

# Neural Feature Fusion Fields: 3D Distillation of Self-Supervised 2D Image Representations

Vadim Tschernezki<sup>1,2</sup> Iro Laina<sup>1</sup> Diane Larlus<sup>2</sup> Andrea Vedaldi<sup>1</sup>

<sup>1</sup> Visual Geometry Group, University of Oxford  
{vadim,iro,vedaldi}@robots.ox.ac.uk

<sup>2</sup> NAVER LABS Europe  
diane.larlus@naverlabs.com

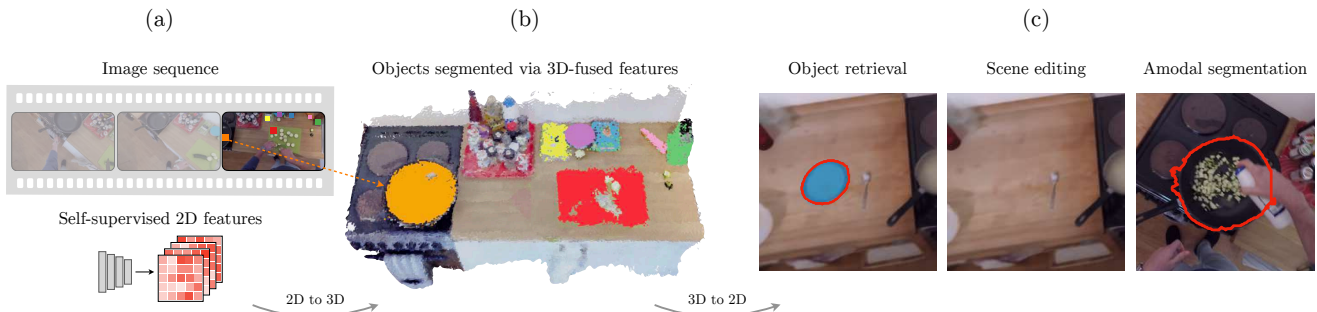


Figure 1: **Neural Feature Fusion Fields:** (a) Given a collection of 2D images and a dense self-supervised feature extractor such as DINO, our method distills the features into a 3D representation via neural rendering. This allows to operate within the learned scene representation through 2D inputs. (b) For example, prompted with the features of a 2D region (any of the colored patches in (a)), our method segments the corresponding object in 3D as shown in the point cloud. (c) We can also render the scene representation and solve image-level tasks such as object retrieval, scene editing, or amodal segmentation.

## Abstract

We present *Neural Feature Fusion Fields (N3F)*, a method that improves dense 2D image feature extractors when the latter are applied to the analysis of multiple images reconstructible as a 3D scene. Given an image feature extractor, for example pre-trained using self-supervision, N3F uses it as a teacher to learn a student network defined in 3D space. The 3D student network is similar to a neural radiance field that distills said features and can be trained with the usual differentiable rendering machinery. As a consequence, N3F is readily applicable to most neural rendering formulations, including vanilla NeRF and its extensions to complex dynamic scenes. We show that our method not only enables semantic understanding in the context of scene-specific neural fields without the use of manual labels, but also consistently improves over the self-supervised 2D baselines. This is demonstrated by considering various tasks, such as 2D object retrieval, 3D segmentation, and scene editing, in diverse sequences, including long ego-centric videos in the EPIC-KITCHENS benchmark. Project page: <https://www.robots.ox.ac.uk/~vadim/n3f/>

## 1. Introduction

With the advent of machine learning two decades ago, computer vision has shifted its focus from 3D reconstruction to interpreting images mostly as 2D patterns. Recently, however, methods such as NeRF [36] have shown that even 3D reconstruction can be cast effectively as a learning problem. However, such methods are often optimized per scene, resulting in low-level representations of appearance and geometry that do not capture high-level semantics. It is thus compelling to rethink how 3D reconstruction can be integrated with semantic analysis at image level to obtain more holistic scene representations.

In this paper, we consider a simple, general and effective approach for achieving such an integration, which we call *Neural Feature Fusion Fields (N3F)*; see Figure 1). The key idea is to map semantic features, initially learned and computed in 2D image space, to equivalent features defined in 3D space. Forward mapping from 3D space to 2D images uses the same neural rendering equations as for view synthesis in prior work. From this, backpropagation can move features from the 2D images back to the 3D model.

Recent methods such as Semantic NeRF [77] and Panoptic NeRF [12, 25] have described a similar process for semantic and instance segmentation via label transfer. Our intuition is that fusion does not need to be limited to image labels, but can be extended to any image features. N3F follows a student-teacher setup, where features of a 2D teacher network are distilled into a 3D student network. We show that distilling the features alongside 3D geometry via neural rendering significantly boosts their consistency, viewpoint independence, and occlusion awareness. As a result, the student “surpasses the teacher” in understanding a particular scene and improves tasks such as object retrieval and segmentation from visual queries.

As a particularly compelling case of this idea, we consider starting from a *self-supervised* feature extractor. Recent work [1, 4, 34, 60] has shown that self-supervised features can be used to identify object categories, parts, and their correspondences in an “open-world” setting, *i.e.* without committing to a specific set of labels and without collecting annotations for them. This is of particular relevance for emerging applications such as egocentric video understanding, where image understanding must work in user-specific and constantly evolving scenarios.

Specifically, we consider two scenarios of increasing difficulty. First, we validate our contribution on simple *static* scenes with only one or a few objects of interest, and combine N3F with the vanilla NeRF model [36]. Second, we consider the more challenging scenario of egocentric videos which include static but also *dynamic* components. We adopt the same setting as NeuralDiff [59] and consider videos from EPIC-KITCHENS [10] which contain long sequences of actors cooking in first-person view.

Given these diverse sets of videos, we use object retrieval (*e.g.*, one-shot recognition) as a proxy to evaluate the quality of the fused image features. Considering an object instance in a single frame, we use it to pool the features from N3F, and then retrieve other occurrences of the same object in the rest of the video, despite severe viewpoint changes or occlusions. We show that, while 2D features already perform well for this task, N3F systematically boosts performance by a large margin. This observation is consistent for several self-supervised and supervised features. We illustrate other benefits of such an integrated model by also showing results for the tasks of 3D object segmentation, amodal segmentation, and scene editing.

## 2. Related Work

We summarize relevant background work in feature extraction, reconstruction and neural rendering.

