Unsupervised Continual Semantic Adaptation through Neural Rendering

Zhizheng Liu[1*] Francesco Milano[1*] Jonas Frey[1,2] Roland Siegwart[1]
Hermann Blum[1†] Cesar Cadena[1†]
[1]ETH Zurich [2]Max Planck ETH Center for Learning Systems

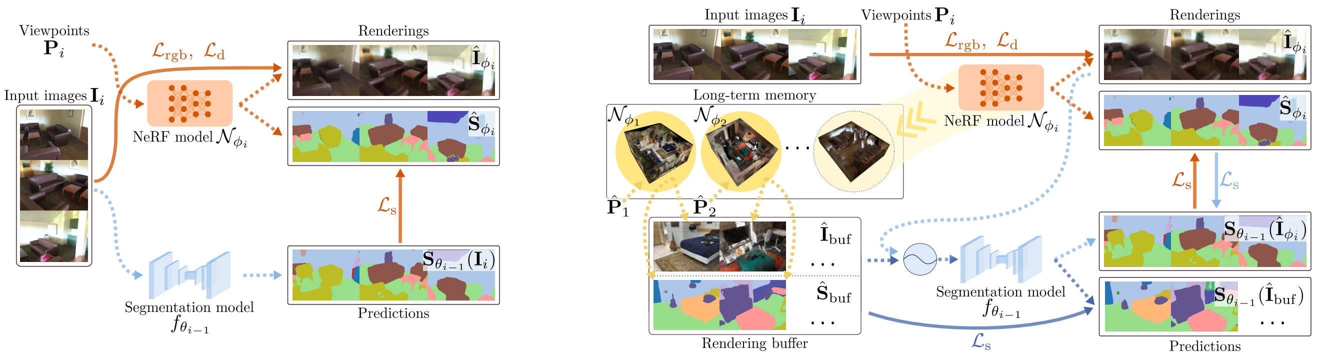

(a) Pseudo-label formation (b) Adaptation through joint training and long-term memory

Figure 1. We propose a method to continually adapt a semantic segmentation model f in an unsupervised fashion across multiple scenes, using neural rendering. For each scene $\mathcal{S}_i$: a) RGB(-D) images $\mathbf{I}_i$ from multiple viewpoints $\mathbf{P}_i$ and their corresponding predictions $\mathbf{S}_{\theta_{i-1}}(\mathbf{I}_i)$ by the latest model $f_{\theta_{i-1}}$ are used to supervise a (Semantic-)NeRF model $\mathcal{N}_{\phi_i}$; b) Adaptation on $\mathcal{S}_i$ is performed through a *joint training*, in which the segmentation network is supervised using the 3D-aware, view-consistent pseudo-labels $\hat{\mathbf{S}}_{\phi_i}$ rendered from $\mathcal{N}_{\phi_i}$ and the NeRF model through the smooth predictions of $f_{\theta_{i-1}}$. For each scene, the NeRF model can be compactly stored in a long-term memory, from which images and pseudo-labels from arbitrary viewpoints $\hat{\mathbf{P}}$ can be rendered into a fixed-size rendering buffer and mixed with the renderings from the current scene to reduce forgetting. Bold and dotted lines denote supervision signals and inputs/outputs, respectively.

Abstract

An increasing amount of applications rely on data-driven models that are deployed for perception tasks across a sequence of scenes. Due to the mismatch between training and deployment data, adapting the model on the new scenes is often crucial to obtain good performance. In this work, we study continual multi-scene adaptation for the task of semantic segmentation, assuming that no ground-truth labels are available during deployment and that performance on the previous scenes should be maintained. We propose training a Semantic-NeRF network for each scene by fusing the predictions of a segmentation model and then using the view-consistent rendered semantic labels as pseudo-labels to adapt the model. Through joint training with the segmentation model, the Semantic-NeRF model effectively enables 2D-3D knowledge transfer. Furthermore, due to its compact size, it can be stored in a long-term memory and subsequently used to render data from arbitrary viewpoints to reduce forgetting. We evaluate our approach on ScanNet, where we outperform both a voxel-based baseline and a state-of-the-art unsupervised domain adaptation method.

*Authors share first authorship. † Authors share senior authorship.

1. Introduction

Data-driven models trained for perception tasks play an increasing role in applications that rely on scene understanding, including, *e.g.*, mixed reality and robotics. When deploying these models on real-world systems, however, mismatches between the data used for training and those encountered during deployment can lead to poor performance, prompting the need for an *adaptation* of the models to the new environment. Oftentimes, the supervision data required for this adaptation can only be obtained through a laborious labeling process. Furthermore, even when such data are available, a naïve adaptation to the new environment results in decreased performance on the original training data, a phenomenon known as *catastrophic forgetting* [19, 23].

In this work, we focus on the task of adapting a semantic segmentation network across multiple indoor scenes, under the assumption that no labeled data from the new environment are available. Similar settings are explored in the literature in the areas of *unsupervised domain adaptation (UDA)* [23, 38] and *continual learning (CL)* [19]. However, works in the UDA literature usually focus on a single source-to-target transfer where the underlying assumption is that the data from both the source and the target domain are available all at once in the respective training stage, and often study the setting in which the knowledge transfer happens between a synthetic and a real environment [6, 31, 32, 38]. On the other hand, the CL community, which generally explores the adaptation of networks across different *tasks*, has established the *class-incremental* setting as the standard for semantic segmentation, in which new classes are introduced across different scenes from the same domain and ground-truth supervision is provided [23]. In contrast, we propose to study network adaptation in a setting that more closely resembles the deployment of semantic networks on real-world systems. In particular, instead of assuming that data from a specific domain are available all at once, we focus on the scenario in which the network is sequentially deployed in multiple *scenes* from a real-world indoor environment (we use the ScanNet dataset [7]), and therefore has to perform multiple *stages* of adaptation from one scene to another. Our setting further includes the possibility that previously seen scenes may be revisited. Hence, we are interested in achieving high prediction accuracy on each new scene, while at the same time preserving performance on the previous ones. Note that unlike the better explored class-incremental CL, in this setting we assume a *closed set* of semantic categories, but tackle the covariate shift across scenes without the need for ground-truth labels. We refer to this setting as *continual semantic adaptation*.

In this work, we propose to address this adaptation problem by leveraging advances in neural rendering [27]. Specifically, in a similar spirit to [11], when deploying a pre-trained network in a new scene, we aggregate the semantic predictions from the multiple viewpoints traversed by the agent into a 3D representation, from which we then render pseudo-labels that we use to adapt the network on the current scene. However, instead of relying on a voxel-based representation, we propose to aggregate the predictions through a semantics-aware NeRF [27, 48]. This formulation has several advantages. First, we show that using NeRFs to aggregate the semantic predictions results in higher-quality pseudo-labels compared to the voxel-based method of [11]. Moreover, we demonstrate that using these pseudo-labels to adapt the segmentation network results in superior performance compared both to [11] and to the state-of-the-art UDA method CoTTA [43]. An even more interesting insight, however, is that due the differentiability of NeRF, we can jointly train the frame-level semantic network and the scene-level NeRF to enforce similarity between the predictions of the former and the renderings of the latter. Remarkably, this joint procedure induces better performance of both labels, showing the benefit of mutual 2D-3D knowledge transfer.

