

Instance-Guided Unsupervised Domain Adaptation for Robotic Semantic Segmentation

Michele Antonazzi, Lorenzo Signorelli, Matteo Luperto, Nicola Basilico

Abstract—Semantic segmentation networks, which are essential for robotic perception, often suffer from performance degradation when the visual distribution of the deployment environment differs from that of the source dataset on which they were trained. Unsupervised Domain Adaptation (UDA) addresses this challenge by adapting the network to the robot’s target environment without external supervision, leveraging the large amounts of data a robot might naturally collect during long-term operation. In such settings, UDA methods can exploit multi-view consistency across the environment’s map to fine-tune the model in an unsupervised fashion and mitigate domain shift. However, these approaches remain sensitive to cross-view instance-level inconsistencies. In this work, we propose a method¹ that starts from a volumetric 3D map to generate multi-view consistent pseudo-labels. We then refine these labels using the zero-shot instance segmentation capabilities of a foundation model, enforcing instance-level coherence. The refined annotations serve as supervision for self-supervised fine-tuning, enabling the robot to adapt its perception system at deployment time. Experiments on real-world data demonstrate that our approach consistently improves performance over state-of-the-art UDA baselines based on multi-view consistency, without requiring any ground-truth labels in the target domain.

I. INTRODUCTION

Semantic segmentation is a key component of robotic vision, as it enables autonomous systems to interpret complex environments and reason about object categories at the pixel level. This capability can boost core robotic tasks such as navigation, planning, and interaction [1]. Typically, segmentation models are trained on source datasets collected under controlled conditions, but their performance degrades when deployed in target environments characterized by different, though related, data distributions [2]. This phenomenon, known as *domain shift*, is particularly critical for mobile robots, which must operate for long periods in previously unseen environments where annotated data are unavailable at training time and prohibitively expensive to acquire during deployment. A common strategy to mitigate domain shift is to collect and manually annotate target data for fine-tuning [3], but this process is costly, time-consuming, and often impractical in real robotic deployments [4].

Unsupervised Domain Adaptation (UDA) addresses the challenge of adapting a model to a new environment without relying on ground-truth labels or humans in the loop [5]. With mobile robots, UDA can leverage domain-specific opportunities: unlike with static computer vision benchmarks, a robot is continuously immersed in its target environment, can acquire multiple observations of the same scene,




Fig. 1. Examples of the improvement provided by our method (third column) to the view-consistent pseudo-labels [6] (middle column).

and can exploit mapping and view consistency as natural proxies for model adaptation. Previous works [6], [7] have shown that aggregating predictions across viewpoints in a 3D environmental representation [8] improves pseudo-labels by enforcing spatial consistency, yet such methods remain vulnerable to instance-level incoherence that might persist across frames. As shown in Fig. 1, the view-consistent pseudo-labels (middle column) obtained from the 3D map present visual artifacts generated by the rendering process and, more importantly, fail to be instance-coherent as multiple categories are assigned to the same object.

In this work, we tackle and mitigate these issues by proposing an instance-guided approach to UDA for robotic semantic segmentation. Starting from multi-view consistent pseudo-labels, we integrate a refinement stage that leverages the zero-shot instance segmentation capabilities of a foundation model, queried via two different and complementary automated prompting strategies. Our method enforces instance-level coherence and mitigates rendering artifacts, yielding more accurate and stable annotations for self-supervised fine-tuning (see Fig. 1, third column). Experiments on real-world data show that our method significantly outperforms existing UDA baselines, while requiring no ground-truth labels in the target domain.

In summary, this work’s contributions are: (i) we propose a novel framework for UDA in robotic semantic segmentation that combines multi-view consistency with instance-aware pseudo-label refinement; (ii) we integrate the SAM foundation model into the refinement stage, devising two automatic and complementary prompting strategies; (iii) we carry out an extensive experimental evaluation on real-world data assessing how our method improves adaptation performance over existing baselines both in terms of pseudo-label quality and segmentation performance.

All authors are with the Department of Computer Science, University of Milan, Milano, Italy. Email name.surname@unimi.it

¹<https://aislab.di.unimi.it/research/instanceuda>

II. RELATED WORKS

UDA approaches for semantic segmentation. Domain shift is caused by a difference in the distribution of the dataset collected in the environments used for training and evaluation. One of the mainstream approaches for unsupervised domain adaptation in semantic segmentation tasks leverages adversarial learning to align the distributions of train and test domains used at input, intermediate features, and/or output levels. The pioneer work of [9] presents a general paradigm based on unpaired image-to-image translation to align the neural network’s embeddings extracted from images from different domains. Each training image is translated from its source domain to the other, thus building a latent space in which the distribution of image features is domain invariant, and semantically similar images are projected into close vicinity. This paradigm is extended by CrDoCo [10], which combines massive image-to-image translation using GANs with adversarial learning and pixel-wise consistency.

Another promising solution is to perform adversarial learning on the neural network’s output space. The rationale behind this is that the spatial distribution of objects in the images from the training and evaluation datasets should be similar, even if their visual features differ; an example of this is the fact that tables and chairs are usually located close to each other. Consequently, when correctly identified, the spatial distribution of labels coming from different domains should be similar. Following this, the authors of [11] propose to train a CNN to fool a discriminator tasked to disambiguate the domain from which its output masks are generated. An extension of this approach is ADVENT [12], which employs an adversarial training scheme to align the entropy distribution of the target images (computed as the pixel-wise categorical probability distribution) with that of the source domain, under the assumption that the entropy is usually low on data from the source domain and high on that from the target one.

The adversarial learning paradigm has some limitations: it is inherently unstable, it focuses only on domain confusion, and it completely ignores domain-specific cues. To overcome these issues, recent works leverage self-supervision with pseudo-labels inferred by the model itself from the target data. However, pseudo-labels from images in the target domain are noisy and inaccurate, and they cannot be directly used to finetune the source model. A promising solution to this is to balance the influence of pseudo-labels during training, promoting high-confident predictions while penalizing noisy annotations [13]. The confidence is estimated as the divergence between the outputs from two independent prediction layers. This uncertainty is then used as a regularization term during fine-tuning. An alternative approach is that of [14], which estimates the confidence of pseudo-labels by using the relative distance of the extracted features from their respective class centroid.