**Self-supervised visual features.** While N3F can work on top of any 2D dense image features, including recent ones based on Vision Transformers (ViT) [11] and vari-

ants [2, 5, 13, 18, 30, 56, 57, 66, 74, 79], of particular interest are self-supervised versions such as [4, 8, 16, 27] as they are more generically applicable and can benefit more from the consistency induced by N3F. Caron *et al.* [4] observed that their method (DINO), trained with self-distillation, learns better localized representations, which can be used to segment salient objects *without any labels*. Subsequently these features have been used for unsupervised object localization [34, 52, 65], semantic segmentation [15, 34, 80], part segmentation [1, 9] and point correspondences [1].

**Neural rendering.** Using implicit representations of geometry in vision dates back at least to level-set methods [38]. Recently, authors have proposed to represent implicit functions with deep neural networks for the representation of geometry [40] and radiance fields [53], fitting the latter to 2D images via differentiable rendering. Neural Radiance Fields (NeRFs) [36] have popularized such ideas by applying them in a powerful manner to a comparatively simple setting: novel view synthesis from a collection of images of a single static scene. They combine radiance fields with internal learning [61] and various architectural improvements such as positional encoding to obtain excellent scene reconstructions. For a comprehensive overview of recent trends in this field see [54, 69].

Among countless extensions of NeRF, of particular interest for our applications are versions tackling dynamic scenes. For example, NSFF [29] models scenes through time-dependent flow fields, which enable novel view synthesis in space *and* time; other methods achieve a similar effect by introducing canonical models [6, 41, 44, 45] or space warping [58]. Used here, NeuralDiff [59] extends the standard NeRF reconstruction of the static part of a scene with two dynamic components, one for transient objects (foreground), and one for the actor in egocentric videos.

**Semantic and object-centric neural rendering.** Radiance fields provide low-level representations of geometry and radiance and lack a higher-level (*e.g.*, semantic or object-centric) understanding of the scene. Several works employ neural rendering to decompose multi-view or dynamic scenes into background and foreground components [39, 47, 51, 59, 73, 75], while others focus on modeling the scenes as compositions of objects [14, 26, 37, 67, 68, 70]. Some authors propose to combine radiance fields with image-language models (*e.g.*, CLIP [46]) to achieve semantically aware synthesis [21, 64].

More related to our work, however, are methods that extend radiance fields to also *predict* semantics [24, 63, 77, 78]. For example, Semantic-NeRF [77] has done so by using differentiable rendering to achieve multi-view semantic fusion of 2D labels akin to [19, 33, 62]. NeSF [63] focuses instead on inferring semantics jointly across various scenes, using density fields as input to a 3D segmentation model. However, it is only demonstrated on synthetic scenes with a

limited number of categories and shapes. Panoptic (*i.e.*, semantic and instance) labels, have also been considered: [12] uses NeRF as a means to integrate coarse 3D and noisy 2D labels and render refined 2D panoptic maps, while [25] proposes an object-aware approach that can handle dynamic scenes, where each 3D instance is modeled by a separate MLP. All of these methods use semantic labels to train their models and in particular the latter two require 3D labels. Instead, our approach builds on self-supervised features and can yield a 3D-consistent semantic segmentation of static and dynamic scenes without *any* labels.

The most related work is the concurrent paper by Kobayashi *et al.* [23] who propose to fuse features in the same manner as we do; they mainly differ in the example applications, including the use of multiple modalities, such as text, image patches and point-and-click seeds, to generate queries for segmentation and, in particular, scene editing.

**Feature distillation.** The motivation behind distillation originates from the task of compressing, or “distilling”, the knowledge of large, complex model ensembles into smaller models, while preserving their performance [3]. Hinton *et al.* [20] have shown that the performance of a distilled (student) model can even improve over the performance of the original model or model ensemble (teacher) when following the teacher-student paradigm. Many methods have since then proposed to use this paradigm on features, tackling the task of feature distillation [17, 42, 43, 48, 55, 72, 76]. While N3F also makes use of the teacher-student paradigm, it differs in that the output of a 2D teacher network is distilled into a student network that implements a 3D feature field, resulting in different domains of the student (2D images) and teacher (3D points).

### 3. Method

We first describe Neural Feature Fusion Fields (N3F) for generic (Section 3.1) and advanced (Section 3.2) neural rendering models, and then introduce a number of applications (Section 3.3) which we use to demonstrate its benefits.

#### 3.1. Neural Feature Fusion Fields

Let  $I \in \mathbb{R}^{3 \times H \times W}$  be an input image defined on the lattice  $\Omega = \{1, \dots, H\} \times \{1, \dots, W\}$  and let  $\Phi$  be a feature extractor, *i.e.* a function mapping the image  $I$  to a representation  $\Phi(I)$ . We assume the representation is in the form of a vector field  $\mathbb{R}^{C \times H \times W}$ , which is in itself an image with  $C$  feature channels. Example features include dense SIFT features [32], convolutional networks [7] and visual transformers [11]. Furthermore, these features can be handcrafted, supervised or unsupervised.

Now suppose that the image is part of a collection  $\{I_t\}_{1 \leq t \leq T}$  and that the camera parameters are given, so that the projection function  $\pi_t$  from world coordinates  $X \in \mathbb{R}^3$

to image coordinates  $u = \pi_t(X) \in \mathbb{R}^2$  is known. A *neural radiance field* is a pair of functions  $(\sigma, c)$  mapping 3D points  $X \in \mathbb{R}^3$  to occupancy values  $\sigma(X) \in \mathbb{R}_+$  and to colors  $c(X) \in \mathbb{R}^3$  respectively. In practice, color also depends on the viewing direction  $d \in \mathbb{R}^3$ , but our notations omit this dependency for brevity. The neural rendering equation reconstructs the color  $I_{tu}$  of pixel  $u$  as

$$\hat{I}_{tu} = \int_0^\infty c(X_{tu}(r))\sigma(X_{tu}(r))e^{-\int_0^r \sigma(X_{tu}(q))dq} dr \quad (1)$$

where  $\{X_{tu}(r)\}_{r>0}$  are points along the ray from the camera center through pixel  $u$  in image  $I_t$ .

The idea of neural rendering is to *learn* the functions  $\sigma$  and  $c$  given only the images  $I_t$  and the camera poses  $\pi_t$  as input. This is done by minimizing the image reconstruction loss  $\sum_t \|\hat{I}_t - I_t\|^2$  for all images in the sequence with respect to the parameters of the models  $\sigma$  and  $c$ .