A further benefit of our method is that after adapting to a new scene, the NeRF encoding the appearance, geometry and semantic content for that scene can be compactly saved in long-term storage, which effectively forms a "memory bank" of the previous experiences and can be useful in reducing catastrophic forgetting. Specifically, by mixing pairs of semantic and color NeRF renderings from a small number of views in the previous scenes and from views in the current scene, we show that our method is able to outperform both the baseline of [11] and CoTTA [43] on the adaptation to the new scene and in terms of knowledge retention on the previous scenes. Crucially, the collective size of the NeRF models is lower than that of the explicit replay buffer required by [11] and of the teacher network used in CoTTA [43] up to several dozens of scenes. Additionally, each of the NeRF models stores a potentially infinite number of views that can be used for adaptation, not limited to the training set as in [11], and removes the need to explicitly keep color images and pseudo-labels in memory.

In summary, the main contributions of our work are the following: (i) We propose using NeRFs to adapt a semantic segmentation network to new scenes. We find that enforcing 2D-3D knowledge transfer by jointly adapting NeRF and the segmentation network on a given scene results in a consistent performance improvement; (ii) We address the problem of continually adapting the segmentation network across a sequence of scenes by compactly storing the NeRF models in a long-term memory and mixing rendered images and pseudo-labels from previous scenes with those from the current one. Our approach allows generating a potentially infinite number of views to use for adaptation at constant memory size for each scene; (iii) Through extensive experiments, we show that our method achieves better adaptation and performance on the previous scenes compared both to a recent voxel-based method that explored a similar setting [11] and to a state-of-the-art UDA method [43].

2. Related work

Unsupervised domain adaptation for semantic segmentation. Unsupervised domain adaptation (UDA) studies the problem of transferring knowledge between a source and a target domain under the assumption that no labeled data for the target domain are available. In the following, we provide an overview of the main techniques used in UDA for semantic segmentation and focus on those which are most closely related to our work; for a more extensive summary we refer the reader to the recent survey of [23].

The majority of the methods rely on auto-encoder CNN architectures, and perform network adaptation either at the level of the input data [3, 14, 20, 45, 46], of the intermediate network representations [4, 10, 14, 29, 46], or of the output predictions [3, 4, 10, 20, 22, 33, 34, 36, 41, 50, 51]. The main strategies adopted consist in: using adversarial learning techniques to enforce that the network representations have similar statistical properties across the two domains [3, 4, 10, 14, 20, 29, 34], performing image-to-image translation to align the data from the two domains [3, 14, 20, 45, 46], learning to detect non-discriminative feature representations for the target domain [18, 33], and using self-supervised learning based either on minimizing the pixel-level entropy in the target domain [41] or on self-training techniques [5, 20, 22, 36, 47, 50, 51]. The latter category of methods is the most related to our setting. In particular, a number of works use the network trained on the source domain to generate semantic predictions on the unlabeled target data; the obtained *pseudo-labels* are then used as a self-supervisory learning signal to adapt the network to the target domain. While our work and the self-training UDA methods both use pseudo-labels, the latter approaches neither exploit the sequential structure of the data nor explicitly enforce multi-view consistency in the predictions on the target data. Furthermore, approaches in UDA mostly focus on single-stage, sim-to-real transfer settings, often for outdoor environments, and generally assume that the data from each domain are available all at once during the respective training stage. In contrast, we focus on a multi-step adaptation problem, in which data from multiple scenes from an indoor environment are available sequentially.

Within the category of self-training methods, a number of works come closer to our setting by presenting techniques to achieve *continuous*, multi-stage domain adaptation. In particular, the recently proposed CoTTA [43] uses a student-teacher framework, in which the student network is adapted to a target environment through pseudo-labels generated by the teacher network, and stochastic restoration of the weights from a pre-trained model is used to preserve source knowledge. ACE [44] proposes a style-transfer-based adaptation with replay of feature statistics from previous domains, but assumes ground-truth source labels and focuses on changes of environmental conditions within the same scene. Finally, related to our method is also the recent work of Frey et al. [11], which addresses a similar problem as ours by aggregating predictions from different viewpoints in a target domain into a 3D voxel grid and rendering pseudo-labels, but does not perform multi-stage adaptation.

Continual learning for semantic segmentation. Continual learning for semantic segmentation (CSS) focuses on the problem of updating a segmentation network in a *class-incremental setting*, in which it is assumed that the domain is available in different *tasks* and that new classes are added over time in a sequential fashion [23]. The main objective consists in performing adaptation to the new task, mostly using only data from the current stage, while preventing forgetting of the knowledge from the previous tasks. The methods proposed in the literature typically adopt a combination of different strategies, including distilling knowledge from a previous model [1, 9, 24, 26], selectively freezing the network parameters [24, 26], enforcing regularization of the latent representations [25], and generating or crawling data from the internet to replay [21, 30]. While similarly to CSS methods we explore a continual setting in which the network is sequentially presented with data from the same domain, we do not tackle the class-incremental problem, and instead focus on a closed-set scenario with shifting distribution of classes and scene appearance. A further important difference is that while CSS methods assume each adaptation step to be supervised, in our setting no ground-truth labels from the current adaptation stage are available.

NeRF-based semantic learning. Since the introduction of NeRF [27], several works have proposed extensions to the framework to incorporate semantic information into the learned scene representation. Semantic-NeRF [48] first proposed jointly learning appearance, geometry, and semantics through an additional multi-layer perceptron (MLP) and by adapting the volume rendering equation to produce semantic logits. Subsequent works have further extended this framework along different directions, including combining NeRF with a feature grid and 3D convolutions to achieve generalization [40], interactively labeling scenes [49], performing panoptic segmentation [12, 17], and using pre-trained Transformer models to supervise few-shot NeRF training [15], edit scene properties [42], or distill knowledge for different image-level tasks [16, 39]. In our work, we rely on Semantic-NeRF, which we use to fuse predictions from a segmentation network and that we jointly train with the latter exploiting differentiability. We include the formed scene representation in a long-term memory and use it to render pseudo-labels to adapt the segmentation network.

