as `NeuralDiff+C`.

Note that the two proposed improvements are complementary and we refer to the method enhanced by both as `NeuralDiff+C+A`.

3.3. Uncertainty and regularization

Uncertainty. The MLPs also predict scalars $\beta_k^p \geq 0$ (where, in practice, $\beta_k^b = 0$ for the background). These are used to express the uncertainty of the color associated to each 3D point r_k for each material p as pseudo-standard deviations (StDs). Following [15], the StD of the rendered color x_{ut} is just the sum of the StDs $\beta_{ut} = \sum_p \beta_{ut}^p$, where β_{ut}^p is obtained via eq. (2) by ‘rendering’ the StDs β_k^p of each 3D material point (it suffices to replace β_k^p for c_k^p in eq. (2)). The StD β_{ut} is used in a Gaussian observation loss as a form of self-calibrated aleatoric uncertainty [20, 12]:

$$\mathcal{L}_{\text{prob}}(f, z_t | x_t, g_t, u) = \frac{\|x_{ut} - f_u(g_t, z_t)\|^2}{2\beta_{ut}^2} + \log \beta_{ut}^2. \quad (5)$$

Sparsity. We follow NeRF-W [15] and further penalize the occupancy of the foreground and actor components using an L^1 penalty:

$$\mathcal{L}_{\text{sparse}}(f, z_t | x_t, g_t, u) = \sum_{p=1}^P \sum_{k=0}^M \sigma_k^p.$$

This is the L^1 norm of the ray occupancies, which encourages the foreground occupancy to be sparse.

Training loss. Finally, the model is trained using the loss $\mathcal{L} = \mathcal{L}_{\text{prob}} + \lambda \mathcal{L}_{\text{sparse}}$ where $\lambda > 0$ is a weight set to 0.01.

3.4. From NeuralDiff to a scene segmentation

Our approach is trained for a reconstruction task, but our primary goal is wide-baseline background subtraction, so we need to extract masks that assign a given pixel to the background, foreground, or actor layers. We do so by constructing three indicator channels and we use eq. (2) to ‘render’ them as a mask $m_{ut} \in \mathbb{R}^3$ — in other words, we simply associate pseudo-colors $(1, 0, 0)$, $(0, 1, 0)$ and $(0, 0, 1)$ to all 3D points of background, foreground and actor, respectively, and use eq. (2) to obtain m_{ut} .

4. EPIC-Diff benchmark

Our goal is to identify any object which moves independently of the camera in a video sequence. We create a suitable benchmark for this task by augmenting the well-known EPIC-KITCHENS dataset [6] with new annotations. EPIC-KITCHENS is an egocentric dataset, with 100 hours of recording, 20M frames, and 90’000 actions performed.



Figure 3: Two frames from EPIC-Diff and their corresponding manual foreground/background segmentation masks.

Data selection. We selected 10 video sequences, each lasts 14 minutes on average. Then we uniformly extracted 1’000 frames from each, and preprocessed them with COLMAP [26] to obtain the camera calibration and extrinsics (motion). We refer to each video sequence as a *scene*. To select the scenes, we considered two constraints. First, the videos should contain a diversity of viewpoints and manipulated objects. Second, COLMAP must successfully reconstruct the scene and register at least 600 frames. We only retain the frames where COLMAP succeeds, which results in an average of 900 frames per scene.

Data annotation. Since our algorithms are unsupervised, we do not need to collect extensive data annotations. We uniformly hold out on average 56 frames for validation (for setting parameters) and 56 for testing. We annotated the latter with segmentation masks to assess background/foreground segmentation. These frames are *not* used to train the model, so they can also be used to assess new-view synthesis. A frame annotation consists of a pixel-level binary segmentation mask, where the foreground contains any pixel that belongs to an object that is observed moving *at any point in time* during the video. Note that the fact that an object is marked as foreground in a frame does not mean it moves in *that* frame; it only means that it moves at least once in the scene. Based on this definition, the foreground mask covers both foreground objects and the actor. We obtain about 560 manual image-level segmentation masks. Examples of these masks are given in Figure 3.

Evaluation. We evaluate the wide-baseline background subtraction task. For this, we use standard segmentation metrics: for each test frame, we sort all pixels based on their foreground score as defined above, and use the ground truth binary mask to compute average precision (AP). We then report mAP by averaging across frames and scenes. We also evaluate the new-view synthesis quality by tasking our model with synthesizing each of the test views, and we measure PSNR. Specifically, we report PSNR for different parts of the scene using the ground-truth segmentation masks: the whole image, the background and the foreground regions.