Another general approach is to use a student network, tasked with adapting to the target domain, supervised by a teacher network that, being gradually updated using an

exponential moving average of the student, provides more stable and consistent pseudo-labels. Following this, the work in [15] rectifies the pseudo-labels using feature clustering on pixel embeddings from the target domain. The label of each pixel is updated based on the assigned cluster, thus providing more coherent and stable annotations that are used to supervise the student training. A similar approach is to use Masked Image Consistency [16], in which the student has access to images in which random patches are masked, and it should be consistent with the teacher network, which uses unmasked ones. An extension of the teacher-student paradigm is presented in [17]. In such a work, a student network is supervised by two teachers: the first one is applied to the outputs while the second, bigger teacher is fixed during training and provides transfer knowledge in the feature space.

Leveraging robot embodiment for UDA. One major limitation of the previously presented approaches is that they operate by processing each frame independently, thus neglecting possible relations with previous and past perceptions. This ignores the fact that indoor mobile robots, while navigating, observe several similar views from the same scene; thus, the images in the sequence are highly correlated. This fact can be leveraged as a natural proxy for consistency, which can be exploited to improve pseudo-label quality, reduce uncertainty, and enhance the recognition of small, occluded, or partially visible objects.

An early work is [6], which introduces the concept of multi-view consistency, leveraging the fact that model’s predictions about the same locations of the environment must exhibit some level of similarity even if taken from multiple viewpoints. The method aggregates single-frame 2D semantic predictions of the robot into a 3D voxel map, which stores the probability for each voxel to belong to a certain semantic class. This dense representation of the environment is then used for rendering pseudo-labels that aggregate information from multiple viewpoints, thus removing the intrinsic noise of the frame-by-frame model’s predictions. The approach is extended in [7], where the authors propose to leverage the robot embodiment for data acquisition in specific environmental locations. The idea is to aggregate in the dense environmental map not only the semantic labels, but also the uncertainty estimation of the model’s predictions. In this way, the robot can collect novel views focusing on areas with high uncertainty to improve the quality of the 3D annotations. In the recent work of [18], the voxel map is substituted with a semantic NeRF, a more compact spatial representation of the environment that allows the rendering of novel views (RGB images) and their respective pseudo-labels to augment the training data.

While these methods present promising results, they still often struggle to resolve coherent misclassifications that persist across multiple frames, as we detail in Section III-C. We advance multi-view consistency by addressing this limitation. We integrate a zero-shot, instance-aware refinement phase after 3D map rendering that enforces single semantic classes per object instance while removing rendering artifacts.

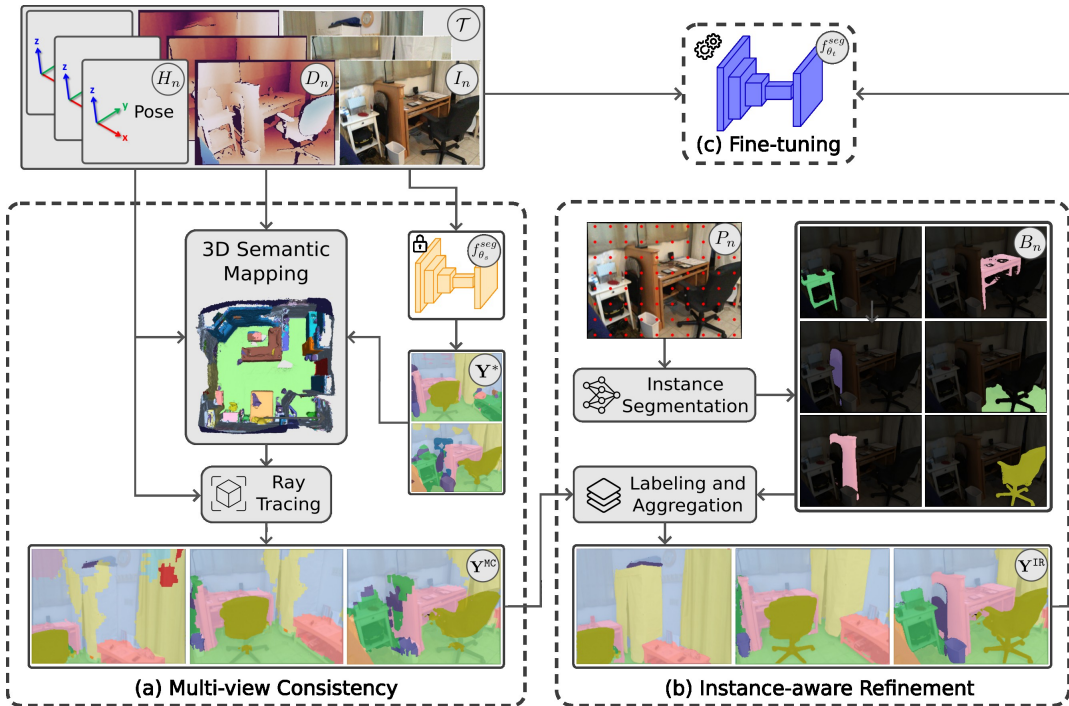


Fig. 2. Our method (a) first aggregates the source model’s predictions \mathbf{Y}^* into a 3D map to obtain multi-view consistent pseudo-labels \mathbf{Y}^{MC} . (b) These are then refined with a foundation model for instance segmentation. (c) The resulting pseudo-labels \mathbf{Y}^{IR} are finally used to fine-tune the source model.

III. METHOD

A. Problem formulation

We consider a reference scenario in which a mobile robot relies on a neural network to perform semantic segmentation from images acquired with an on-board RGB camera, while depth information is obtained either from stereo or an RGB-D sensor. Formally, we describe the task with a function $Y^* = f_{\theta_s}^{\text{seg}}(I)$ that predicts from an RGB image I a semantic mask Y^* , where each pixel is assigned to an object category $o \in \mathcal{O}$. The network’s parameters θ_s are pre-trained on a *source* dataset $\mathcal{S} = \{\mathbf{I}^s, \mathbf{Y}^s\}$, composed of a set of images (\mathbf{I}^s) and their corresponding ground-truth semantic annotations (\mathbf{Y}^s) coming from multiple indoor scenes. We aim to mitigate the domain shift, that is the performance degradation that occurs when a robot’s perception system encounters visual conditions, layouts, or object appearances that differ from its training data, experienced by the robot when deployed in a novel and previously unseen *target* environment t , where no ground-truth semantic labels are available. From that environment, we suppose to have a sequence of N perceptions $\mathcal{T} = \{\mathbf{I}^t, \mathbf{D}^t, \mathbf{H}^t\}$ acquired by the robot from multiple viewpoints, where each colored image $I_n \in \mathbf{I}^t$ and its corresponding depth $D_n \in \mathbf{D}^t$ are localized using the camera extrinsics $H_n \in \mathbf{H}^t$ (with $n \in \{1, \dots, N\}$). Our goal is to adapt $f_{\theta_s}^{\text{seg}}$ for the target environment t with no ground-truth and obtain a new semantic network $f_{\theta_t}^{\text{seg}}$ with improved performance in t .