3. Continual Semantic Adaptation

3.1. Problem definition

In our problem setting, which we refer to as *continual semantic adaptation*, we assume we are provided with a segmentation model f_{θ_0}, with parameters θ_0, that was pre-trained on a dataset $\mathcal{P} = \left(\mathbf{I}_{\text{pre}}, \mathbf{S}^{\star}_{\text{pre}}\right)$. Here $\mathbf{I}_{\text{pre}}$ is a set of input color images (potentially with associated depth information) and $\mathbf{S}^{\star}_{\text{pre}}$ are the corresponding pixel-wise ground-truth semantic labels. We aim to adapt f_{θ_0} across a sequence of N scenes $\mathcal{S}_i,\ i \in \{1, \dots, N\}$ for each of which a set $\mathbf{I}_i$ of color (and depth) images, are collected from different viewpoints, but no ground-truth semantic labels are available. We assume that the input data $\{\mathbf{I}_{\text{pre}}, \mathbf{I}_1, \dots, \mathbf{I}_N\}$ originate from similar indoor environments (for instance, we do not consider simultaneously synthetic and real-world data) and that the classes to be predicted by the network belong to a closed set and are all known from the pre-training. For each scene $\mathcal{S}_i,\ i \in \{1, \dots, N\}$, the objective is to find a set of weights θ_i of the network, starting from θ_{i-1}, such that the performance of f_{θ_i} on $\mathcal{S}_i$ is higher than that of $f_{\theta_{i-1}}$. Additionally, it is desirable to preserve the performance of f_{θ_i} on the previous scenes $\{\mathcal{S}_1, \dots, \mathcal{S}_{i-1}\}$, in other words mitigate catastrophic forgetting.

The proposed setting aims to replicate the scenario of the deployment of a segmentation network on a real-world perception system (for instance a robot, or an augmented reality platform), where multiple sequential experiences are collected across similar scenes, and only limited data of the previous scenes can be stored on an on-board computing unit. During deployment, environments might be revisited over time, rendering the preservation of previously learned knowledge essential for a successful deployment.

3.2. Methodology

We present a method to address continual semantic adaptation in a self-supervised fashion (Fig. 1). In the following, $\boldsymbol{I}_i^k$ and $\boldsymbol{P}_i^k$ are the k-th RGB(-D) image collected in scene $\mathcal{S}_i$ and its corresponding camera pose, where $k \in \{1, \ldots, |\mathbf{I}_i|\}$. We further denote with $\mathbf{S}_\theta(\boldsymbol{I}_i^k)$ the prediction produced by f_θ for $\boldsymbol{I}_i^k$[1]. With a slight abuse of notation, we use $\mathbf{S}_\theta(\mathbf{I}_i)$ in place of $\{\mathbf{S}_\theta(\boldsymbol{I}_i^k),\ \boldsymbol{I}_i^k \in \mathbf{I}_i\}$ and similarly for other quantities that are a function of elements in a set.

For each new scene $\mathcal{S}_i$, we train a Semantic-NeRF [48] model $\mathcal{N}_{\phi_i}$, with learnable parameters ϕ_i, given for each viewpoint $\boldsymbol{P}_i^k$ the corresponding semantic label $\mathbf{S}_{\theta_j}(\boldsymbol{I}_i^k)$ predicted by a previous version f_{θ_j}, $j < i$, of the segmentation model. From the trained Semantic-NeRF model $\mathcal{N}_{\phi_i}$ we render semantic pseudo-labels $\hat{\mathbf{S}}_{\phi_i}$ and images $\hat{\mathbf{I}}_{\phi_i}$. The key observation at the root of our self-supervised adaptation is that semantic labels should be *multi-view consistent*, since they are constrained by the scene geometry that defines them. While the predictions of f often do not reflect this constraint because they are produced for each input frame independently, the NeRF-based pseudo-labels are by construction multi-view consistent. Inspired by [11], we hypothesize that this consistency constitutes an important prior that can be exploited to guide the adaptation of the network to the scene. Therefore, we use the renderings from $\mathcal{N}_i$ to adapt the segmentation network on scene $\mathcal{S}_i$, by minimizing a cross-entropy loss between the pseudo-labels and the network predictions. Crucially, we can use the NeRF and segmentation network predictions to supervise each other, allowing for joint optimization and adaptation of the two networks, which we find further improves the performance of both models.

To continually adapt the segmentation network f to multiple scenes in a sequence $\mathcal{S}_1 \rightarrow \mathcal{S}_2 \rightarrow \cdots \rightarrow \mathcal{S}_N$ and prevent catastrophic forgetting, we leverage the compact representation of NeRF by storing the corresponding model weights ϕ_i after adaptation in a long-term memory for each scene $\mathcal{S}_i$. Given that a trained NeRF can be queried from any viewpoint, this formulation allows generating for each scene a theoretically infinite number of views for adaptation, at the fixed storage cost given by the size of ϕ_i. For each previous scene $\mathcal{S}_j$, images $\hat{\mathbf{I}}_{\phi_j}$ and pseudo-labels $\hat{\mathbf{S}}_{\phi_j}$ from both previously seen and novel viewpoints can be rendered and used in an experience replay strategy to mitigate catastrophic forgetting on the previous scenes. An overview of our method is shown in Fig. 1.

NeRF-based pseudo-labels. We train for each scene a NeRF [27] model, which implicitly learns the geometry and appearance of the environment from a sparse set of posed images and can be used to render photorealistic novel views. More specifically, we extend the NeRF formulation by adding a semantic head as in Semantic-NeRF [48], and we render semantic labels $\hat{\mathbf{S}}_\phi$ by aggregating through the learned density function the semantic-head predictions for M sample points along each camera ray $\mathbf{r}$, as follows:

$$\hat{\mathbf{S}}_\phi(\mathbf{r}) = \sum_{i=1}^{M} T_i \alpha_i \mathbf{s}_i, \tag{1}$$

where $\alpha_i = 1 - e^{-\sigma_i \delta_i}, T_i = \prod_{j=1}^{i-1}(1-\alpha_j)$, with δ_i being the distance between adjacent sample points along the ray, and σ_i and $\mathbf{s}_i$ representing the predicted density and semantic logits at the i-th sample point along the ray, respectively.

We observe that if Semantic-NeRF is directly trained on the labels predicted by a pre-trained segmentation network on a new scene, the lack of view consistency of these labels can severely degrade the quality of the learned geometry, which in turn hurts the performance of the rendered semantic labels. To alleviate the influence of the inconsistent labels on the geometry, we propose to adopt several modifications. First, we stop the gradient flow from the semantic head into the density head. Second, we use depth supervision, as introduced in [8], to regularize the depth $\hat{d}(\mathbf{r}) = \sum_{i=1}^{N} T_i \alpha_i \delta_i$ rendered by NeRF via ℓ_1 loss with respect to the ground-truth depth $d(\mathbf{r})$:

$$\mathcal{L}_{\mathrm{d}}(\mathbf{r}) = \left\| \hat{d}(\mathbf{r}) - d(\mathbf{r}) \right\|_1. \tag{2}$$

Through ablations in the Supplementary, we show that this choice is particularly effective at improving the quality of both the geometry and the rendered labels. Additionally, we note that since the semantic logits $\mathbf{s}_i$ of each sampled point are unbounded, the logits $\hat{\mathbf{S}}_\phi(\mathbf{r})$ of the ray $\mathbf{r}$ can be dominated by a sampled point with very large semantic logits instead of one that is near the surface of the scene. This could cause the semantic labels generated by the NeRF model to overfit the initial labels of the segmentation model and lose multi-view consistency even when the learned geometry is correct. To address this issue, we instead first apply softmax to the logits of each sampled point, so these are normalized and contribute to the final aggregated logits through the weighting induced by volume rendering, as follows:

$$\hat{\mathbf{S}}'_\phi(\mathbf{r}) = \sum_{i=1}^{N} T_i \alpha_i {\cdot} \mathrm{softmax}(\mathbf{s}_i),\ \hat{\mathbf{S}}_\phi(\mathbf{r}) = \hat{\mathbf{S}}'_\phi(\mathbf{r}) / \|\hat{\mathbf{S}}'_\phi(\mathbf{r})\|_1. \tag{3}$$

The final normalized $\hat{\mathbf{S}}_\phi(\mathbf{r})$ is then a categorical distribution

[1] Note that in our experiments f_θ does not use the depth channel of $\boldsymbol{I}_i^k$.