5. Experiments

5.1. Experimental settings

Implementation details. All reported experiments are based on a PyTorch implementation of NeRF [18] extended with several NeRF-W components [15]. The architecture

combines several interconnected MLPs to model density, color and uncertainty of background, foreground and actor. Their architecture’s details are given in the supplementary material. We reuse the same motion coefficients $z_k^f = z_k^a = z_k \in \mathbb{R}^{17}$ for both foreground and actor (see section 3.2). Positional encoding is used to encode position, motion coefficients, and viewing direction using respectively 10, 10 and 4 frequencies. All models are trained using the Adam optimizer with an initial learning rate of 5×10^{-4} using a cosine annealing schedule [14]. For more details regarding this implementation and differences to NeRF and NeRF-W, see the supplementary material.

Baselines. We compare against the following baselines: (1) **NeRF** [18] uses a single stream. As it has been designed to only capture the static part of a scene (background), we use the per-pixel prediction error as a pseudo foreground score (the larger the error, the more likely for the pixel to belong to the foreground, because NeRF cannot explain it with its static model). (2) **NeRF-BF** trains two parallel NeRF models, one for the background (B) and one for the foreground (F). The F stream is further conditioned on time, but passing a positional-encoded version of the time variable t (frame number) as input, instead of the learnable dynamic frame encoding of our method. (3) **NeRF-W** [15] also contains two interlinked background and foreground streams. Because it was initially proposed for image collections and not videos, by default the foreground parameters z_t^f are learned independently for each frame.¹ This design limits the applicability of NeRF-W, as it is unable to render novel views for frames not available in the training set. We redesign NeRF-W as a baseline for our task by adjusting the frame specific code such that it produces the foreground using the code of its closest neighboring frame from the training set. More precisely, given a test frame t , we set the code for this test view $z_t^f := z_j^f$ where j is the closest frame to t out of all training frames \mathcal{I} (*i.e.* $j = \operatorname{argmin}_i(\{|i - t| : i \in \mathcal{I}\})$). We follow a similar strategy for the frame specific code for appearance (responsible for capturing the photometric variations).

5.2. Results

We compare our method and the different baselines on the EPIC-Diff benchmark presented in the previous section.

Quantitative results. Table 1 compares the different methods according to the four evaluation metrics of EPIC-Diff. This allows to evaluate their capacity to discover and segment 3D objects but also to reconstruct dynamic scenes. We make the following observations.

First, all flavors of our approach largely outperform the existing NeRF and NeRF-W. They also outperform the

¹In NeRF-W, the foreground is designed to capture the transient part of the scene that should be ignored (*e.g.* persons occluding a landmark).




Figure 4: For three different scenes, reconstruction produced by NeRF [18], NeRF-W [15], NeuralDiff and NeuralDiff+C+A. Our method has less ghosting artifacts, captures most moving objects, and shows more details.

Method	mAP	PSNR	PSNR _b	PSNR _f
NeRF [18]	47.8	20.9	22.8	17.6
NeRF-W [15]	59.2	23.2	26.4	18.9
NeRF-BF	64.4	23.8	26.8	19.6
NeuralDiff (ours)	66.7	24.0	27.2	19.8
NeuralDiff+A (ours)	69.1	24.1	27.3	19.9
NeuralDiff+C (ours)	67.4	24.1	27.2	19.9
NeuralDiff+C+A (ours)	67.8	24.2	27.3	20.0

Table 1: **EPIC-Diff**. Average mAP scores, which evaluate foreground segmentation, and average PSNR scores for the full scene (PSNR), the background part (PSNR_b), and the foreground part (PSNR_f) according to the segmentation.

naive NeRF-BF which simply combines two NeRF models respectively modeling the foreground and the background.

Second, including temporal information as an input to neural rendering (either with a low-rank expansion of the trajectory states as in NeuralDiff or using a two stream architecture that processes time transformed with positional encoding as in NeRF-BF) proves to be essential for improving the performance over NeRF and NeRF-W.

Third, we observe that each of the proposed improvements, NeuralDiff+A and NeuralDiff+C, outperform the vanilla version of our approach. The third stream for modeling the actor brings +2.4 mAP points to the segmentation task while the better color model brings +0.7 mAP points. They also slightly improve the frame reconstruc-

tions (PSNR scores).

Finally, we see that combining the two proposed improvements (NeuralDiff+C+A) further improve new-view synthesis but not the foreground segmentation quality. Note however that the segmentation metric does not reflect the ability of the full model to separate foreground into objects and actor (since they are merged together in the annotated masks). This ability is illustrated in Figure 1.