In N3F, we propose to generalize this model by also reconstructing *feature* images  $\Phi(I_t)$  instead of just color images  $I_t$ . For this, we also minimize the loss:

$$\sum_t \|\hat{\Phi}_t - \Phi(I_t)\|^2. \quad (2)$$

In order to do so, we modify Equation (1) to generate, in addition to a color image  $\hat{I}_t$ , a feature image  $\hat{\Phi}_t \in \mathbb{R}^{C \times H \times W}$ . This is obtained by modifying the range of the function  $c = (c_{\text{rgb}}, c_\Phi)$  to be  $\mathbb{R}^{3+C}$ , where  $C$  is the number of feature channels. We call the pair  $(\sigma, c_\Phi)$  a *neural feature field* to distinguish it from the neural radiance field  $(\sigma, c_{\text{rgb}})$  typically used for view synthesis.

In the context of neural networks, this approach can be understood as a 2D-teacher-3D-student model. The teacher is the feature network  $\Phi$ , which is defined in image space, and the student is the network implementing the function  $c_\Phi$ , defined in 3D world space. This is illustrated in Figure 2. The final training loss for the student network is simply the sum of the image reconstruction loss and the feature reconstruction loss weighed by a factor  $\lambda$ :

$$\sum_t \|\hat{I}_t - I_t\|^2 + \lambda \|\hat{\Phi}_t - \Phi(I_t)\|^2. \quad (3)$$

The key benefits of this approach are twofold. First, knowledge from the teacher network is *distilled* into the student network in a manner that correctly reflects the 3D geometry of the scene, which has a smoothing effect and helps to regularize feature prediction. As we show later, this results in higher quality features that are more consistent across viewpoints. Second, distilling features of general-purpose feature extractors pre-trained on large external datasets—with or without supervision—brings open-world knowledge into the 3D representation, which is otherwise scene-specific and lacks semantic understanding.




Figure 2: **Overview of our approach.** N3F follows a student-teacher setting where features computed from individual images are distilled into a 3D student network. The student network extends NeRF-like models such that a ray from a selected view is mapped to a color value  $\hat{I}_{tu}$  and a corresponding feature vector  $\hat{\Phi}_{tu}$  through volumetric rendering. The teacher network, which is learned with self-supervision (SSL), predicts the 2D image features  $\Phi(I_t)_u$  to be distilled. The student is trained to optimize both image and feature reconstruction objectives, whereas the teacher is not trained further (stop gradient or ‘sg’). While the student network solely learns from 2D features, the resulting representation can operate either in 2D or in 3D.

### 3.2. Distillation with advanced NeRF architectures

In N3F, we are free to implement the neural field  $(\sigma, c)$  in any of the many variants that have been proposed in the literature. In this paper, we showcase the approach on two different scenarios: simple static scenes, as typically handled by NeRF (presented before), and egocentric videos that are significantly more complex, for which standard neural rendering models are insufficient. Again, we stress that many other variants would also apply.

The challenge of egocentric videos is that they contain a mixture of static background objects, foreground objects that are manipulated by the actor, and body parts of the actor themselves (*e.g.*, hands). We handle them by adopting NeuralDiff [59], a NeRF-like architecture that automatically decomposes a dynamic scene into these three components, combining a static field representing the background, a dynamic field (with a dependency on time  $t$  in addition to space  $X$ ) representing the foreground, and another dynamic field anchored at the camera representing the actor. We can adapt NeuralDiff to support N3F simply by considering a feature prediction head in addition to the color and density heads for *each* of the three components (MLPs).

### 3.3. Applications of N3F

In addition to employing N3F on top of different neural rendering models, we demonstrate its versatility by considering various downstream applications: 2D object retrieval, 3D object segmentation, 3D scene editing, and amodal segmentation. For all these tasks and for ease of evaluation, we assume that a 2D region is provided as a query for a single given frame  $I_t$ . As a particular use case for providing queries, one can think of the user introducing an object of interest which can be then localized, *e.g.*, across a

video. However, we note that providing such annotations is not strictly necessary and one could also consider direct clustering of the distilled features to obtain segmentations of objects without manual input, as shown in [1, 34].

**2D object retrieval.** Given a collection of images, and given any object from a single reference frame, we would like to find all the occurrences of the same object in the rest of the collection, despite significant viewpoint changes, occlusions, and various dynamic effects. In particular, given a region  $R_t \subset \Omega$  of the image containing a fully or partially visible object at time  $t$ , or even just a patch, we pool a feature descriptor as the mean of the region’s features:

$$\Phi(I_t)_{R_t}^{\text{avg}} = \frac{1}{|R_t|} \sum_{u \in R_t} \Phi(I_t)_u. \quad (4)$$

To localize the object in another image  $I_{t'}, t' \neq t$ , we return as matching region  $\tilde{R}_{t'}$  the set of pixels whose features are sufficiently close to the mean descriptor according to a threshold  $\tau$ :

$$\tilde{R}_{t'} = \{u \in \Omega : \|\eta(\Phi(I_{t'})_u) - \eta(\Phi(I_t)_{R_t}^{\text{avg}})\| \leq \tau\},$$

where  $\eta(a) = a/\|a\|$  normalizes the input vector  $a$ .

Performance on this task directly depends on the quality of the matched features. Despite the 2D character of this task, the above equations are directly applicable to N3F, by simply replacing  $\Phi(I_t)$  with the distilled feature map  $\hat{\Phi}_t$  obtained *after rendering* the 3D features back to the  $t$ -th view, as explained above. Hereby, we denote the *distilled* mean feature vector corresponding to region  $R_t$  as  $\hat{\Phi}_{tR_t}^{\text{avg}}$ .

**3D object segmentation.** Since N3F predicts a 3D field of features, these features can be used directly, *i.e.* prior to rendering, to segment a queried object along with its geometry in the 3D space, rather than retrieving it in a series of

2D images. Formally, given features  $\hat{\Phi}_{tR_t}^{\text{avg}}$  extracted from a single 2D annotation  $R_t$  of the object in image  $I_t$ , we retrieve the 3D region  $\{X \in \mathbb{R}^3 : \|c_\Phi(X) - \hat{\Phi}_{tR_t}^{\text{avg}}\| \leq \tau_\Phi \wedge \sigma(X) \geq \tau_\sigma\}$ , where  $\tau_\Phi$  and  $\tau_\sigma$  denote thresholds for the features of interest and densities respectively. We note that this application is seamlessly enabled by N3F, while not possible to address with the 2D teacher network.