$(\hat{S}(\mathbf{r})_1, \cdots, \hat{S}(\mathbf{r})_C)$ over the C semantic classes predicted by NeRF, and we use a negative log-likelihood loss to supervise the rendered semantic labels with the predictions of the semantic network:

$$\mathcal{L}_{\mathrm{s}}(\mathbf{r}) = -\sum_{c=1}^{C} \log(\hat{S}(\mathbf{r})_c) \cdot \mathbb{1}_{c=c(\mathbf{r})}, \tag{4}$$

where $c(\mathbf{r})$ is the semantic label predicted by the segmentation network f_θ. We train the NeRF model by randomly sampling rays from the training views and adding together the losses in (2) and (4), as well as the usual ℓ_2 loss $\mathcal{L}_{\mathrm{rgb}}(\mathbf{r})$ on the rendered color [27], as follows:

$$\mathcal{L} = \sum_{i=1}^{R} \mathcal{L}_{\mathrm{rgb}}(\mathbf{r_i}) + w_{\mathrm{d}} \mathcal{L}_{\mathrm{d}}(\mathbf{r_i}) + w_{\mathrm{s}} \mathcal{L}_{\mathrm{s}}(\mathbf{r_i}), \tag{5}$$

where R is the number of rays sampled for each batch and w_{d}, w_{s} are the weights for the depth loss and the semantic loss, respectively. After training the NeRF model, we render from it both color images $\hat{\mathbf{I}}_\phi$ and semantic labels $\hat{\mathbf{S}}_\phi$, as *pseudo-labels* for adapting the segmentation network.

Being able to quickly fuse the semantic predictions and generate pseudo-labels might be of particular importance in applications that require fast, possibly online adaptation. To get closer to this objective, we adopt the multi-resolution hash encoding proposed in Instant-NGP [28], which significantly improves the training and rendering speed compared to the original NeRF formulation. In the Supplementary, we compare the quality of the Instant-NGP-based pseudo-labels and those obtained with the original implementation from [48], and show that our method is agnostic to the specific NeRF implementation chosen.

Adaptation through joint 2D-3D training. To adapt the segmentation network f_{θ_j} on a given scene $\mathcal{S}_i$ (where $i > j$), we use the rendered pseudo-labels $\hat{\mathbf{S}}_{\phi_i}$ as supervisory signal by optimizing a cross-entropy loss between the network predictions $\mathbf{S}_{\theta_j}$ and $\hat{\mathbf{S}}_{\phi_i}$, similarly to previous approaches in the literature [11, 43, 44]. However, we propose two important modifications enabled by our particular setup and by its end-to-end differentiability. First, rather than adapting via the segmentation predictions for the ground-truth input images $\mathbf{I}_i$, we use $\mathbf{S}_{\theta_j}(\hat{\mathbf{I}}_{\phi_i})$, that is, we feed the *rendered* images as input to f. This removes the need for explicitly storing images for later stages, allows the adaptation to use novel viewpoints for which no observations were made, and as we show in our experiments, results in improved performance over the use of ground-truth images.

Second, we propose to *jointly train* $\mathcal{N}_{\phi_i}$ and f_{θ_j} by iteratively generating labels from one and back-propagating the cross-entropy loss gradients through the other in each training step. In practice, to initialize the NeRF pseudo-labels we first pre-train $\mathcal{N}_{\phi_i}$ with supervision of the ground-truth input images $\mathbf{I}_i$ and of the associated segmentation predictions $\mathbf{S}_{\theta_j}(\mathbf{I}_i)$, and then jointly train $\mathcal{N}_{\phi_i}$ and f_{θ_j} as described above. We demonstrate the positive influence of this joint adaptation in the experiments, where we show in particular that this 2D-3D knowledge transfer effectively produces improvements in the visual content of both the network predictions and the pseudo-labels.

Continual NeRF-based replay. A simple but effective approach to alleviate catastrophic forgetting as the adaptation proceeds across scenes is to *replay* previous experiences, *i.e.*, storing the training data of each newly-encountered scene in a memory buffer, and for each subsequent scene, training the segmentation model using both the data from the new scene and those replayed from the buffer, as done for instance in [11]. In practice, the size of the replay buffer is often limited due to memory and storage constraints, thus one can only store a subset of the data for replay, resulting in a loss of potentially useful information. Unlike previous methods that save explicit data into a buffer, we propose storing the NeRF models in a long-term memory. The advantages of this choice are multifold. First, the memory footprint of multiple NeRF models is significantly smaller than that of explicit images and labels (required by [11]) or of the weights of the segmentation network, stored by [43]. Second, since the NeRF model stores both color and semantic information and attains photorealistic fidelity, it can be used to render a theoretically infinite amount of training views at a fixed storage cost (unlike [11], which fits semantics in the map, and could not produce photorealistic renderings even if texture was aggregated in 3D). Therefore, the segmentation network can be provided with images rendered from NeRF as input. As we show in the experiments, by rendering a small set of views from the NeRF models stored in the long-term memory, our method is able to effectively mitigate catastrophic forgetting.

4. Experiments

4.1. Experimental settings

Dataset. We evaluate our proposed method on the ScanNet [7] dataset. The dataset includes 707 unique indoor scenes, each containing RGB-D images with associated camera poses and manually-generated semantic annotations. In all the experiments we resize the images to a resolution of 320×240 pixels. Similarly to [11], we use scenes 11-707 in ScanNet to pre-train the semantic segmentation network, taking one image every 100 frames in each of these scenes, for a total of approximately 25 000 images. The pre-training dataset is randomly split into a training set of 20k frames and a validation set of 5k frames. We use scene 1-10 to adapt the pre-trained model (cf. Sec. 4.3, 4.4, 4.5); if the dataset contains more than one video sequence for a given scene, we select only the first one. We select the first 80% of the frames (we refer to them as *training views*) from each sequence to generate predictions with the segmentation network and fuse these into a 3D representation, both by training our Semantic-NeRF model and with the baseline of [11]. The last 20% of the frames (*validation views*) are instead used to test the adaptation performance of the semantic segmentation model on the scene. We stress that this pre-training-training-testing setup is close to a real-world application scenario of the segmentation model, in which in an initial stage the network is trained offline on a large dataset, then some data collected during deployment may be used to adapt the model in an unsupervised fashion, and finally the model performance is tested during deployment on a different trajectory.