Qualitative results. First, we verify the quality of the views reconstructed by our method. Figure 4 compares three frames from three different scenes, with their reconstruction by NeRF, NeRF-W, NeuralDiff and NeuralDiff+C+A. As already observed [23, 13], NeRF struggles with the dynamic components in the scene and produces blurry reconstructions. NeRF-W obtains sharper reconstructions of the static regions, but does not handle well the dynamic regions. NeuralDiff produces sharper results, especially for the moving objects. See for example the plates. Finally NeuralDiff+C+A captures more details, such as the arms in the first scene, or the spoon in the second.

Next, Figure 5 illustrate success and failure cases for the segmentation task. In the best cases for our method (top), NeRF-W produces noisy predictions which barely capture the plate, the pasta colander, and the actor’s body. NeuralDiff improves over these results and captures all detachable objects and more body parts, but classifies part of the floor as foreground. NeuralDiff+A successfully identifies the floor as static and better predicts the shape of the actor’s body (second and third rows). For the failure cases

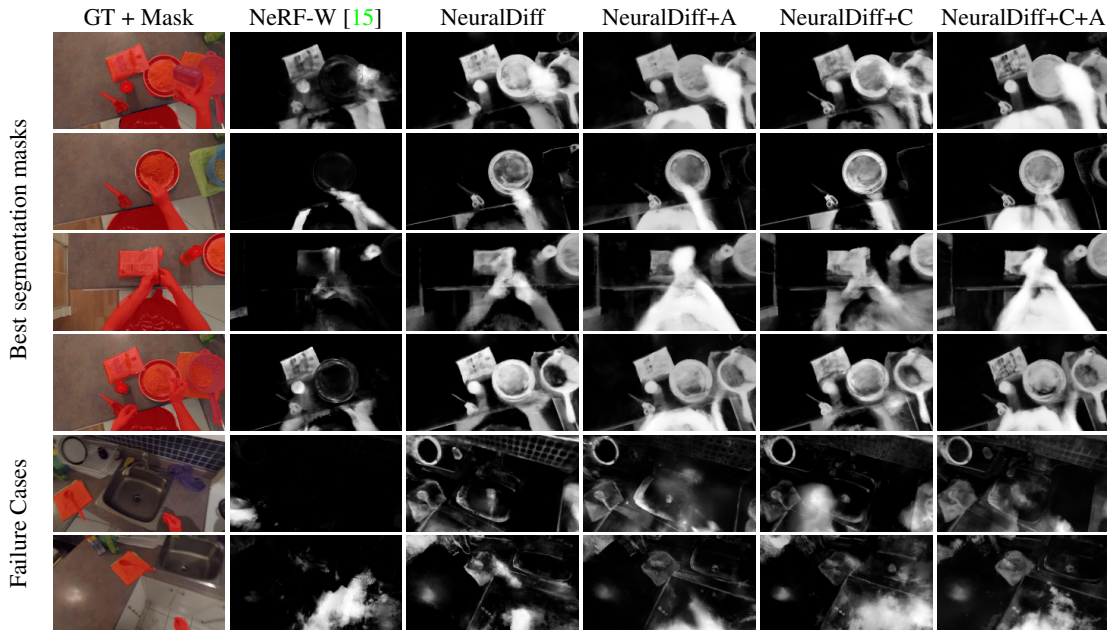


Figure 5: **Segmentation masks.** For a given sequence, segmentation results of the 5 (resp. 3) frames for which NeuralDiff+A got the highest (resp. lowest) AP, for NeRF-W [15], NeuralDiff, NeuralDiff+C, and NeuralDiff+C+A.

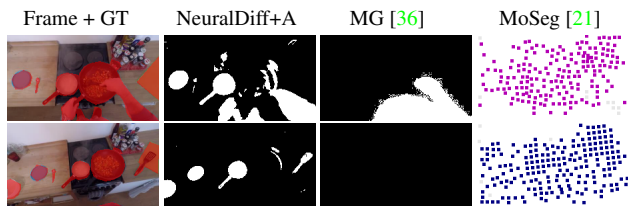


Figure 6: **Qualitative results.** Our method captures the body parts, even when currently not moving, and foreground objects, even if temporarily idle. MotionGroup (MG) [36] only captures the actor when moving. The traditional MoSeg [21] is not designed for this challenging egocentric scenario and fails entirely.

(bottom), while the segmentation mAP scores are generally low, NeuralDiff outperforms NeRF-W even more significantly. Specifically, NeRF-W fails to identify foreground objects. NeuralDiff improves over that but incorrectly classifies some of the background as foreground (see *e.g.* rows 6 to 8 where the drying rack, static in this scene, is predicted as dynamic). Such errors are reduced by NeuralDiff+C, NeuralDiff+A and NeuralDiff+C+A.