**Scene editing.** Instead of extracting a 3D object, we can also suppress it, *i.e.* remove it from the scene. To achieve this, we can simply set the occupancy  $\sigma(X)$  to zero for all 3D points belonging to an object, *i.e.* all points  $X$  such that  $\|c_\Phi(X) - \hat{\Phi}_{tR_t}^{\text{avg}}\| \leq \tau_\Phi$ . Once again, in our experiments, the object to be removed is identified using a query region in one of the views (object patch or region).

**Amodal segmentation.** We can adjust the querying and retrieval process of our method to handle occlusions in two different ways. The first corresponds to the 2D object retrieval task; in this case, due to the rendering process, features (just like colors) are “blocked” from reaching the camera if they are occluded, for example by the actor in egocentric videos or by other objects. However, our approach makes it possible to also *see through occluders*, by disabling the occupancies  $\sigma$  for regions of the 3D space that contain features dissimilar to the query descriptor (Equation (4)). In practice, this amounts to rendering the 3D features after obtaining a segmentation of the object in 3D, as described above. In this manner, it is possible to obtain a mask of the full extent of the object, as if occluders are removed, which is often referred to as amodal segmentation [28].

## 4. Experiments

In this section, we evaluate the features produced by N3F for the tasks introduced in Section 3.3 for static and dynamic scenes. Section 4.1 gives the experimental details and Section 4.2 reports the results for the different tasks. Section 4.3 discusses limitations of our approach.

### 4.1. Experimental setup

**Datasets.** We consider scenes from the LLFF dataset [35] and a subset of the EPIC-KITCHENS dataset [10]. The former contains images of static scenes while the latter contains egocentric videos of people cooking in different kitchens, and interacting with a large number of different objects, such as food or kitchen utensils. For the former, we implement N3F on top of the vanilla NeRF [36] architecture, and for the latter on top of the more complex NeuralDiff [59] architecture, as described next.

**NeRF.** For the experiments on the LLFF dataset [35], we use the NeRF PyTorch implementation [71] with the default hyperparameters. We adapt the architecture with an additional feature prediction head, consisting of a single linear layer with  $\tanh$  as activation function. We use the pre-

trained models supplied with the implementation and continue training for 5k iterations, freezing all but the feature prediction head for the first 1k iterations. The weight for the feature distillation loss  $\lambda = 0.001$ . The features are rendered similarly to pixel colors as described in Section 3.1.

**NeuralDiff.** We build on the model proposed in [59] for the experiments on EPIC-KITCHENS. We extend the three-stream architecture with feature prediction heads (followed by  $\tanh$ ), one for each component (static, dynamic, actor). The model is first trained for RGB reconstruction (10 epochs with 20k iterations each and a batch size of 1024). Training for a single scene takes approximately 24 hours on an NVIDIA Tesla P40. We then finetune the model to distill the pre-computed teacher features for 20 epochs, 500 iterations each, with the same batch size (approx. 2 hours) and again freezing all but the feature prediction heads for the first 1k steps (training from scratch yields similar results, but is slower). We down-sample images to  $480 \times 270$  pixels and upscale the 2D features with nearest neighbour interpolation. We set  $\lambda = 1.0$  for the feature distillation loss. The models are trained using Adam optimization [22], an initial learning rate of  $5 \times 10^{-4}$  and a cosine annealing schedule [31].

**2D teacher features.** We consider four transformer-based feature extractors: DINO [4] with patch size 8 and 16, MoCo-v3 [8] and DeiT [56]. DINO and MoCo-v3 are self-supervised whereas DeiT is trained with supervision (image labels). Features on all scenes are pre-computed using the publicly available weights (pre-trained on ImageNet [49]), which are not further updated during the distillation process. The features are then  $L^2$ -normalized and reduced with PCA to 64 dimensions before distillation.

**Evaluation metric.** Each sequence used for quantitative evaluation in the next section has  $K$  objects annotated in  $N$  different frames, corresponding to ground truth regions  $\{R_{kn}\}_{1 \leq k \leq K, 1 \leq n \leq N}$ . These annotations are used for evaluation only, and never considered during training. Details on the annotation process can be found in the supplementary material. We consider the task of 2D object retrieval for quantitative evaluation and divide the annotated frames into two non-overlapping sets, a query set  $Q$  and a gallery set  $G$  ( $Q \cup G = \{1, \dots, N\}$ ). Each region  $R_{kq}$  ( $q \in Q$ ) is in turn used as a query, searching for the corresponding object in each annotated frame from the gallery set. In order to avoid fixing a threshold  $\tau$  as in Equation (4), for each target frame  $I_g$  ( $g \in G$ ), pixels  $u$  are sorted by increasing distance to the mean feature  $\Phi(I_g)_{R_{kq}}^{\text{avg}}$  ( $\hat{\Phi}_{R_{kq}}^{\text{avg}}$  for the N3F-distilled features) and labeled as positive if they belong to the ground truth region  $R_{kq}$  and as negative otherwise. The sorted labels are used to compute the Average Precision (AP)  $AP_{kqg}$ . The AP values are then averaged across videos, objects and queries to obtain a mean AP (mAP).

Method	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10	Average (abs gain)
DINO [4] [ViT-B/8]	75.75	57.25	56.46	63.11	70.56	65.81	52.28	78.28	58.19	65.79	64.35
N3F (DINO)	83.64	67.19	69.21	80.23	78.17	77.57	64.32	83.85	76.24	82.17	<b>76.26 (+11.91)</b>
DINO [4] [ViT-B/16]	77.37	53.21	48.91	57.44	68.32	60.39	40.39	74.07	53.22	62.19	60.15
N3F (DINO)	88.61	66.99	69.90	87.02	78.66	78.97	70.57	85.17	77.59	84.93	<b>78.84 (+18.69)</b>
MoCo-v3 [8] [ViT-B/16]	70.73	54.02	48.02	52.89	67.18	57.34	43.54	73.45	47.85	60.12	57.51
N3F (MoCo-v3)	86.67	68.95	68.53	82.93	75.74	78.00	65.63	83.58	68.26	83.21	<b>76.15 (+18.64)</b>
DeiT [56] [ViT-B/16]	55.27	40.78	38.02	42.76	54.01	51.70	37.72	61.53	40.88	52.48	47.51
N3F (DeiT)	86.02	62.47	66.69	81.22	72.93	77.88	61.63	83.73	69.59	83.12	<b>74.53 (+26.82)</b>

Table 1: **2D object retrieval.** We compare the features learned by our approach (NeuralDiff-N3F) with the 2D teacher features on the task of retrieving 2D objects for 10 scenes of the EPIC-KITCHENS dataset. We consider features from three self-supervised models, two flavors of DINO [4] and MoCo-v3 [8], and a supervised one, DeiT [56]. We report per scene mean average precision (mAP) results and the overall Average.