Fig. 2 presents a general overview of our proposed approach, which performs a self-supervised fine-tuning of $f_{\theta_s}^{\text{seg}}$ on data from the target scene t using directly its predictions as pseudo-labels. Our method aims at improving

the quality of these pseudo-labels to prevent the performance degradation typically caused by re-training with the raw model’s outputs, which are particularly noisy due to domain shifts. Following [6], we start by aggregating the per-frame predictions of $f_{\theta_s}^{\text{seg}}$ in a 3D environmental representation, which is used for rendering multi-view consistent pseudo-labels (Section III-B, Fig. 2 (a)). Then, we add a novel instance refinement step to ensure instance-coherent semantic annotations and fix the artifacts produced by the rendering procedure (Section III-C, Fig. 2 (b)). Finally, the refined pseudo-labels are used to fine-tune $f_{\theta_s}^{\text{seg}}$, obtaining $f_{\theta_t}^{\text{seg}}$ with improved performance in t (Fig. 2 (c)).

B. Multi-view consistent pseudo-labels

To retrieve multi-view consistent pseudo-labels we follow the method described in [6]. From the sequence \mathcal{T} we extract the set $\mathbf{Y}^* = \bigcup_{n=1}^N \{Y_n^* | Y_n^* = f_{\theta_s}^{\text{seg}}(I_n)\}$ containing the single-frame model’s prediction for each image in \mathbf{I}^t . Then, we use Kimera Semantics [8] to aggregate \mathcal{T} together with \mathbf{Y}^* in a 3D semantically-annotated representation of the environment, in which the geometry is represented by a TSDF volume [19] and the semantics are stored in an additional voxel layer. Each voxel maintains a distribution probability over the object categories, which is updated every frame n by matching the 3D semantic points obtained by projecting the semantic pixels of Y_n^* using the depth D_n and the camera pose H_n . After integrating all measurements, we export a high-resolution mesh and ray-trace it from each camera pose H_n to find the first 3D intersection per pixel. The corresponding voxel provides the most likely class, yielding the set \mathbf{Y}^{MC} of multi-view consistent pseudo-labels for all $H_n \in \mathbf{H}^t$ (see the example reported in Fig. 2 (a)).

C. Instance-aware refinement

Despite the rendered pseudo-labels \mathbf{Y}^{MC} show improved quality with respect to single-frame predictions (as they are persistent in consecutive viewpoints), they have some limitations that motivate our additional refinement step. This is a result of how robots perceive the environment while performing navigation. Robots often perceive objects only partially during navigation, leading to systematic errors in label aggregation. Take as an example a robot observing a cabinet from multiple views. When first observing it, the robot may only capture small corners or partially occluded views, which are difficult to recognize and thus are frequently misclassified as other categories, such as `wall`, as happens in the examples of Fig. 2 (a). Because these partial views are encountered repeatedly, their incorrect labels end up dominating in the voxel-wise majority of the resulting 3D map. Later, once the cabinet is fully visible, the robot correctly classifies it, but those few correct predictions remain outnumbered. As a result, the final aggregated label is wrong, despite the robot having seen the cabinet clearly. Our instance-aware refinement step resolves this issue by leveraging SAM to group pixels into coherent object instances. Within each instance, the correct labels from full observations can outweigh the scattered errors from partial ones, allowing our method to propagate the right semantic category to the entire object, as shown in Fig. 2 (b). Additionally, the multi-view consistent pseudo-labels are affected by artifacts introduced by the 3D reconstruction (such as missing depth in reflective surfaces) and by the rendering process (e.g., inaccurate boundaries caused by the voxel discretization).

To address these limitations, we propose to further refine the pseudo-labels of \mathbf{Y}^{MC} according to the object instances contained in their corresponding RGB images in \mathbf{I}^t . To do this, our method leverages the zero-shot capabilities of Segment Anything (SAM) [20], [21], a prompt-based foundation model for instance segmentation. From an RGB image I_n and a list of J prompts P_n (in which each prompt $p_{n,j} \in P_n$ is expressed with a point and/or a bounding box), we use SAM to obtain a set of binary masks

$$B_n = \bigcup_{j=1}^J \{b_{n,j} | b_{n,j} = f_{\theta}^{\text{sam}}(I_n, p_{n,j})\},$$

where f_{θ}^{sam} denotes the SAM model with pre-trained parameters θ . Inside each $b_{n,j}$, a pixel u assumes the value $b_{n,j}(u) = 1$ if it lies inside the instance identified by SAM from the relative prompt $p_{n,j}$, otherwise $b_{n,j}(u) = 0$.

For prompting SAM, we propose two alternative strategies, which we later compare. The first one, named `grid`, uses a fixed prompt, which is created once and used for every frame. In this strategy, P_n is defined as a list of points uniformly arranged in a grid structure on the image plane, where d specifies the distance between adjacent points in the horizontal and vertical axes. In the second one, called `informed`, the prompt is procedurally extracted for each frame using its rendered pseudo-label. At first, a pseudo-label Y_n^{MC} is partitioned into clusters, where each cluster is a

connected region of pixels sharing the same semantic class. Then, for each cluster larger than a percentage a of the image area, it calculates the bounding box and its centroid. P_n is finally composed by aggregating the bounding box/centroid pairs derived from all clusters.