Networks. We use DeepLabv3 [2] with a ResNet-101 [13] backbone as our semantic segmentation network. To implement our Semantic-NeRF network, we rely on an open-source PyTorch implementation [37] of Instant-NGP [28]. Further details about the architectures of both networks can be found in the Supplementary. For brevity, in the following Sections we refer to Semantic-NeRF as "NeRF".

Baselines. As there are no previous works that explicitly tackle the continual semantic adaptation problem, we compare our proposed method to the two most-closely related approaches. The first one [11] uses per-frame camera pose and depth information to aggregate predictions from a segmentation network into a voxel map and then renders semantic pseudo-labels from the map to adapt the network. We implement the method using the framework of [35] and use a voxel resolution of $5\,\mathrm{cm}$, as done in [11], which yields a total map size comparable to the memory footprint of the NeRF parameters (cf. Supplementary for further details). The second approach, CoTTA [43], focuses on continual test-time domain adaptation and proposes a student-teacher framework with label augmentation and stochastic weight restoration to gradually adapt the semantic segmentation model while keeping the knowledge on the source domain. We use the official open-source implementation, which we adapt to test its performance on the proposed setting.

Metric. For all the experiments, we report mean intersection over union (mIoU, in percentage values) as a metric.

4.2. Pre-training of the segmentation network

We pre-train DeepLab for 150 epochs to minimize the cross-entropy loss with respect to the ground-truth labels $\mathbf{S}^{\star}_{\mathrm{pre}}$. We apply common data augmentation techniques, including random flipping/orientation and color jitter. After pre-training, we select the model with best performance on the validation set for adaptation to the new scenes.

4.3. Pseudo-label formation

We train the NeRF network by minimizing (5) for 60 epochs using the training views. While with our method we can render pseudo-labels from any viewpoint, to allow a controlled comparison against [11] in Sec. 4.4 and 4.5, we generate the pseudo-labels from our NeRF model using the same training viewpoints. While the pseudo-labels of [11] are deterministic, to account for the stochasticity of NeRF, we run our method with 10 different random seeds and report the mean and variance over these. As shown in Tab. 1, the pseudo-labels produced by our method outperform on average those of [11]. A further improvement can be obtained by jointly training NeRF and the DeepLab model, which we discuss in the next Section.

	Pre-train	Mapping [11]	Ours	Ours Joint Training
Scene 1	41.1	48.9	48.8±0.7	**54.8**±1.8
Scene 2	35.5	33.9	36.2±0.8	**38.3**±0.4
Scene 3	23.5	25.1	**27.1**±0.9	26.4±1.8
Scene 4	62.8	**65.3**	62.9±0.5	65.0±1.1
Scene 5	49.8	49.3	**55.5**±1.3	46.6±0.2
Scene 6	48.9	**51.7**	50.4±0.4	50.9±0.4
Scene 7	39.7	41.2	40.4±0.5	**41.7**±2.0
Scene 8	31.6	34.8	34.0±0.4	**39.0**±4.6
Scene 9	31.7	33.8	**35.6**±0.4	31.3±0.4
Scene 10	52.5	55.8	**56.4**±0.6	56.2±1.0
Average	41.7	44.0	44.7±0.7	**45.0**±1.4

Table 1. Pseudo-label performance averaged over the training views and 10 different seeds for Ours pseudo-labels. "Pre-train" denotes the performance of the segmentation model f_{θ_0}.

4.4. One-step adaptation

As a first adaptation experiment, we evaluate the performance of the different methods when letting the segmentation network f_{θ_0} adapt in a single stage to each of the scenes 1-10. This setup is similar to that of one-stage UDA, and we thus compare to the state-of-the-art method CoTTA [43].

We evaluate our method in two different settings. In the first one, which we refer to as *fine-tuning*, we simply use the pseudo-labels rendered as in Sec. 4.3 to adapt the segmentation network through cross-entropy loss on its predictions. In the second one, we *jointly train* NeRF and DeepLab via iterative mutual supervision. For a fair comparison, in both settings we optimize the pre-trained NeRF for the same number of additional epochs, while maintaining supervision through color and depth images. In fine-tuning, we perform NeRF pre-training for 60 epochs, according to Sec. 4.3. In joint training, we instead first pre-train NeRF for 10 epochs, and then train NeRF concurrently with DeepLab for 50 epochs. We run each method 10 times and report mean and standard deviation across the runs. Given that the baselines do not support generating images from novel viewpoints, both in fine-tuning and in joint training we use images from the training viewpoints as input to DeepLab. Additionally, since our method allows *rendering* images, we evaluate the difference between feeding ground-truth images vs. NeRF renderings from the same viewpoints to DeepLab.

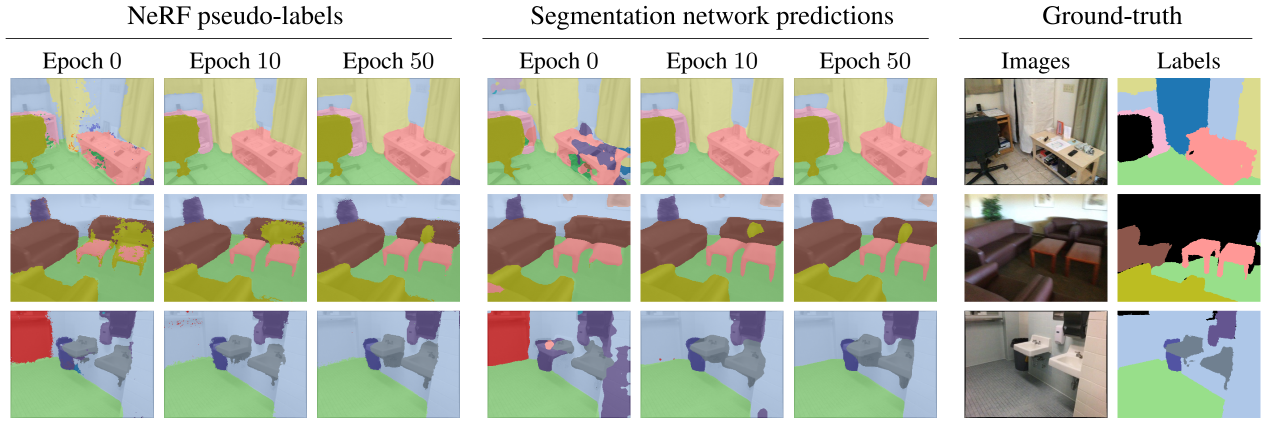

Figure 2. Effect of joint training over the pseudo-labels and the predictions of the segmentation network (DeepLab). Color-coded labels are overlaid on the corresponding color images. Black pixels in the ground-truth labels denote missing annotation. First scene: The noisy predictions of DeepLab are corrected and the segmentation results conform much better to the geometry of the scene. Second scene: The geometric details can be better recovered even for the legs of the table. Third scene: By enforcing multi-view consistency, the initial wrong predictions on the wall are corrected through the predictions from other views. Note that the obtained labels adhere accurately to the scene geometry, often even better than in the ground-truth annotations.