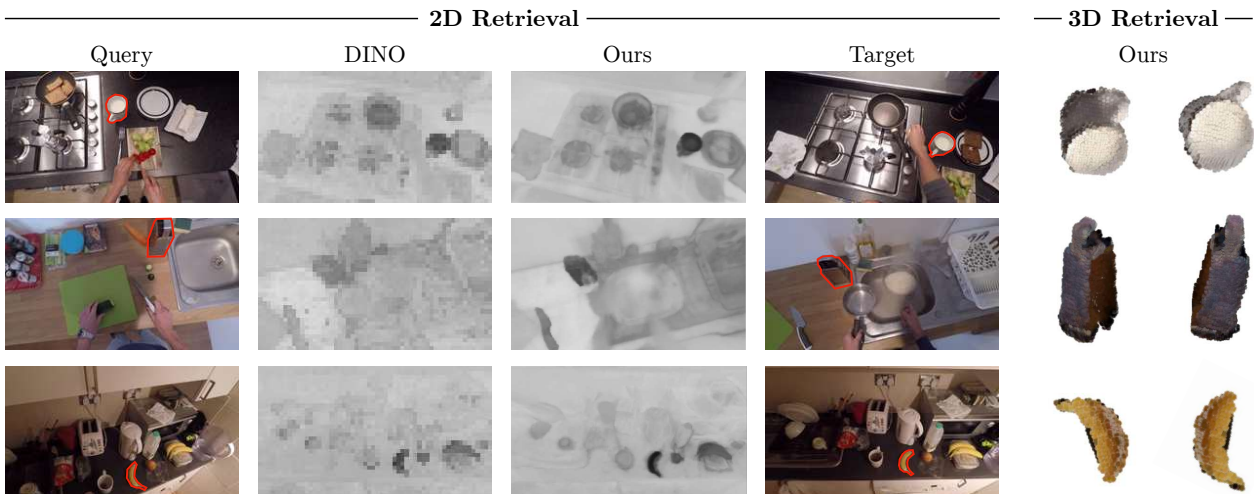


Figure 3: **Retrieving (segmenting) objects in 2D and 3D.** Given a feature descriptor obtained by pooling features from a given region (Query) in a reference frame, we retrieve similar regions in another frame (Target) of a video sequence. This can be achieved with either features from a teacher network (DINO) or features learned by our model (NeuralDiff-N3F). We show that N3F features are less affected by viewpoint dependent changes such as reflectance, as can be seen for the grater, which has a non-Lambertian surface. Additionally, our model can compute the densities and colors of 3D features for a given 2D query, which allows us to extract the full 3D extent of objects (seen as point clouds on the right).

## 4.2. Results

We present our results for the different tasks mentioned in Section 3.3, namely 2D and 3D object retrieval and segmentation, amodal segmentation, and scene editing.

**2D object retrieval.** In Table 1 we present quantitative evaluation results for scenes of the EPIC-KITCHENS dataset. We report the mAP value over different queries for each scene, and the average performance over all scenes. We compare the distilled features learned by NeuralDiff-N3F to those of the corresponding 2D teacher networks. We observe that the 2D features alone perform already well on this task, with self-supervised features (DINO, MoCo-v3)

surpassing supervised ones (DeiT). This is likely due to the fact that models trained with self-supervision have better generalization properties [50]. When comparing the 2D features with the distilled features, we observe significant improvements across all feature extractors and all scenes. The smallest increase occurs when distilling the already strong DINO features, resulting in an absolute difference of 11.9 mAP. The potential for improvement is larger when distilling DeiT features and we observe a larger performance gap, with our model reaching an mAP of 74.5 vs. 47.5.

We also present qualitative results on both EPIC-KITCHENS (Figure 3) and LLFF (Figure 7), comparing the features learned by NeuralDiff-N3F and NeRF-N3F re-



Figure 4: **Scene editing.** Our approach allows us to edit a scene in 3D given 2D queries. Given a 2D segment and its corresponding feature vector, we extract a 3D region of matching features, suppress its occupancies and render the view without the object, *i.e.* removing the banana (first row), lid (second row), pot (third row) and package (fourth row).

spectively with those directly obtained from a 2D teacher (DINO). In Figure 3, we show objects queried in a given frame by selecting an object mask, followed by the resulting distance map in *feature* space for a different frame of the same scene. Overall, we observe that N3F increases the clarity and correctness of the maps, resulting in sharper boundaries and higher confidence for the target objects. For example, in Figure 3 (second row), DINO struggles to recognize the grater in the target frame, possibly due to metallic reflections present in the query and a strong change in appearance. We observe similar results for NeRF-N3F in Figure 7, where our approach retrieves the whole object from a small user-provided patch, extracting a more detailed as well as complete segmentation of the objects compared to vanilla DINO. In both scenarios, our approach improves over the 2D teacher by encouraging multi-view-consistency, a property then captured by the distilled 3D features.

**3D object segmentation.** Besides retrieving objects in 2D space, our approach also allows to extract the geometry (*e.g.*, as a point cloud) of a queried region, as detailed in Section 3.3. We can thus obtain segmentations of various




Figure 5: **Amodal segmentation.** We compare NeuralDiff-N3F-distilled features to 2D DINO features for the task of amodal segmentation, *i.e.* segmenting objects through occlusions. Given a query in a reference frame, N3F allows us to retrieve the whole object in a target frame despite occluders, by comparing features in 3D and suppressing the occupancies of dissimilar regions prior to rendering.

objects in 3D, without requiring any 3D labels to train our models. This is illustrated in Figure 3 (3D Retrieval). While details are limited due to the precision of the model and partiality of the observations, the recovered shape is broadly correct. We also note that this task lies outside the capabilities of the original teacher network and is only enabled by the fusion of 2D features into the 3D field.