Independently from the prompting strategy, our method proceeds by combining the object instances of B_n with the classes contained in the corresponding pseudo-label Y_n^{MC} . To do this, for each instance mask $b_{n,j} \in B_n$, we extract the most frequent object category according Y_n^{MC} , formally defined as

$$o_{n,j}^* = \arg \max_{o \in \mathcal{O}} \sum_{u \in \Omega} b_{n,j}(u) \mathbb{1}_{\{Y_n^{\text{MC}}(u)=o\}},$$

where $\mathbb{1}_{\{\cdot\}}$ is the indicator function and Ω are the pixels indices. All the binary masks $b_{n,j} \in B_n$ are then merged, together with their relative object category $o_{n,j}^*$, in a single instance-aware pseudo-label Y_n^{IR} . At the beginning, we set $Y_n^{\text{IR}} = Y_n^{\text{MC}}$. Then, for each mask $b_{n,j} \in B_n$, we update Y_n^{IR} by overriding the pixels's category inside the region defined by $b_{n,j}$ with $o_{n,j}^*$, formally

$$Y_n^{\text{IR}}(u) = \begin{cases} o_{n,j}^* & \text{if } b_{n,j}(u) = 1 \\ Y_n^{\text{IR}}(u) & \text{if } b_{n,j}(u) = 0 \end{cases} \quad \forall u \in \Omega.$$

Note that this step is sequentially applied for all $j \in \{1, \dots, J\}$, meaning that previously added semantic labels can, in principle, be replaced in later stages. When B_n is derived using the `informed` prompting strategy, this replacement is not possible: the object instances are disjoint by construction since they are extracted from non-overlapping clusters of labels. In contrast, when the `grid` strategy is used, the order in which the instances $b_{n,j}$ are aggregated might impact on the quality of the pseudo-labels. This is because B_n contains a high number of overlapping instances, meaning that small objects may be completely overwritten by larger ones aggregated later. To avoid this, we sort the B_n in descending order of instance dimension, thus preserving the fine-grained details identified by SAM when prompted with the `grid` strategy. This entire instance-oriented refinement procedure is performed for all the images/pseudo-labels pairs in \mathbf{I}^t and \mathbf{Y}^{MC} . The resulting instance-aware annotations (for all the measurements in \mathcal{T}) are contained in the set \mathbf{Y}^{IR} . Finally, we use refined pseudo-labels \mathbf{Y}^{IR} to finetune the model, thus doing UDA, as shown in Fig. 2 (c).

IV. EVALUATION

A. Experimental setting

Dataset For the evaluation of our proposed method, we rely on ScanNet [4], a real-world dataset containing 1513 sequences acquired in 707 different indoor environments. Each sequence includes RGB-D image pairs with their relative camera poses, estimated using BundleFusion [22]. The dataset also provides the per-frame ground-truths, which are manually annotated with the NYU40 object categories [23]. In all the experiments, we rescale the images to a resolution

of 320×240 pixels, to be compliant with the hardware setup of a low-powered mobile robot.

Network pre-training. Similarly to [6], [18], the scenes from 11 to 707 represent our source domain and are used to initialize the weights of the neural network for semantic segmentation $f_{\theta_s}^{seg}$. For simplicity, we rely on the DeepLabV3 [24] (with a ResNet-101 [25] as backbone) provided by [18], which was pre-trained on a dataset \mathcal{S} counting $\approx 25k$ images/ground-truths pairs from environments 11–707, sampling one frame every 100. \mathcal{S} was randomly split in $\approx 20k$ and $\approx 5k$ samples for training and validation, respectively. The network is optimized using Adam for 150 epochs, with a batch size of 4. To measure the quality of the predictions, we use the Intersection over Union averaged over the objects categories (mIoU).

Network adaptation. Scenes 1–10, that represent the target environments, are used to obtain $f_{\theta_t}^{seg}$ by fine-tuning $f_{\theta_s}^{seg}$ with the pseudo-labels provided by our method. We consider the first sequence for each environment, from which the first 80% of the frames are used to produce the pseudo-labels and fine-tuning $f_{\theta_s}^{seg}$, while the last 20% is used for testing. This is the same method adopted in [6] and [18], which we replicate but also extend to achieve a more representative evaluation. Although training and testing frames are temporally disjoint, they are sampled from the same sequence within a given environment. In contrast, real-world deployments typically involve two distinct runs: one to collect data for unsupervised adaptation and another to evaluate performance in a different trajectory. To better reflect this setting, we perform an additional evaluation in which two separate sequences are used for each environment: one for pseudo-label generation and fine-tuning, and the other for testing. (We cannot run this protocol on environments 5, 8, and 9, as only a single sequence is available in the dataset.) Overall, this experimental design aims at better reflecting the practical conditions of robotic operation: a model $f_{\theta_s}^{seg}$ is first pre-trained on a large source dataset, the robot then collects an initial set of observations (e.g., while mapping the environment) to generate pseudo-labels and adapt $f_{\theta_s}^{seg}$ in an unsupervised fashion. Finally, the adapted model $f_{\theta_t}^{seg}$ is tested on new trajectories.

Instance-aware refinement. To generate the multi-view consistent pseudo-labels \mathbf{Y}^{MC} , we leveraged the open-source implementation of [6]. We set the voxel size of Kimera Semantics [8] to 3cm and 5cm, to evaluate the performance under different resolutions, following the works of [6], [18]. More precisely, setting the voxel size to 3cm (5cm) creates a more (less) detailed 3D representation, requiring higher (lower) processing power for real-time robot deployment. The object instances B_n are obtained using the Ultralytics’ implementation of Segment Anything v2 [21]. For prompting SAM, the points of the `grid` strategy are generated with a distance $d = 32$, while the cluster of labels smaller than $a = 0.1\%$ of the image size are not considered in the `informed` approach. To evaluate the performance improvements related to the two prompting strategies, we compare three versions of $f_{\theta_t}^{seg}$ for each environment: (i) using only `grid`, (ii)

TABLE I
PSEUDO-LABELS PERFORMANCE

Env	\mathbf{Y}^*	Baseline (5cm)		Our method			
		\mathbf{Y}^{MC}	$\Delta \mathbf{Y}^*$	\mathbf{Y}_G^{IR}	$\Delta \mathbf{Y}^{MC}$	\mathbf{Y}_I^{IR}	$\Delta \mathbf{Y}^{MC}$
1	41,2	48,1	16,7%	52,3	8,7%	49,1	2,1%
2	34,8	28,8	-17,2%	32,5	12,8%	31,5	9,4%
3	23,7	26	9,7%	25,4	-2,3%	26,5	1,9%
4	62,8	63,9	1,8%	65,2	2,0%	64,7	1,3%
5	49,8	42,6	-14,5%	45,5	6,8%	44,9	5,4%
6	48,7	49,3	1,2%	<u>52,6</u>	6,7%	53,3	8,1%
7	40,3	48,4	20,1%	48,8	0,8%	50,8	5,0%
8	31,4	34,8	10,8%	35,4	1,7%	36,4	4,6%
9	31,8	32,8	3,1%	30,8	-6,1%	32,1	-2,1%
10	52,1	55,8	7,1%	59,6	6,8%	<u>58,8</u>	5,4%
Avg	41,7	43,1	3,9%	<u>44,8</u>	3,8%	44,8	4,1%

mIoU (best and second best values) of the pseudo-labels from our method (\mathbf{Y}_G^{IR} and \mathbf{Y}_I^{IR}) against the baseline (\mathbf{Y}^{MC}) and the pre-trained model (\mathbf{Y}^*).

using only `informed`, and (iii) combining the two. For the single-method settings, $f_{\theta_s}^{seg}$ is optimized using \mathbf{Y}^{IR} and their corresponding RGB images for 10 epochs using Adam with a learning rate of 10^{-5} , and batch size of 4. When the methods are combined, the model processes the same RGB image twice inside the same epoch, each time with different pseudo-labels. To maintain a constant training budget, epochs are halved from 10 to 5. Random flipping and color jitter are used as data augmentation.