Comparison to motion segmentation. Finally, we compare to MoSeg [21] as a traditional motion segmentation method and MotionGroup [36] as a recent one. We report intersection over union (IoU) scores, as these methods produce binary masks. Quantitative and qualitative results are shown in Table 2 and Figure 6. We see that none of these approaches is suited for the task that our method is designed for, as these techniques are unable to capture objects that move infrequently and suffer from the heavy occlusions of

Method	IoU
NeuralDiff + A (ours)	43.1
MotionGroup [36]*	19.8
MotionGroup [36]**	28.7

Table 2: Comparison with MotionGroup [36]. Average of per-frame IoU scores for all 10 sequences. * indicates that we use their pretrained model, and ** that we retrained their model on one of our scenes.

the actor. The method of [21] is particularly ill-suited for the task, so we did not evaluate it quantitatively.

6. Conclusions

We have shown that recent neural rendering techniques can successfully be applied to the *analysis* of egocentric videos. We have done so by introducing NeuralDiff, a triple-stream neural renderer separating the background, foreground and actor via appropriate inductive biases. Together with other improvements such as smooth dynamics and more principled color mixing, NeuralDiff significantly outperforms baselines such as NeRF-W for our task.

We have used NeuralDiff to identify objects that move (*i.e.* are detached) in long and complex video sequences from EPIC-KITCHENS, even if the objects are small and move only little and sporadically. We believe that these results will inspire further research on the use of neural rendering for unsupervised image understanding.

Acknowledgments. Andrea Vedaldi was partially sponsored by ERC 638009-IDIU.

References

- [1] Pia Bideau and Erik Learned-Miller. It's moving! a probabilistic model for causal motion segmentation in moving camera videos. In *Proc. ECCV*, 2016. 3
- [2] Thierry Bouwmans. Traditional and recent approaches in background modeling for foreground detection: An overview. *Computer science review*, 11:31–66, 2014. 2
- [3] Thomas Brox and Jitendra Malik. Object segmentation by long term analysis of point trajectories. In *Proc. ECCV*, 2010. 3
- [4] Jianchuan Chen, Ying Zhang, Di Kang, Xuefei Zhe, Linchao Bao, and Huchuan Lu. Animatable neural radiance fields from monocular rgb video. *arXiv.cs*, abs/2106.13629, 2021. 3
- [5] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The epic-kitchens dataset. In *Proc. ECCV*, 2018. 2
- [6] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Jian Ma, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Rescaling egocentric vision. *CoRR*, abs/2006.13256, 2020. 5
- [7] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. *arXiv.cs*, abs/2105.06468, 2021. 3
- [8] J. J. Gibson. *The Ecological Approach to Visual Perception*. Houghton Mifflin, 1986. 2
- [9] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, Miguel Martin, Tushar Nagarajan, Ilija Radosavovic, Santhosh Kumar Ramakrishnan, Fiona Ryan, Jayant Sharma, Michael Wray, Mengmeng Xu, Eric Zhongcong Xu, Chen Zhao, Siddhant Bansal, Dhruv Batra, Vincent Cartillier, Sean Crane, Tien Do, Morrie Doulaty, Akshay Erapalli, Christoph Feichtenhofer, Adriano Fragomeni, Qichen Fu, Christian Fuegen, Abraham Gebreselasie, Cristina Gonzalez, James Hillis, Xuhua Huang, Yifei Huang, Wenqi Jia, Weslie Khoo, Jachym Kolar, Satwik Kottur, Anurag Kumar, Federico Landini, Chao Li, Yanghao Li, Zhenqiang Li, Karttikeya Mangalam, Raghava Modhugu, Jonathan Munro, Tullie Murrell, Takumi Nishiyasu, Will Price, Paola Ruiz Puentes, Merey Ramazanov, Leda Sari, Kiran Somasundaram, Audrey Southerland, Yusuke Sugano, Ruijie Tao, Minh Vo, Yuchen Wang, Xindi Wu, Takuma Yagi, Yunyi Zhu, Pablo Arbelaez, David Crandall, Dima Damen, Giovanni Maria Farinella, Bernard Ghanem, Vamsi Krishna Ithapu, C. V. Jawahar, Hanbyul Joo, Kris Kitani, Haizhou Li, Richard Newcombe, Aude Oliva, Hyun Soo Park, James M. Rehg, Yoichi Sato, Jianbo Shi, Mike Zheng Shou, Antonio Torralba, Lorenzo Torresani, Mingfei Yan, and Jitendra Malik. Ego4d: Around the world in 3,000 hours of egocentric video, 2021. 2
- [10] Suyog Dutt Jain, Bo Xiong, and Kristen Grauman. Fusionseg: Learning to combine motion and appearance for fully automatic segmentation of generic objects in videos. In *Proc. CVPR*, 2017. 3
- [11] Rudrika Kalsotra and Sakshi Arora. A comprehensive survey of video datasets for background subtraction. *IEEE Access*, 7:59143–59171, 2019. 2
- [12] Alex Kendall and Yarin Gal. What uncertainties do we need in Bayesian deep learning for computer vision? *Proc. NeurIPS*, 2017. 5
- [13] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Proc. CVPR*, 2021. 3, 7
- [14] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv.cs*, abs/1608.03983, 2017. 6
- [15] Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. In *Proc. CVPR*, 2021. 2, 3, 4, 5, 6, 7, 8
- [16] Jana Mattheus, Hans Grobler, and Adnan M Abu-Mahfouz. A review of motion segmentation: Approaches and major challenges. In *IMITEC*, 2020. 3
- [17] Kevin Matzen and Noah Snavely. Scene chronology. In *Proc. ECCV*, 2014. 3
- [18] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tanck, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *Proc. ECCV*, 2020. 2, 3, 4, 6, 7
- [19] Anton Mitrokhin, Zhiyuan Hua, Cornelia Fermuller, and Yiannis Aloimonos. Learning visual motion segmentation using event surfaces. In *Proc. CVPR*, 2020. 3
- [20] David Novotný, Diane Larlus, and Andrea Vedaldi. Learning 3D object categories by looking around them. In *Proc. ICCV*, 2017. 5
- [21] P. Ochs, J. Malik, and T. Brox. Segmentation of moving objects by long term video analysis. *PAMI*, 36(6):1187 – 1200, Jun 2014. Preprint. 3, 8
- [22] Anestis Papazoglou and Vittorio Ferrari. Fast object segmentation in unconstrained video. In *Proc. ICCV*, 2013. 3
- [23] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. *arXiv.cs*, abs/2011.12948, 2020. 3, 7
- [24] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Animatable neural radiance fields for human body modeling. *arXiv preprint arXiv:2105.02872*, 2021. 3
- [25] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proc. CVPR*, 2021. 3
- [26] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proc. CVPR*, 2016. 4, 6
- [27] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *Proc. ECCV*, 2016. 4