**Scene editing.** Figure 4 shows examples of images rendered with NeuralDiff-N3F before (left) and after (right) editing. Given a 2D query region, we find its location in 3D by matching features and suppress its occupancies (setting them to zero), thus removing selected objects. Note that images are correctly ‘inpainted’ under the object because of the holistic scene knowledge implicitly contained in the radiance field. This is especially true in the case of the EPIC-KITCHENS data, and dynamic scenes in general, as objects appear at different locations for different time steps. Thus, removing objects results in valid backgrounds, because the background was observed at some point. In comparison, scene editing in NeRF-N3F (Figure 6) results in partially hallucinated background, since part of it is occluded for all viewing directions provided during training.

**Amodal segmentation.** Figure 5 shows qualitative results for the task of amodal segmentation, *i.e.* segmenting the full extent of an object, including both visible and occluded parts. For reference, the figure also shows “ground truth” segmentations for these objects, but note that these are man-

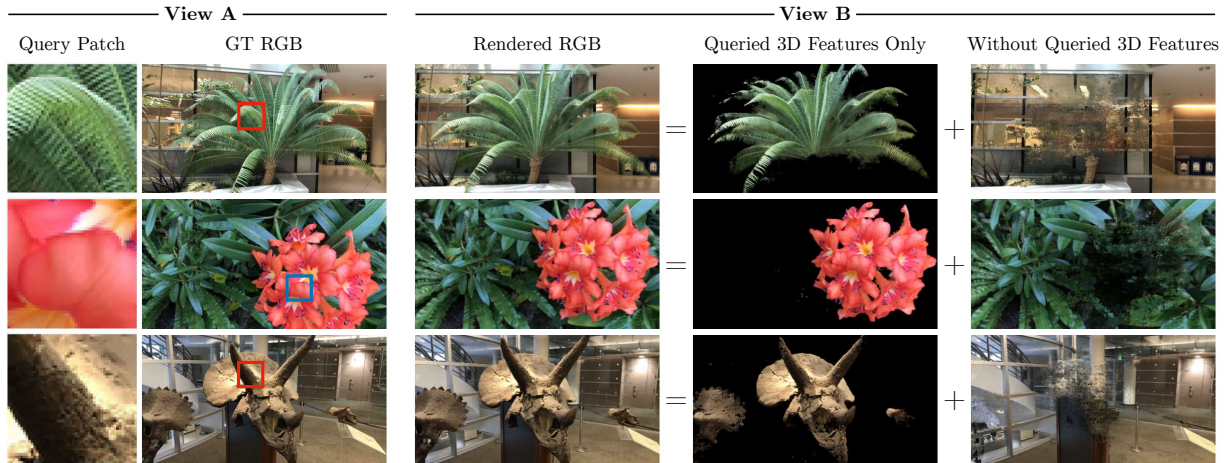


Figure 6: **Scene editing.** Given a query patch from one (unseen) view, NeRF-N3F renders an image from another (unseen) view while separating foreground from background by matching the fused 3D features to the query patch features. These results also highlight the close relationship to the concurrent work of Kobayashi *et al.* [23].

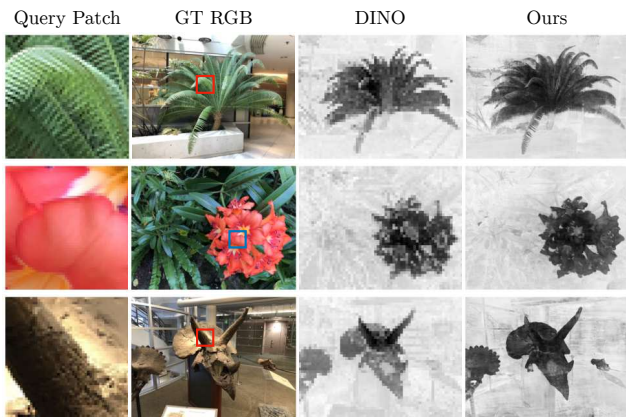


Figure 7: **2D object retrieval.** We calculate feature distance maps with DINO and with our model (NeRF-N3F) for unseen views from three LLFF scenes. Our model predicts features for these views through its 3D representation.

ually extrapolated in case of occlusions (since the object is not visible). Owing to its 3D awareness, our model is able to accurately segment, *e.g.*, the cutting board (first column), even though it is barely visible behind the actor’s arm. In comparison, the teacher network (DINO) cannot segment occluded parts, since it is limited to 2D representations.

### 4.3. Limitations and ethical considerations

N3F inherits some of the limitations of the source features. For instance, self-supervised features such as DINO tend to group semantically related objects. In the EPIC-KITCHENS dataset, we have observed this behavior for objects such as fruits and vegetables or the handles of utensils (pans, kettles, *etc.*), which are often close in feature space.

This might be undesirable in scenarios where a specific object *instance* should be tracked across a video sequence.

Another limitation is the quality of the 3D reconstruction. Reconstruction can fail catastrophically in some videos. In general, details of small or thin objects can be difficult to reconstruct, making it impossible to segment some 3D objects even if they are separated correctly by the 2D features. An example is the cutting board in Figure 3 because of its thinness and proximity to the underlying table.

Besides general caveats on the reliability of unsupervised machine learning, there do not appear to be significant ethical concerns specific to this project. EPIC-KITCHENS contains personal data (hands), but was collected with consent, and it is used in a manner compatible with their terms.

## 5. Conclusions

We have presented N3F, an approach to boost the 3D consistency of 2D image features within sets of images that can be reconstructed in 3D via neural rendering. We have shown that N3F works with various neural rendering models and scenarios, including static objects and harder egocentric videos of dynamic scenes. Our experiments illustrate the benefit of our approach for the tasks of object retrieval, segmentation and editing. Future work includes integrating N3F in the self-supervised process that learns the 2D features in the first place (*e.g.*, DINO) and fusing multiple videos to establish cross-instance correspondences (*e.g.*, by matching similar utensils in different kitchens).

**Acknowledgments.** We are grateful for support by NAVER LABS, ERC 2020-CoG-101001212 UNION, and EPSRC VisualAI EP/T028572/1. We thank the anonymous reviewers for their feedback that helped to improve our paper.