Baseline. The baseline we use to evaluate our approach is the multi-view consistency [6], that is also an important building block of our pipeline, as detailed in Section III-B. For a fair comparison, we replicate the results by re-generating the pseudo-labels from scratch using the aforementioned pre-trained DeepLabV3. With the multi-view consistent annotations \mathbf{Y}^{MC} and their relative RGB images, we fine-tune a model $f_{\theta_t}^{seg}$ for each target environment, using the Adam optimizer with a learning rate of 10^{-5} and batch size of 4 for 10 epochs.

B. Pseudo-labels improvement

This section evaluates the quality of the instance-aware pseudo-labels produced by our method compared to those of the baseline. The quantitative results are reported in Table I, which shows the precision of the pseudo-labels (measured in mIoU) obtained from the first 80% of the first sequence for each environment, using a voxel size of 5cm. We refer to the refined pseudo-labels as \mathbf{Y}_G^{IR} or \mathbf{Y}_I^{IR} if obtained using the `grid` or the `informed` prompting strategy.

The metrics show that our instance-refinement step improves the multi-view consistent annotations \mathbf{Y}^{MC} of the baseline, on average by 3.8% and 4.1% for `grid` and `informed`, respectively. Despite the baseline improves the raw model’s predictions \mathbf{Y}^* by 3,9% on average, its impact is not stable and changes according to the environment. In some cases, the multi-view consistency yields substantial improvements (such as in environments 1 and 7 with an increment of 16,7% and 20.1%), while in others it strongly degrades the quality of the per-frame predictions (e.g., in environments 2 and 5 with a decrease of -17,2% and -14,5%). In contrast, our instance-aware refinement step increases the quality of \mathbf{Y}^{MC} in all environments, with only marginal decreases observed in scenes 3 and 9. As shown in

TABLE II

PERFORMANCE OF THE ADAPTED MODEL WHEN TRAINING – TESTING USE THE 80% – 20% OF THE SAME SEQUENCE OF EACH ENVIRONMENT

Env	PT	Baseline (5cm)				Our method					Baseline (3cm)				Our method			
		MC	Δ_{PT}	IR _G	Δ_{MC}	IR _I	Δ_{MC}	IR _{G_I}	Δ_{MC}		PT	MC	Δ_{PT}	IR _G	Δ_{MC}	IR _I	Δ_{MC}	IR _{G_I}
1	44,2	48,2	9,0%	51	5,8%	49,3	2,3%	50,9	5,6%	44,2	47,6	7,7%	50,1	5,3%	50,4	5,9%	51	7,1%
2	36,2	31,3	-13,5%	33,6	7,3%	34,4	9,9%	34,8	11,2%	36,2	34,4	-5,0%	34,5	0,3%	34,9	1,5%	33,8	-1,7%
3	22,9	21	-8,3%	21,4	1,9%	23,2	10,5%	22,9	9,0%	22,9	21,7	-5,2%	22,6	4,1%	23,1	6,5%	23,2	6,9%
4	50,1	51,1	2,0%	52,5	2,7%	52,5	2,7%	53,4	4,5%	50,1	52,4	4,6%	52,9	1,0%	49,5	-5,5%	53,7	2,5%
5	39,7	40,1	1,0%	44,5	11,0%	44,4	10,7%	44,7	11,5%	39,7	45	13,4%	47,2	4,9%	46,1	2,4%	47,3	5,1%
6	35	31,9	-8,9%	33,9	6,3%	34,1	6,9%	34,9	9,4%	35	33,5	-4,3%	39,5	17,9%	38,8	15,8%	40,1	19,7%
7	56,7	63,5	12,0%	60,7	-4,4%	63,5	0,0%	60,9	-4,1%	56,7	63	11,1%	60,9	-3,3%	56,7	-10,0%	59,5	-5,6%
8	29,5	25,1	-14,9%	25,2	0,4%	26,3	4,8%	27,6	10,0%	29,5	25,1	-14,9%	24,1	-4,0%	27	7,6%	25,7	2,4%
9	55,9	65,3	16,8%	70,2	7,5%	69	5,7%	69,2	6,0%	55,9	66,9	19,7%	70,9	6,0%	69,8	4,3%	69,9	4,5%
10	73,4	71,9	-2,0%	75,3	4,7%	74,7	3,9%	75	4,3%	73,4	73,2	-0,3%	75,7	3,4%	75,6	3,3%	75,3	2,9%
Avg	44,4	44,9	-0,7%	46,8	4,3%	46,8	5,7%	47,4	6,7%	44,4	46,3	2,7%	47,8	3,6%	47,2	3,2%	48,0	4,4%

