

Efficient 3D Instance Mapping and Localization with Neural Fields

<https://gtangg12.github.io/iML>

George Tang, Krishna Murthy Jatavallabhula, Antonio Torralba
Massachusetts Institute of Technology

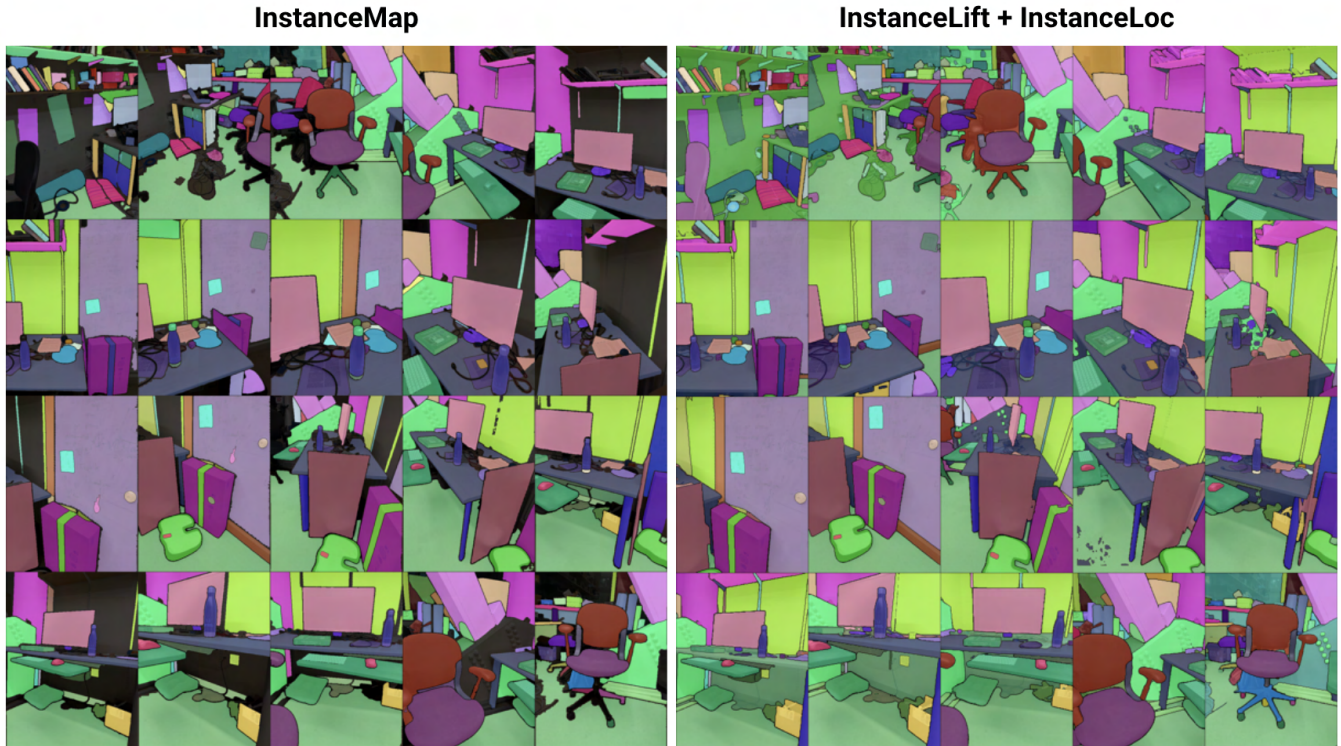


Fig. 1: Our approach, *3DIML*, learns an implicit representation that represents a scene as a composition of object instances via lifting 2D view-inconsistent instance labels from off-the-shelf 2D segmentation models (such as the Segment Anything family of models) into 3D view-consistent instance labels. The images above show results for the in-the-wild scan `postdoc_office` generated using *3DIML*, which is composed of InstanceMap (left) and InstanceLift, then refined with InstanceLoc (right). Each identified 3D label is shown in a different color. Notice how thin and partially occluded objects are accurately delineated across the entire sequence.

Abstract—We tackle the problem of learning an implicit scene representation for 3D instance segmentation from a sequence of posed RGB images. Towards this, we introduce *3DIML*, a novel framework that efficiently learns a *label field* that may be rendered from novel viewpoints to produce view-consistent instance segmentation masks. *3DIML* significantly improves upon training and inference runtimes of existing implicit scene representation based methods. Opposed to prior art that optimizes a neural field in a self-supervised manner, requiring complicated training procedures and loss function design, *3DIML* leverages a two-phase process. The first phase, InstanceMap, takes as input 2D segmentation masks of the image sequence generated by a frontend instance segmentation model, and associates corresponding masks across images to 3D labels. These almost view-consistent pseudolabel masks are then used in the second phase, InstanceLift, to supervise the training of a neural label field, which interpolates regions missed by InstanceMap and resolves ambiguities. Additionally, we introduce InstanceLoc, which enables near realtime localization of instance masks given a trained label field and an off-the-shelf image segmentation model by fusing outputs from both. We evaluate *3DIML* on sequences from the Replica and ScanNet

datasets and demonstrate *3DIML*'s effectiveness under mild assumptions for the image sequences. We achieve a 14-24 \times speedup over existing implicit scene representation methods with comparable quality, showcasing its potential to facilitate faster and more effective 3D scene understanding.