## References

- [1] Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel. Deep vit features as dense visual descriptors. *arXiv preprint arXiv:2112.05814*, 2021. [2](#), [4](#)
- [2] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. BEit: BERT pre-training of image transformers. In *Proc. ICLR*, 2022. [2](#)
- [3] Cristian Buciluă, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 535–541, 2006. [3](#)
- [4] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proc. ICCV*, pages 9650–9660, 2021. [2](#), [5](#), [6](#)
- [5] Chun-Fu Richard Chen, Quanfu Fan, and Rameswar Panda. Crossvit: Cross-attention multi-scale vision transformer for image classification. In *Proc. ICCV*, pages 357–366, 2021. [2](#)
- [6] Jianchuan Chen, Ying Zhang, Di Kang, Xuefei Zhe, Linchao Bao, Xu Jia, and Huchuan Lu. Animatable neural radiance fields from monocular rgb videos. *arXiv preprint arXiv:2106.13629*, 2021. [2](#)
- [7] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *PAMI*, 40(4):834–848, 2017. [3](#)
- [8] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proc. ICCV*, pages 9640–9649, 2021. [2](#), [5](#), [6](#)
- [9] Subhabrata Choudhury, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Unsupervised part discovery from contrastive reconstruction. *Proc. NeurIPS*, 34:28104–28118, 2021. [2](#)
- [10] 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*, pages 720–736, 2018. [2](#), [5](#)
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. [2](#), [3](#)
- [12] Xiao Fu, Shangzhan Zhang, Tianrun Chen, Yichong Lu, Lanyun Zhu, Xiaowei Zhou, Andreas Geiger, and Yiyi Liao. Panoptic nerf: 3d-to-2d label transfer for panoptic urban scene segmentation. *arXiv preprint arXiv:2203.15224*, 2022. [2](#), [3](#)
- [13] Benjamin Graham, Alaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Hervé Jégou, and Matthijs Douze. Levit: a vision transformer in convnet’s clothing for faster inference. In *Proc. ICCV*, pages 12259–12269, 2021. [2](#)
- [14] Michelle Guo, Alireza Fathi, Jiajun Wu, and Thomas Funkhouser. Object-centric neural scene rendering. *arXiv preprint arXiv:2012.08503*, 2020. [2](#)
- [15] Mark Hamilton, Zhoutong Zhang, Bharath Hariharan, Noah Snavely, and William T. Freeman. Unsupervised semantic segmentation by distilling feature correspondences. In *Proc. ICLR*, 2022. [2](#)
- [16] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proc. CVPR*, pages 16000–16009, 2022. [2](#)
- [17] Byeongho Heo, Jeessoo Kim, Sangdoon Yun, Hyojin Park, Nojun Kwak, and Jin Young Choi. A comprehensive overhaul of feature distillation. In *Proc. ICCV*, pages 1921–1930, 2019. [3](#)
- [18] Byeongho Heo, Sangdoon Yun, Dongyoon Han, Sanghyuk Chun, Junsuk Choe, and Seong Joon Oh. Rethinking spatial dimensions of vision transformers. In *Proc. ICCV*, pages 11936–11945, 2021. [2](#)
- [19] Alexander Hermans, Georgios Floros, and Bastian Leibe. Dense 3d semantic mapping of indoor scenes from rgb-d images. In *Proc. ICRA*, pages 2631–2638, 2014. [2](#)
- [20] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NeurIPS Deep Learning and Representation Learning Workshop*, 2015. [3](#)
- [21] Ajay Jain, Matthew Tancik, and Pieter Abbeel. Putting nerf on a diet: Semantically consistent few-shot view synthesis. In *Proc. ICCV*, pages 5885–5894, 2021. [2](#)
- [22] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proc. ICLR*, 2015. [5](#)
- [23] Sosuke Kobayashi, Eiichi Matsumoto, and Vincent Sitzmann. Decomposing NeRF for editing via feature field distillation. *arXiv preprint arXiv:2205.15585*, 2022. [3](#), [8](#)
- [24] Amit Pal Singh Kohli, Vincent Sitzmann, and Gordon Wetstein. Semantic implicit neural scene representations with semi-supervised training. In *Proc. 3DV*, pages 423–433. IEEE, 2020. [2](#)
- [25] Abhijit Kundu, Kyle Genova, Xiaoqi Yin, Alireza Fathi, Caroline Pantofaru, Leonidas J Guibas, Andrea Tagliasacchi, Frank Dellaert, and Thomas Funkhouser. Panoptic neural fields: A semantic object-aware neural scene representation. In *Proc. CVPR*, pages 12871–12881, 2022. [2](#), [3](#)
- [26] Verica Lazova, Vladimir Guzov, Kyle Olszewski, Sergey Tulyakov, and Gerard Pons-Moll. Control-NeRF: Editable feature volumes for scene rendering and manipulation. *arXiv preprint arXiv:2204.10850*, 2022. [2](#)
- [27] Chunyuan Li, Jianwei Yang, Pengchuan Zhang, Mei Gao, Bin Xiao, Xiyang Dai, Lu Yuan, and Jianfeng Gao. Efficient self-supervised vision transformers for representation learning. In *Proc. ICLR*, 2022. [2](#)
- [28] Ke Li and Jitendra Malik. Amodal instance segmentation. In *Proc. ECCV*, pages 677–693, 2016. [5](#)
- [29] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Proc. CVPR*, pages 6498–6508, 2021. [2](#)
- [30] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proc. ICCV*, pages 10012–10022, 2021. [2](#)
- [31] Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradi-