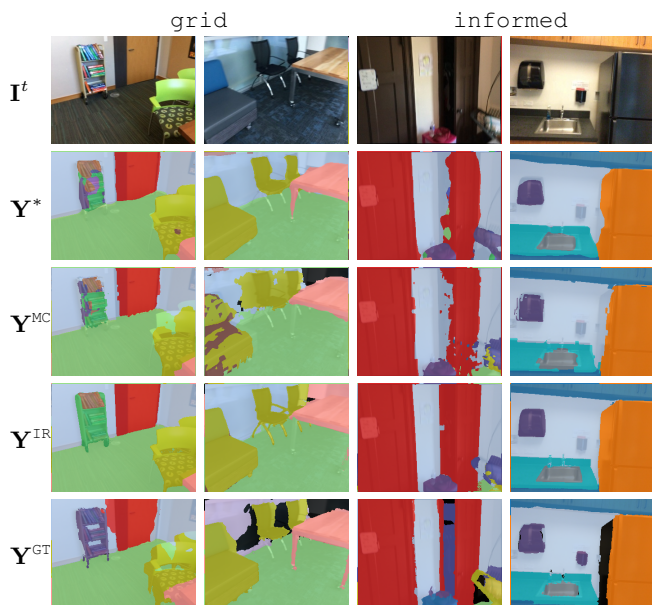
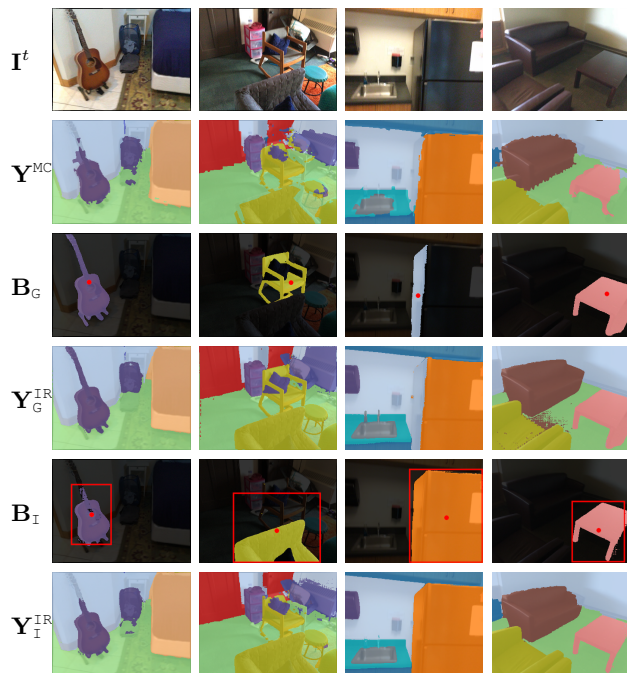
mIoU (best and second best values) of $f_{\theta_t}^{seg}$ trained with our method (IR_G, IR_I, and IR_{G_I}) against the baseline (MC), and the pre-trained model (PT).Fig. 3. Instance-aware pseudo-labels (Y^{IR}) generated with *grid* (left) and *informed* (right), compared to multi-view consistent (Y^{MC}) and raw (Y^*) ones. I^t and Y^{GT} denote the RGB images and ground truths.

Fig. 3, our approach substantially improves the precision of the pseudo-labels by propagating the object classes according to instances identified in the RGB perception, and fixes the rendering artifacts produced by the ray tracing.

Another interesting outcome from the results of Table I is that, although the two prompting strategies achieve the same average precision, their performances vary across environments. This can be explained by their markedly different behaviors, which are, to some extent, complementary.

The *grid* approach exploits all the foundational model’s strengths in instance segmentation by using a fixed grid of points for each frame. This setting makes this method particularly effective in challenging situations where the objects have a complex shape or the multi-view consistent pseudo-labels are particularly noisy. Examples are reported in Fig. 4, respectively in the first and second columns, where we can see that the *grid* approach precisely identifies the instances of a guitar and a chair, which is partially occluded by a pillow. Compared to the automatic approach, the *informed* prompts derived from Y^{MC} produce less

Fig. 4. Comparison of pseudo-labels from our method using *grid* (Y_G^{IR}) and *informed* (Y_I^{IR}) prompts, with representative SAM instances and their prompts (B_G , B_I). I^t and Y^{MC} denote the RGB images and multi-view annotations. The first two columns highlight the strengths of *grid*, while the last two show the improved precision of *informed*.

precise object instances. In the first case, the bounding box cuts the guitar’s neck, while in the second one, the couch and the chair are grouped in the same prompt (because they form a single cluster), but only the couch is segmented.

In contrast, *informed* is a more conservative approach because the prompts are extracted from the multi-view consistent annotations. This is particularly useful to prevent the over-segmentation, a well-known limitation of SAM that often splits a single object into multiple instances, especially when the prompt is particularly dense (as in the case of *grid*). See, for example, the third column of Fig. 4, where *informed* identifies the fridge as a unique object, while *grid* considers its left side as a separate instance, which is then wrongly labeled. Another important effect of *informed* is to guide the predictions of SAM when objects have similar textures, a situation in which an automatic segmentation is particularly challenging. This is illustrated

TABLE III

PERFORMANCE OF THE ADAPTED MODEL WHERE TRAINING AND TESTING USE DIFFERENT SEQUENCES OF THE SAME ENVIRONMENT

Env	PT	Baseline (5cm)			Our method					PT	Baseline (3cm)			Our method				
		MC	Δ_{PT}	IR_G	Δ_{MC}	IR_I	Δ_{MC}	IR_{GI}	Δ_{MC}		MC	Δ_{PT}	IR_G	Δ_{MC}	IR_I	Δ_{MC}	IR_{GI}	Δ_{MC}
1	47,1	49,6	5,3%	55	10,9%	51,2	3,2%	54,3	9,5%	47,1	48,9	3,8%	50,6	3,5%	49,5	1,2%	50,8	3,9%
2	44,1	35,2	-20,2%	39,6	12,5%	39,6	12,5%	39,3	11,6%	44,1	42,1	-4,5%	41,8	-0,7%	41,6	-1,2%	41,6	-1,2%
3	30,6	30,5	-0,3%	32,8	7,5%	34,4	12,8%	34,7	13,8%	30,6	31,5	2,9%	36,5	15,9%	35,1	11,4%	36,3	15,2%
4	55,3	55,8	0,9%	57,2	2,5%	56,2	0,7%	57,7	3,4%	55,3	57,1	3,3%	57,4	0,5%	56,8	-0,5%	57,2	0,2%
6	49,2	50,5	2,6%	52,8	4,6%	54,7	8,3%	53	5,0%	49,2	51,2	4,1%	54,8	7,0%	55,8	9,0%	54,8	7,0%
7	52	54,4	4,6%	55,5	2%	57,2	5,1%	56,2	3,3%	52	54,1	4%	54,8	1,3%	55,1	1,8%	55,3	2,2%
10	57,5	63,2	9,9%	68,7	8,7%	67,6	7,0%	68,7	8,7%	57,5	65,2	13,4%	69,1	6,0%	68,3	4,8%	70,1	7,5%
Avg	48	48,5	0,4%	51,7	7%	51,6	7,1%	52,0	7,9%	48	50	3,9%	52,1	4,8%	51,7	3,8%	52,3	5,0%

mIoU (best and second best values) of $f_{\theta_t}^{seg}$ trained with our method (IR_G , IR_I , and IR_{GI}) against the baseline (MC), and the pre-trained model (PT).