I. INTRODUCTION

Intelligent agents require scene understanding at the object level to effectively carry out context-specific actions such as navigation and manipulation. While segmenting objects from images has seen remarkable progress with scalable models trained on internet-scale datasets [1], [2], extending such capabilities to the 3D setting remains challenging.

In this work, we tackle the problem of learning a 3D scene representation from posed 2D images that *factorizes* the underlying scene into its set of constituent objects. Existing approaches to tackle this problem have focused on training class-agnostic 3D segmentation models [3], [4], requiring large amounts of annotated 3D data, and operating directly over explicit 3D scene representations (e.g., pointclouds). An

alternate class of approaches [5], [6] has instead proposed to directly *lift* segmentation masks from off-the-shelf instance segmentation models into implicit 3D representations, such as neural radiance fields (NeRF) [7], enabling them to render 3D-consistent instance masks from novel viewpoints.

However, the neural field-based approaches have remained notoriously difficult to optimize, with [5] and [6] taking several hours to optimize for low-to-mid resolution images (e.g., 300×640). In particular, Panoptic Lifting [5] scales cubically with the number of objects in the scene preventing it from being applied to scenes with hundreds of objects, while Contrastive Lifting [6] requires a complicated, multi-stage training procedure, hindering practicality for use in robotics applications.

To this end, we propose *3DIML*, an efficient technique to learn 3D-consistent instance segmentation from posed RGB images. *3DIML* comprises two phases: InstanceMap and InstanceLift. Given view-inconsistent 2D instance masks extracted from the RGB sequence using a frontend instance segmentation model [2], InstanceMap produces a sequence of view-consistent instance masks. To do so, we first associate masks across frames using keypoint matches between similar pairs of images. We then use these potentially noisy associations to supervise a neural label field, InstanceLift, which exploits 3D structure to interpolate missing labels and resolves ambiguities. Unlike prior work, which requires multistage training and additional loss function engineering, we use a single rendering loss for instance label supervision, enabling the training process to converge significantly faster. The total runtime of *3DIML*, including InstanceMap, takes 10-20 minutes, as opposed to 3-6 hours for prior art.

In addition, we devise InstaLoc, a fast localization pipeline that takes in a novel view and localizes all instances segmented in that image (using a fast instance segmentation model [8]) by sparsely querying the label field and fusing the label predictions with extracted image regions. Finally, *3DIML* is extremely modular, and we can easily swap components of our method for more performant ones as they become available.

To summarize, our contributions are:

- An efficient neural field learning approach that factorizes a 3D scene into its constituent objects
- A fast instance localization algorithm that fuses sparse queries to the trained label field with performant image instance segmentation models to generate 3D-consistent instance segmentation masks
- An overall training time improvement of $14\text{-}24\times$ over prior art, over scenes comprising hundreds of objects, on a single GPU (NVIDIA RTX 3090)

II. BACKGROUND

2D segmentation: The prevalence of vision transformer architecture and the increasing scale of image datasets have resulted in a series of state-of-the-art image segmentation models. Panoptic and Contrastive Lifting both lift panoptic segmentation masks produced by Mask2Former [1] to 3D by learning a neural field. Towards open-set segmentation,

segment anything (SAM) [2] achieves unprecedented performance by training on a billion masks over 11 million images. HQ-SAM [9] improves upon SAM for fine-grained masks. FastSAM [8] distills SAM into a CNN architecture and achieves similar performance while being orders of magnitude faster. In this work, we use GroundedSAM [10], [11], which refines SAM to produce object-level, as opposed to part-level segmentation masks.

Neural fields for 3D instance segmentation: NeRFs are implicit scene representations that can accurately encode complex geometry, semantics, and other modalities, as well as resolve viewpoint inconsistent supervision [12]. Panoptic lifting [5] constructs semantics and instances branches on an efficient variant of NeRF, TensorRF [13], utilizing a Hungarian matching loss function to assign learned instance masks to surrogate object IDs given reference view-inconsistent masks. This scales poorly with increasing number of objects (owing to the cubic complexity of Hungarian matching). Contrastive lifting [6] addresses this by instead employing contrastive learning on scene features, with positive and negative relations determined by whether or not they project onto the same mask. In addition, contrastive lifting requires a slow-fast clustering-based loss for stable training, leading to faster performance than panoptic lifting but requires multiple stages of training, leading to slow convergence. Concurrently to us, Instance-NeRF [14] directly learn a label field, but they based their mask association on utilizing NeRF-RPN [15] to detect objects in a NeRF. Our approach, on the contrary, allows scaling to very high image resolutions while requiring only a small number (40-60) of neural field queries to render segmentation masks.

Structure from Motion: During mask association in InstanceMap, we take inspiration from scalable 3D reconstruction pipelines such as hLoc [16], including the use of visual descriptors for matching image viewpoints first, then applying keypoint matching as a preliminary for mask association. We utilize LoFTR [17] for keypoint extraction and matching.