Table 2 presents the adaptation performance of the different methods on the validation views, which provides a measure of the knowledge transfer induced within the scene by the self-training. Since our method is unsupervised, in the Supplementary we additionally report the improvement in performance on the training views, which is indicative of the effectiveness of the self-supervised adaptation and is of practical utility for real-world deployment scenarios where a scene might be revisited from similar viewpoints.

As shown in Tab. 2, fine-tuning with our pseudo-labels results in improved performance compared to the pre-trained model, and outperforms both baselines for most of the scenes. Interestingly, using rendered images (NI + NL) consistently produces better results than fine-tuning with the ground-truth images (GI + NL). We hypothesize that this is due to the small image artifacts introduced by the NeRF rendering acting as an augmentation mechanism. We further observe that the mIoU can vary largely across the scenes. This can be explained with the variability in the room types and lighting conditions, which is also reflected in the scenes with more extreme illumination (and hence more challenging for NeRF to reconstruct the geometry) having a larger variance with our approach. However, the main observation is that jointly training NeRF and DeepLab (using rendered images as input) results in remarkably better adaptation on almost all the scenes. This improvement can be attributed to the positive knowledge transfer induced between the frame-level predictions of DeepLab and the 3D-aware NeRF pseudo-labels. As shown in Fig. 2, this strategy allows effectively resolving local artifacts in the NeRF pseudo-labels through the smoothing effect of the DeepLab labels, while at the same time addressing inconsistencies in the per-frame outputs of the segmentation network due to its lack of view consistency.

4.5. Multi-step adaptation

To evaluate our method in the full scenario of continual semantic adaptation, we perform multi-step adaptation across scenes 1-10, where in the i-th step the segmentation network $f_{\theta_{i-1}}$ gets adapted on scene $\mathcal{S}_i$, resulting in f_{θ_i}, and the NeRF model $\mathcal{N}_i$ is added to the long-term memory at the end of the stage. For steps $i \in \{2, \dots, 10\}$, to counteract forgetting on the previous scenes we render images and pseudo-labels for each of the $\mathcal{N}_j$ models $(1 \leq j \leq i-1)$ in the long-term memory. In practice, we construct a memory buffer of fixed size 100, to which at stage i each of the previous models $\mathcal{N}_j$ contribute equally with images $\hat{\mathbf{I}}_{\mathrm{buf}}$ and pseudo-labels $\hat{\mathbf{S}}_{\mathrm{buf}}$ rendered from $\lfloor 100/(i-1) \rfloor$ randomly chosen training views. Following [11], we additionally randomly select 10% of the pre-training data and combine them to the data from the previous scenes, which acts as prior knowledge and prevents the model from overfitting to the new scenes and losing its generalization performance. This has a similar effect to the regularization scheme used by CoTTA [43] to preserve previous knowledge, namely storing the network parameters for the initial pre-trained model and the teacher network. Note that both the size of our memory buffer (14 MB) and that of the replayed pre-training data (65 MB) are much smaller than the size of two sets of DeepLab weights (2×225 MB), so our method actually requires less storage space than CoTTA [43]. A detailed analysis of the memory footprint of the different approaches is presented in the Supplementary; we show in particular that since our method is agnostic to the specific NeRF implementation, with the slower but lighter implementation of Semantic-NeRF [48] the storage comparison is in our favor up to 90 scenes. We deem this to be a realistic margin for real-world deployment scenarios (*e.g.*, it is hardly the case that an agent sequentially visits more than a few scenes dur-

	Pre-train	CoTTA [43]	Fine-tuning (GI + ML)	Ours Fine-tuning (GI + NL)	Ours Fine-tuning (NI + NL)	Ours Joint Training
Scene 1	43.9	44.0±0.0	46.3±0.3	46.2±1.0	47.1±1.0	**50.0**±1.3
Scene 2	41.3	41.2±0.0	39.4±0.3	39.5±1.0	44.2±1.0	**47.1**±1.2
Scene 3	**23.0**	22.8±0.0	21.6±0.1	21.9±0.7	21.5±1.0	19.9±2.3
Scene 4	50.2	50.3±0.0	52.4±0.2	51.5±0.5	52.8±0.8	**53.7**±2.4
Scene 5	40.1	40.1±0.0	49.4±0.5	50.6±2.4	**52.8**±2.9	42.7±1.0
Scene 6	37.6	37.6±0.0	33.7±0.3	36.2±1.6	37.1±2.4	**40.8**±1.2
Scene 7	55.8	55.9±0.0	50.7±0.5	50.7±1.8	52.1±1.3	**56.5**±4.8
Scene 8	**27.9**	**27.9**±0.0	24.7±0.2	23.8±0.4	25.3±0.8	25.7±2.9
Scene 9	54.9	54.9±0.0	62.2±1.3	57.6±5.3	52.1±2.7	**63.7**±3.3
Scene 10	73.5	73.5±0.0	**73.8**±0.2	**73.8**±0.2	73.5±0.4	73.7±0.5
Average	44.8	44.8±0.0	45.4±0.4	45.2±1.5	45.9±1.4	**47.4**±2.1

Table 2. Performance of the segmentation network on the validation set of each scene after one-step adaptation. GI and NI denote respectively ground-truth color images and NeRF-rendered color images. ML and NL indicate adaptation using pseudo-labels formed respectively with the method of [11] and with our approach. In joint training, we use NeRF-based renderings and pseudo-labels.