- [28] Karl Stelzner, Kristian Kersting, and Adam R. Kosiorek. Decomposing 3d scenes into objects via unsupervised volume segmentation. *arXiv.cs*, abs/2104.01148, 2021. 3
- [29] Narayanan Sundaram, Thomas Brox, and Kurt Keutzer. Dense point trajectories by gpu-accelerated large displacement optical flow. In *Proc. ECCV*, 2010. 3
- [30] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290, 2000. 3
- [31] Pavel Tokmakov, Karteek Alahari, and Cordelia Schmid. Learning motion patterns in videos. In *Proc. CVPR*, 2017. 3
- [32] Edgar Tretschik, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. *arXiv.cs*, abs/2012.12247, 2021. 3
- [33] Wenguan Wang, Xiankai Lu, Jianbing Shen, David J Crandall, and Ling Shao. Zero-shot video object segmentation via attentive graph neural networks. In *Proc. ICCV*, 2019. 3
- [34] Christopher Xie, Keunhong Park, Ricardo Martin-Brualla, and Matthew Brown. Fig-nerf: Figure-ground neural radiance fields for 3d object category modelling. *arXiv.cs*, abs/2104.08418, 2021. 3
- [35] Christopher Xie, Yu Xiang, Zaid Harchaoui, and Dieter Fox. Object discovery in videos as foreground motion clustering. In *Proc. CVPR*, 2019. 3
- [36] Charig Yang, Hala Lamdouar, Erika Lu, Andrew Zisserman, and Weidi Xie. Self-supervised video object segmentation by motion grouping. In *Proc. ICCV*, 2021. 3, 8
- [37] Gengshan Yang and Deva Ramanan. Learning to segment rigid motions from two frames. In *Proc. CVPR*, 2021. 3
- [38] Wentao Yuan, Zhaoyang Lv, Tanner Schmidt, and Steven Lovegrove. Star: Self-supervised tracking and reconstruction of rigid objects in motion with neural rendering. In *Proc. CVPR*, 2021. 3
- [39] Luca Zappella, Xavier Lladó, and Joaquim Salvi. Motion segmentation: A review. In *Proc. CCIA*, 2008. 3
- [40] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. NeRF++: Analyzing and improving neural radiance fields. *arXiv.cs*, abs/2010.07492, 2020. 4