- ent descent with warm restarts. In *Proc. ICLR*, 2017. 5
- [32] David G Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004. 3
- [33] John McCormac, Ankur Handa, Andrew Davison, and Stefan Leutenegger. Semanticfusion: Dense 3d semantic mapping with convolutional neural networks. In *Proc. ICRA*, pages 4628–4635, 2017. 2
- [34] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. Deep spectral methods: A surprisingly strong baseline for unsupervised semantic segmentation and localization. In *Proc. CVPR*, 2022. 2, 4
- [35] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Trans. on Graphics (TOG)*, 2019. 5
- [36] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *Proc. ECCV*, pages 405–421, 2020. 1, 2, 5
- [37] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proc. CVPR*, pages 11453–11464, 2021. 2
- [38] Stanley Osher, Ronald Fedkiw, and K Piechor. Level set methods and dynamic implicit surfaces. *Appl. Mech. Rev.*, 57(3):B15–B15, 2004. 2
- [39] Julian Ost, Fahim Mannan, Nils Thuerey, Julian Knodt, and Felix Heide. Neural scene graphs for dynamic scenes. In *Proc. CVPR*, pages 2856–2865, 2021. 2
- [40] Jeong Joon Park, Peter Florence, Julian Straub, Richard A. Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proc. CVPR*, 2019. 2
- [41] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proc. CVPR*, pages 5865–5874, 2021. 2
- [42] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *Proc. CVPR*, 2019. 3
- [43] Nikolaos Passalis and Anastasios Tefas. Learning deep representations with probabilistic knowledge transfer. In *Proc. ECCV*, pages 268–284, 2018. 3
- [44] 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. 2
- [45] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proc. CVPR*, pages 10318–10327, 2021. 2
- [46] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *Proc. ICML*, pages 8748–8763, 2021. 2
- [47] Zhongzheng Ren, Aseem Agarwala, Bryan Russell, Alexander G Schwing, and Oliver Wang. Neural volumetric object selection. In *Proc. CVPR*, pages 6133–6142, 2022. 2
- [48] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In Yoshua Bengio and Yann LeCun, editors, *Proc. ICLR*, 2015. 3
- [49] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015. 5
- [50] Mert Bulent Sariyildiz, Yannis Kalantidis, Diane Larlus, and Karteek Alahari. Concept generalization in visual representation learning. In *Proc. ICCV*, pages 9629–9639, 2021. 6
- [51] Prafull Sharma, Ayush Tewari, Yilun Du, Sergey Zakharov, Rares Ambrus, Adrien Gaidon, William T Freeman, Fredo Durand, Joshua B Tenenbaum, and Vincent Sitzmann. Seeing 3d objects in a single image via self-supervised static-dynamic disentanglement. *arXiv preprint arXiv:2207.11232*, 2022. 2
- [52] Oriane Siméoni, Gilles Puy, Huy V. Vo, Simon Roburin, Spyros Gidaris, Andrei Bursuc, Patrick Pérez, Renaud Marlet, and Jean Ponce. Localizing objects with self-supervised transformers and no labels. In *Proc. BMVC*, November 2021. 2
- [53] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3D-structure-aware neural scene representations. *Proc. NeurIPS*, 2019. 2
- [54] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, W Yifan, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, et al. Advances in neural rendering. In *Computer Graphics Forum*, volume 41, pages 703–735, 2022. 2
- [55] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. In *Proc. ICLR*, 2020. 3
- [56] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *Proc. ICML*, pages 10347–10357, 2021. 2, 5, 6
- [57] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. In *Proc. ICCV*, pages 32–42, 2021. 2
- [58] Edgar Tretschk, 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. In *Proc. CVPR*, pages 12959–12970, 2021. 2
- [59] Vadim Tschernezki, Diane Larlus, and Andrea Vedaldi. NeuralDiff: Segmenting 3D objects that move in egocentric videos. In *Proc. 3DV*, pages 910–919, 2021. 2, 4, 5
- [60] Narek Tumanyan, Omer Bar-Tal, Shai Bagon, and Tali Dekel. Splicing vit features for semantic appearance transfer. In *Proc. CVPR*, pages 10748–10757, 2022. 2
- [61] Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky. Deep image prior. In *Proc. CVPR*, 2018. 2
- [62] Vibhav Vineet, Ondrej Miksik, Morten Lidegaard, Matthias Nießner, Stuart Golodetz, Victor A Prisacariu, Olaf Kähler, David W Murray, Shahram Izadi, Patrick Pérez, et al. Incremental dense semantic stereo fusion for large-scale semantic scene reconstruction. In *Proc. ICRA*, pages 75–82, 2015. 2

- [63] Suhani Vora, Noha Radwan, Klaus Greff, Henning Meyer, Kyle Genova, Mehdi S. M. Sajjadi, Etienne Pot, Andrea Tagliasacchi, and Daniel Duckworth. NeSF: Neural semantic fields for generalizable semantic segmentation of 3d scenes. *Transactions on Machine Learning Research*, 2022. 2
- [64] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In *Proc. CVPR*, pages 3835–3844, 2022. 2
- [65] Yangtao Wang, Xi Shen, Shell Xu Hu, Yuan Yuan, James L. Crowley, and Dominique Vaufreydaz. Self-supervised transformers for unsupervised object discovery using normalized cut. In *Proc. CVPR*, pages 14543–14553, 2022. 2
- [66] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. Cvt: Introducing convolutions to vision transformers. In *Proc. ICCV*, pages 22–31, 2021. 2
- [67] Qianyi Wu, Xian Liu, Yuedong Chen, Kejie Li, Chuanxia Zheng, Jianfei Cai, and Jianmin Zheng. Object-compositional neural implicit surfaces. *arXiv preprint arXiv:2207.09686*, 2022. 2
- [68] Christopher Xie, Keunhong Park, Ricardo Martin-Brualla, and Matthew Brown. Fig-nerf: Figure-ground neural radiance fields for 3d object category modelling. In *Proc. 3DV*, pages 962–971. IEEE, 2021. 2
- [69] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond. *Computer Graphics Forum*, 2022. 2
- [70] Bangbang Yang, Yinda Zhang, Yinghao Xu, Yijin Li, Han Zhou, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. Learning object-compositional neural radiance field for editable scene rendering. In *Proc. ICCV*, pages 13779–13788, 2021. 2
- [71] Lin Yen-Chen. Nerf-pytorch. <https://github.com/yenchenlin/nerf-pytorch/>, 2020. 5
- [72] Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *Proc. CVPR*, pages 4133–4141, 2017. 3
- [73] Hong-Xing Yu, Leonidas Guibas, and Jiajun Wu. Unsupervised discovery of object radiance fields. In *Proc. ICLR*, 2022. 2
- [74] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zi-Hang Jiang, Francis EH Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *Proc. ICCV*, pages 558–567, 2021. 2
- [75] 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. 2
- [76] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *Proc. ICLR*, 2017. 3
- [77] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew Davison. In-place scene labelling and understanding with implicit scene representation. In *Proc. ICCV*, pages 15838–15847, 2021. 2
- [78] Shuaifeng Zhi, Edgar Sucar, Andre Mouton, Iain Haughton, Tristan Laidlow, and Andrew J Davison. ilabel: Interactive neural scene labelling. *arXiv preprint arXiv:2111.14637*, 2021. 2
- [79] Daquan Zhou, Bingyi Kang, Xiaojie Jin, Linjie Yang, Xiaochen Lian, Zihang Jiang, Qibin Hou, and Jiashi Feng. Deepvit: Towards deeper vision transformer. *arXiv preprint arXiv:2103.11886*, 2021. 2
- [80] Adrian Ziegler and Yuki M Asano. Self-supervised learning of object parts for semantic segmentation. In *Proc. CVPR*, pages 14502–14511, 2022. 2