		Step 1	Step 2	Step 3	Step 4	Step 5	Step 6	Step 7	Step 8	Step 9	Step 10	Average
Pre-train		43.9	41.3	23.0	50.2	40.1	37.6	55.8	**27.9**	54.9	73.5	44.8
Adapt	CoTTA [43]	44.0±0.0	40.9±0.0	22.7±0.0	50.2±0.1	40.0±0.0	37.5±0.0	56.0±0.1	26.9±0.0	54.5±0.0	**73.8**±0.0	44.7±0.0
	Mapping [11]	46.8±0.4	42.1±2.0	23.6±0.7	**50.6**±2.6	**44.0**±0.1	35.8±0.5	56.7±1.3	26.5±1.8	68.3±1.4	72.7±1.0	46.7±1.2
	Ours ($\mathbf{I}_{\text{pre}}$ replay only)	53.3±0.7	**48.0**±2.4	20.5±0.1	49.0±1.5	43.4±0.0	39.0±1.4	**62.1**±6.2	26.7±3.0	65.7±5.6	73.0±0.5	**48.1**±2.1
	Ours	**53.7**±1.3	46.3±0.7	**24.3**±2.0	49.1±0.9	43.7±0.3	**40.4**±1.5	55.8±0.8	26.2±0.9	**68.9**±3.2	72.5±1.6	**48.1**±1.3
Previous	CoTTA [43]	–	44.0±0.0	42.2±0.0	35.6±0.0	39.3±0.0	39.4±0.0	39.1±0.0	41.5±0.0	39.7±0.0	41.3±0.0	40.2±0.0
	Mapping [11]	–	46.5±0.1	42.8±1.0	37.3±0.9	40.4±0.6	40.9±0.7	39.9±1.1	42.2±0.5	40.0±0.4	42.8±0.7	41.4±0.7
	Ours ($\mathbf{I}_{\text{pre}}$ replay only)	–	52.3±0.3	47.5±1.1	38.6±1.1	40.8±0.7	42.4±0.3	41.5±0.7	**44.3**±1.4	41.4±0.6	43.9±0.9	43.6±0.8
	Ours	–	**53.2**±0.9	**48.2**±0.8	**41.5**±0.8	**42.8**±0.8	**43.2**±0.8	**42.2**±0.8	44.1±0.2	**41.7**±0.2	**44.3**±0.3	**44.6**±0.6

Table 3. Multi-step performance evaluated on the validation set of each scene. At Step i, $\mathrm{Pre\text{-}train}$ and Adapt denote respectively the performance of the pre-trained network f_{θ_0} and of the adapted network f_{θ_i} on the current scene $\mathcal{S}_i$, while $\mathrm{Previous}$ represents the average performance of f_{θ_i} on scenes $\mathcal{S}_1$ to $\mathcal{S}_{i-1}$. All Ours are with *joint training*.

ing the same mission). For the baseline of [11] we use the same setup as our method, but with mapping-based pseudo-labels and ground-truth images in the memory buffer, due to its inability to generate images. For a fair comparison, we use *training* views for replay also for our method. The latter, however, also allows generating data from *novel* viewpoints for replay; very interestingly, we find this to yield better knowledge retention (cf. Sec. D.2 of the Supplementary).

The multi-step adaptation results are shown in Tab. 3, where for each method the mean and standard deviation across 3 runs are reported. To better show the effect of NeRF-based replay, we also run our adaptation method with only replay from the pre-training dataset, without replaying from the old NeRF models (Ours ($\mathbf{I}_{\text{pre}}$ $\mathrm{replay\ only}$)). Our method achieves the best average adaptation performance (Adapt) across the new scenes in the multi-step setting, improving by $\sim 3\%$ mIoU over the pre-trained model. Note that this improvement is consistent with the one observed in one-step adaptation (Tab. 2), which validates that our method can successfully adapt across multiple scenes, without the performance dropping after a specific number of steps. At the same time, while NeRF-based replay of the old scenes on average does not induce a positive forward transfer in the adaptation to the new scenes (Adapt), its usage can significantly alleviate forgetting compared to the case with no replay. As a result, when using NeRF-based replay, our method is able to maintain in almost all the adaptation steps the best average performance over the previous scenes ($\mathrm{Previous}$). Further in-detail results for each scene and after each adaptation step are reported in the Supplementary.

5. Conclusion

In this work, we present a novel approach for unsupervised continual adaptation of a semantic segmentation network to multiple novel scenes using neural rendering. We exploit the fact that the new scenes are observed from multiple viewpoints and jointly train in each scene a Semantic-NeRF model and the segmentation network. We show that the induced 2D-3D knowledge transfer results in improved unsupervised adaptation performance compared to state-of-the-art methods. We further propose a NeRF-based replay strategy which allows efficiently mitigating catastrophic forgetting and enables rendering a potentially infinite number of images for adaptation at constant storage cost. We believe this opens up interesting avenues for replay-based adaptation, particularly for use on real-world perception systems, which can compactly store collected experiences on board and generate past data as needed. We discuss the limitations of our method in the Supplementary.

Acknowledgements: This work has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No. 101017008 (Harmony), the Max Planck ETH Center for Learning Systems, and the HILTI group. We thank Marco Hutter for his guidance and support.

References

[1] Fabio Cermelli, Massimiliano Mancini, Samuel Rota Buló, Elisa Ricci, and Barbara Caputo. Modeling the Background for Incremental Learning in Semantic Segmentation. In *CVPR*, 2020. 3

[2] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking Atrous Convolution for Semantic Image Segmentation. *CoRR:1706.05587*, 2017. 6

[3] Yun-Chun Chen, Yen-Yu Lin, Ming-Hsuan Yang, and Jia-Bin Huang. CrDoCo: Pixel-Level Domain Transfer With Cross-Domain Consistency. In *CVPR*, 2019. 2, 3

[4] Yi-Hsin Chen, Wei-Yu Chen, Yu-Ting Chen, Bo-Cheng Tsai, Yu-Chiang Frank Wang, and Min Sun. No More Discrimination: Cross City Adaptation of Road Scene Segmenters. In *ICCV*, 2017. 2, 3

[5] Jaehoon Choi, Taekyung Kim, and Changick Kim. Self-Ensembling With GAN-Based Data Augmentation for Domain Adaptation in Semantic Segmentation. In *ICCV*, 2019. 3

[6] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The Cityscapes Dataset for Semantic Urban Scene Understanding. In *CVPR*, 2016. 2

[7] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. ScanNet: Richly-Annotated 3D Reconstructions of Indoor Scenes. In *CVPR*, 2017. 2, 5

[8] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised NeRF: Fewer Views and Faster Training for Free. In *CVPR*, 2022. 4

[9] Arthur Douillard, Yifu Chen, Arnaud Dapogny, and Matthieu Cord. PLOP: Learning without Forgetting for Continual Semantic Segmentation. In *CVPR*, 2021. 3

[10] Liang Du, Jingang Tan, Hongye Yang, Jianfeng Feng, Xiangyang Xue, Qibao Zheng, Xiaoqing Ye, and Xiaolin Zhang. SSF-DAN: Separated Semantic Feature Based Domain Adaptation Network for Semantic Segmentation. In *ICCV*, 2019. 2, 3

[11] Jonas Frey, Hermann Blum, Francesco Milano, Roland Siegwart, and Cesar Cadena. Continual Adaptation of Semantic Segmentation using Complementary 2D-3D Data Representations. *IEEE Robot. Autom. Lett.*, 2022. 2, 3, 4, 5, 6, 7, 8

[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. In *3DV*, 2022. 3

[13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016. 6

[14] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. CyCADA: Cycle-Consistent Adversarial Domain Adaptation. In *ICML*, 2018. 2, 3

[15] Ajay Jain, Matthew Tancik, and Pieter Abbeel. Putting NeRF on a Diet: Semantically Consistent Few-Shot View Synthesis. In *ICCV*, 2021. 3

[16] Sosuke Kobayashi, Eiichi Matsumoto, and Vincent Sitzmann. Decomposing NeRF for Editing via Feature Field Distillation. In *NeurIPS*, 2022. 3