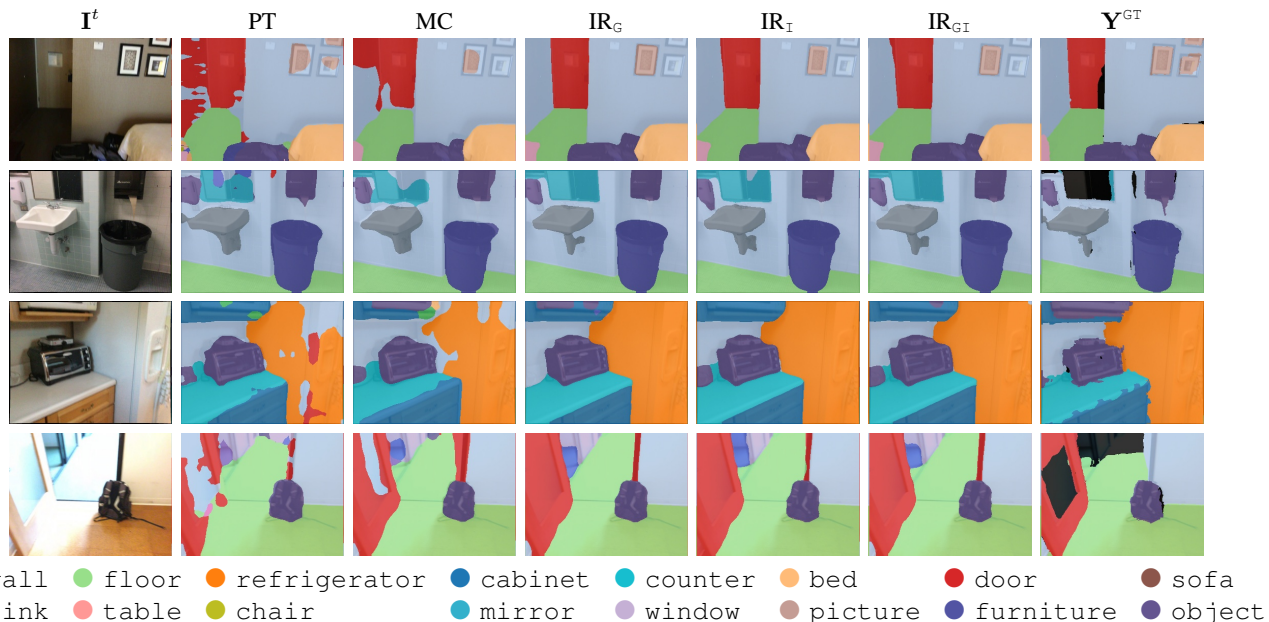


Fig. 5. Segmentation improvements with different prompting strategies (IR_G , IR_I , IR_{GI}) compared to the baselines (MC and PT). The first two rows use a 5cm voxel size, the last two a 3cm. Black areas in Y^{GT} are due to rendering errors or missing labels; these segmentation errors are improved by our IR.

in the last column of Fig. 4, where *informed* identifies improved and more precise instances than *grid*, which partially mixes the table and the couches with the floor.

C. Performance of the adapted model

In this section, we present the experimental results evaluating the performance of our method on the semantic segmentation task. Table II shows the results obtained when fine-tuning and testing inside the same sequence, using the first 80% and the last 20% of the frames, respectively (as done in [6] and [18]). We run all the experiments setting the voxel size of Kimera Semantics to 3cm and 5cm.

The results demonstrate that our instance-aware refinement remarkably increases the model's performance with respect to the multi-view consistency baseline of [6] (MC), and the raw predictions of the pre-trained $f_{\theta_s}^{seg}$ (PT). Considering a voxel size of 5cm, our method improves the baseline MC by 4,3% and 5,7% using the *grid* (IR_G) and *informed* (IR_I) prompt strategies. As observed in the previous section, the impact of the multi-view consistency on the pre-trained model PT varies a lot across environments, registering remarkable improvements (e.g., in environments 1, 7, and 9), but also important degradations (such as in environments 2,

6, and 8). Starting from this, our method improves the performance of MC in both cases: the instance-aware refinement step not only increases the segmentation accuracy when MC is effective, but also mitigates and fixes errors of the multi-view consistency in case of failure. See, for example, scenes 2 and 6: while MC degrades PT by $\approx -14\%$ and -9% , our method IR_I substantially improves the quality of the baseline $\approx 10\%$ and 7% , obtaining segmentation accuracies close to PT. Table II also shows the benefits of fine-tuning the source model $f_{\theta_s}^{seg}$ combining the pseudo-labels obtained with the *grid* and *informed* prompting strategies (IR_{GI}). Mixing the two different types of instance-aware annotations during training mitigates the negative intrinsic effects of *grid* and *informed* (previously described in Section IV-B), helping the model in providing more accurate outputs. More specifically, IR_{GI} reaches an average increment of 6,7% over the baseline MC, achieving the best or the second best performance across all environments (except for scene 7 where IR_I and MC perform better). All these outcomes are confirmed using a voxel size of 3cm. As expected, the quality of the baseline MC is improved from 44,9 to 46,3 mIoU. This is because a smaller voxel size limits the discretization of the world encoded in the 3D map, thus reducing the

artifacts produced by the ray-tracing procedure. Despite this, our method increases the performance of the baseline MC by 3,6% and 3,2% with IR_G and IR_I , respectively. Again, the best performances are obtained by mixing the prompting strategies (IR_{GI}) during training, with an average increment of 4,4% over MC. Interestingly, while the baseline MC is strongly influenced by the voxel size of Kimera Semantics, our method achieves comparable performance regardless of the voxels' resolution (0.6 mIoU difference obtained IR_{GI} with 3cm and 5cm). This is particularly useful in practical deployment scenarios, where the voxel dimension can be set to a higher value, thus reducing the computational overhead of the robot in integrating the frames in Kimera Semantics.