[17] Abhijit Kundu, Kyle Genova, Xiaoqi Yin, Alireza Fathi, Caroline Pantofaru, Leonidas Guibas, Andrea Tagliasacchi, Frank Dellaert, and Thomas Funkhouser. Panoptic Neural Fields: A Semantic Object-Aware Neural Scene Representation. In *CVPR*, 2022. 3

[18] Seungmin Lee, Dongwan Kim, Namil Kim, and Seong-Gyun Jeong. Drop to Adapt: Learning Discriminative Features for Unsupervised Domain Adaptation. In *ICCV*, 2019. 3

[19] Timothée Lesort, Vincenzo Lomonaco, Andrei Stoian, Davide Maltoni, David Filliat, and Natalia Díaz-Rodríguez. Continual Learning for Robotics: Definition, Framework, Learning Strategies, Opportunities and Challenges. *Information Fusion*, 58, 2020. 1

[20] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional Learning for Domain Adaptation of Semantic Segmentation. In *CVPR*, 2019. 2, 3

[21] Andrea Maracani, Umberto Michieli, Marco Toldo, and Pietro Zanuttigh. RECALL: Replay-based Continual Learning in Semantic Segmentation. In *ICCV*, 2021. 3

[22] Umberto Michieli, Matteo Biasetton, Gianluca Agresti, and Pietro Zanuttigh. Adversarial Learning and Self-Teaching Techniques for Domain Adaptation in Semantic Segmentation. *IEEE TIV*, 5, 2020. 2, 3

[23] Umberto Michieli, Marco Toldo, and Pietro Zanuttigh. Domain adaptation and continual learning in semantic segmentation. In *Advanced Methods and Deep Learning in Computer Vision*, chapter 8, pages 275–303. Elsevier, 2022. 1, 2, 3

[24] Umberto Michieli and Pietro Zanuttigh. Incremental Learning Techniques for Semantic Segmentation. In *ICCVW*, 2019. 3

[25] Umberto Michieli and Pietro Zanuttigh. Continual Semantic Segmentation via Repulsion-Attraction of Sparse and Disentangled Latent Representations. In *CVPR*, 2021. 3

[26] Umberto Michieli and Pietro Zanuttigh. Knowledge Distillation for Incremental Learning in Semantic Segmentation. *J. Comput. Vis. Image Understanding*, 205, 2021. 3

[27] 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 *ECCV*, 2020. 2, 3, 4, 5

[28] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM TOG*, 41(4):102:1–102:15, 2022. 5, 6

[29] Zak Murez, Soheil Kolouri, David Kriegman, Ravi Ramamoorthi, and Kyungnam Kim. Image to Image Translation for Domain Adaptation. In *CVPR*, 2018. 2, 3

[30] Mathieu Pagé Fortin and Brahim Chaib-draa. Continual Semantic Segmentation Leveraging Image-level Labels and Rehearsal. In *IJCAI*, 2022. 3

[31] Stephan R. Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for Data: Ground Truth from Computer Games. In *ECCV*, 2016. 2

[32] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M. Lopez. The SYNTHIA Dataset: A Large Collection of Synthetic Images for Semantic Segmentation of Urban Scenes. In *CVPR*, 2016. 2

[33] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Adversarial Dropout Regularization. In *ICLR*, 2018. 2, 3

[34] Swami Sankaranarayanan, Yogesh Balaji, Arpit Jain, Ser Nam Lim, and Rama Chellappa. Learning From Synthetic Data: Addressing Domain Shift for Semantic Segmentation. In *CVPR*, 2018. 2, 3

[35] Lukas Schmid, Jeffrey Delmerico, Johannes Schönberger, Juan Nieto, Marc Pollefeys, Roland Siegwart, and Cesar Cadena. Panoptic Multi-TSDFs: a Flexible Representation for Online Multi-resolution Volumetric Mapping and Long-term Dynamic Scene Consistency. In *ICRA*, 2022. 6

[36] Teo Spadotto, Marco Toldo, Umberto Michieli, and Pietro Zanuttigh. Unsupervised Domain Adaptation with Multiple Domain Discriminators and Adaptive Self-Training. In *ICPR*, 2021. 2, 3

[37] Jiaxiang Tang. Torch-ngp: a PyTorch implementation of instant-ngp. `https://github.com/ashawkey/torch-ngp`, 2022. 6

[38] Marco Toldo, Andrea Maracani, Umberto Michieli, and Pietro Zanuttigh. Unsupervised Domain Adaptation in Semantic Segmentation: a Review. *Technologies*, 8, 2020. 1, 2

[39] Vadim Tschernezki, Iro Laina, Diane Larlus, and Andrea Vedaldi. Neural Feature Fusion Fields: 3D Distillation of Self-Supervised 2D Image Representations. In *3DV*, 2022. 3

[40] 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. *TMLR*, 2022. 3

[41] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Perez. ADVENT: Adversarial Entropy Minimization for Domain Adaptation in Semantic Segmentation. In *CVPR*, 2019. 2, 3

[42] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. CLIP-NeRF: Text-and-Image Driven Manipulation of Neural Radiance Fields. In *CVPR*, 2022. 3

[43] Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continual Test-Time Domain Adaptation. In *CVPR*, 2022. 2, 3, 5, 6, 7, 8

[44] Zuxuan Wu, Xin Wang, Joseph E. Gonzalez, Tom Goldstein, and Larry S. Davis. ACE: Adapting to Changing Environments for Semantic Segmentation. In *ICCV*, 2019. 3, 5

[45] Yanchao Yang and Stefano Soatto. FDA: Fourier Domain Adaptation for Semantic Segmentation. In *CVPR*, 2020. 2, 3

[46] Yiheng Zhang, Zhaofan Qiu, Ting Yao, Dong Liu, and Tao Mei. Fully Convolutional Adaptation Networks for Semantic Segmentation. In *CVPR*, 2018. 2, 3

[47] Zhedong Zheng and Yi Yang. Rectifying Pseudo Label Learning via Uncertainty Estimation for Domain Adaptive Semantic Segmentation. *IJCV*, 2021. 3

[48] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew J. Davison. In-Place Scene Labelling and Understanding with Implicit Scene Representation. In *ICCV*, 2021. 2, 3, 4, 5, 7

[49] Shuaifeng Zhi, Edgar Sucar, Andre Mouton, Iain Haughton, Tristan Laidlow, and Andrew J. Davison. iLabel: Interactive Neural Scene Labelling. *CoRR:2111.14637*, 2021. 3

[50] Yang Zou, Zhiding Yu, B.V.K. Vijaya Kumar, and Jinsong Wang. Unsupervised Domain Adaptation for Semantic Segmentation via Class-Balanced Self-Training. In *ECCV*, 2018. 2, 3

[51] Yang Zou, Zhiding Yu, Xiaofeng Liu, B.V.K. Vijaya Kumar, and Jinsong Wang. Confidence Regularized Self-Training. In *ICCV*, 2019. 2, 3