While the aforementioned results of Table II are collected in a subportion of the environment, Table III reports the performance of $f_{\theta_t}^{seg}$ considering the entire environment, as training and testing are performed in different sequences. The results of this additional experiment further strengthen the findings of the previous evaluation. Using a voxel size of 5cm (3cm), our method outperforms the MC baseline on average by $\approx 7\%$ (4%), obtaining a remarkable performance increment in all scenes. The only case where the pre-trained model (PT) performs worse after adaptation is in scene 2, which is a problematic environment since almost the 30% of the frames have missing positions and a large portion of the scene is not annotated. Despite this, our instance refinement strongly increase the baseline's performance of $\approx 12\%$ with respect to MC. Again, the best averaging performances are obtained by IR_{GI} , which reaches almost the same performance using 5cm and 3cm as voxel sizes, respectively 52 and 52,3 mIoU. The improvements provided by our method can be observed in the qualitative examples of Fig. 5. Despite the baseline (MC) effectively removes the noise of the raw model's predictions (PT), the $f_{\theta_t}^{seg}$ learns the errors that are persistent errors between multiple frames, like the doors in the first and last rows that are confused with the wall. In contrast, our method remarkably increases the quality of the segmentation masks produced by $f_{\theta_t}^{seg}$, solving also some artifacts of the ground-truth. This can be seen by observing the precision of the masks related to the sink and the furniture in the second and third rows.

V. CONCLUSIONS

This paper presents a novel approach for unsupervised domain adaptation in robotics applications based on self-supervision with instance-aware pseudo-labels. In future investigations, we plan to embed the instance segmentation inside the 3D map and leverage foundation models to fix the semantic class of object instances.

ACKNOWLEDGEMENTS

Part of this project was funded by National Plan for NRRP Complementary Investments (PNC) in the call for the funding of research initiatives for technologies and innovative trajectories in the health—project n. PNC0000003—AdvaNced Technologies for Human-centred Medicine (project acronym: ANTHEM).

REFERENCES

- [1] S. Garg, N. Sünderhauf, F. Dayoub, D. Morrison *et al.*, "Semantics for robotic mapping, perception and interaction: A survey," *Foundations and Trends® in Robotics*, vol. 8, no. 1–2, pp. 1–224, 2020.
- [2] G. Wilson and D. J. Cook, "A survey of unsupervised deep domain adaptation," *ACM TRANS INTELL SYST TECHNOL*, vol. 11, no. 5, pp. 1–46, 2020.
- [3] M. Antonazzi, M. Luperto, N. A. Borghese, and N. Basilico, "Development and adaptation of robotic vision in the real-world: the challenge of door detection," *J FIELD ROB*, 2025.
- [4] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Niessner, "Scannet: Richly-annotated 3d reconstructions of indoor scenes," in *Proc. of CVPR*, July 2017.
- [5] M. Schwonberg, J. Niemeijer *et al.*, "Survey on unsupervised domain adaptation for semantic segmentation for visual perception in automated driving," *IEEE Access*, vol. 11, pp. 54 296–54 336, 2023.
- [6] J. Frey, H. Blum, F. Milano, R. Siegwart, and C. Cadena, "Continual adaptation of semantic segmentation using complementary 2d-3d data representations," *IEEE ROBOT AUTOM LETT*, 2022.
- [7] R. Zurbrügg, H. Blum, C. Cadena, R. Siegwart, and L. Schmid, "Embodied active domain adaptation for semantic segmentation via informative path planning," *IEEE ROBOT AUTOM LETT*, vol. 7, no. 4, pp. 8691–8698, 2022.
- [8] A. Rosinol, M. Abate, Y. Chang, and L. Carlone, "Kimera: an open-source library for real-time metric-semantic localization and mapping," in *Proc. of ICRA*, 2020, pp. 1689–1696.
- [9] Z. Murez, S. Kolouri, D. Kriegman, R. Ramamoorthi, and K. Kim, "Image to image translation for domain adaptation," in *Proc. of CVPR*, June 2018.
- [10] Y.-C. Chen, Y.-Y. Lin, M.-H. Yang, and J.-B. Huang, "Crdoco: Pixel-level domain transfer with cross-domain consistency," in *Proc. of CVPR*, June 2019.
- [11] Y.-H. Tsai, W.-C. Hung, S. Schulter, K. Sohn, M.-H. Yang, and M. Chandraker, "Learning to adapt structured output space for semantic segmentation," in *Proc. of CVPR*, June 2018.
- [12] T.-H. Vu, H. Jain, M. Bucher, M. Cord, and P. Perez, "Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation," in *Proc. of CVPR*, June 2019.
- [13] Z. Zheng and Y. Yang, "Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation," *INT J COMPUT VIS*, vol. 129, no. 4, pp. 1106–1120, 2021.
- [14] P. Zhang, B. Zhang, T. Zhang *et al.*, "Prototypical pseudo label denoising and target structure learning for domain adaptive semantic segmentation," in *Proc. of CVPR*, June 2021.
- [15] R. Li, S. Li, C. He, Y. Zhang, X. Jia, and L. Zhang, "Class-balanced pixel-level self-labeling for domain adaptive semantic segmentation," in *Proc. of CVPR*, June 2022, pp. 11 593–11 603.
- [16] L. Hoyer, D. Dai, H. Wang, and L. Van Gool, "Mic: Masked image consistency for context-enhanced domain adaptation," in *Proc. of CVPR*, June 2023, pp. 11 721–11 732.
- [17] S. Jeong, J. Kim, S. Kim, and D. Min, "Revisiting domain-adaptive semantic segmentation via knowledge distillation," *IEEE T IMAGE PROCESS*, vol. 33, pp. 6761–6773, 2024.
- [18] Z. Liu, F. Milano, J. Frey, R. Siegwart, H. Blum, and C. Cadena, "Unsupervised continual semantic adaptation through neural rendering," in *Proc. of CVPR*, June 2023, pp. 3031–3040.
- [19] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning," in *Proc. of IROS*, 2017, pp. 1366–1373.
- [20] A. Kirillov, E. Mintun *et al.*, "Segment anything," in *Proc. of ICCV*, October 2023, pp. 4015–4026.
- [21] N. Ravi, V. Gabeur, Y.-T. Hu, R. Hu, C. Ryal, T. Ma, H. Khedr, R. Rädle, C. Rolland, L. Gustafson, E. Mintun, J. Pan, K. V. Alwala, N. Carion, C.-Y. Wu, R. Girshick, P. Dollár, and C. Feichtenhofer, "Sam 2: Segment anything in images and videos," 2024.
- [22] A. Dai, M. Nießner, M. Zollhöfer, S. Izadi, and C. Theobalt, "Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface reintegration," *ACM T GRAPHIC*, vol. 36, no. 4, 2017.
- [23] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus, "Indoor segmentation and support inference from rgb-d images," in *Proc. of ECCV*, 2012, pp. 746–760.
- [24] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, "Rethinking atrous convolution for semantic image segmentation," 2017.
- [25] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. of CVPR*, 2016.