III. METHOD

Given a sequence of N posed RGB images, (I_i, T_i) where I denotes the image and T pose, we first extract view-inconsistent instance masks M_i using a generic instance segmentation model such as Mask2Former or SAM.

A. Mask Association

We first generate pseudolabel masks with InstanceMap. Formally, define $\phi(M, r)$ to map a mask M and region r to a consistent label for the same 3D object across different masks and regions. We extend the popular hLoc [16] framework for scalable 3D reconstruction to mask association as follows:

Correspondence Generation: For each image I_i in the input sequence, we compute NetVLAD [16] feature descriptors. We then compute the cosine similarity among all pairs of images using these descriptors, defining a visual match as a similarity above a threshold of τ_{global} . For each visual pair, we extract dense correspondences (i.e. set of keypoints

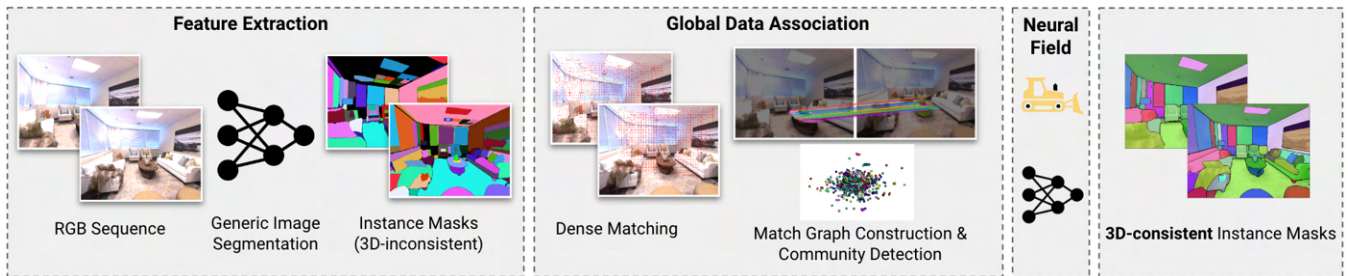


Fig. 2: **Overview of 3DIML.** A sequence of color images is segmented into object instances by an image segmentation backbone. The resulting masks produced are fed into InstanceMap, which produces instance masks consistent over all frames. These pseudo instance masks and their respective camera poses are used to supervise an instance label NeRF, which further improves consistency and resolves ambiguity present in the InstanceMap outputs. The *feature extraction* and *global data association* blocks together form InstanceMap.

in both images and a correspondence mapping from one set to the other) using LoFTR [17]. Since LoFTR computes keypoints independently per image pair, we aggregated keypoints from all pairs. Denote image I_i 's keypoints k_i and matches between I_i, I_j as m_{ij} . These correspondences are used in the next step to associate masks.

Since NetVLAD and LoFTR don't have 3D information, 3DIML only performs well if each image in the scan sequence contains enough context for these models. We observe empirically a good rule of thumb is to have at least one other recognizable landmark for frames containing near-identical objects.

Mask Association Graph: Insofar, our approach produces instance masks and dense pixel correspondences among images that share a visual overlap. However, segmentation models such as SAM [2] suffer multiple issues: (a) segmentations of the same object need not be consistent across images, owing to viewpoint and appearance variations; and (b) owing to over-segmentation of objects, there isn't usually a one-one correspondence among masks.

To associate instance masks across the entire scene, we construct a *mask association graph*, where the nodes are all instances detected across all images in the sequence. We then proceed to populate edges in this graph, corresponding to each pair of object instances that may be deemed a potential match. For each visual pair (i, j) , for all regions in M_i , we find a best-match region in M_j if one exists. Let r_u, r_v denote regions in M_i, M_j , $k_i(r_u), k_j(r_v)$ to be the keypoints of I_i, I_j contained in regions r_u, r_v , respectively, and $m_{ij}(r_u, r_v)$ be the matches between the regions. An edge is added for r_u, r_v if the *matching score* exceeds some threshold

$$\min \left(\frac{m_{ij}(r_u, r_v)}{|k_i(r_u)|}, \frac{m_{ij}(r_u, r_v)}{|k_j(r_v)|} \right) > \tau_{\text{local}}$$

Community Detection: Given this association graph, to remove noisy edges, we run community detection over this graph and merge all masks in the same community (determined by threshold size $\tau_{\text{community}}$; the remaining masks are discarded). We also filter out all labels that do not feature in a minimum number of frames, as there are insufficient observations to accurately associate them across views. To account for potentially overlapping instances, a common

occurrence due to perspective image formation, we compute the average mask area across all frames an instance exists in, which serves as a rasterization order for InstanceLift (instance masks with the largest area averaged over all frames are rasterized first). Example outputs are shown in Fig. 1.

B. Mask Refinement

$\phi(M, r)$ is inherently noisy due to varying segmentation hierarchies for different instance masks due to differing viewpoints as well as design specifics. To address this, in InstanceLift we feed the pseudolabel masks to a label NeRF, which resolves some ambiguities. Still, NeRF cannot handle extreme cases of label ambiguity, to which we devise a fast post-processing method that determines and merges colliding labels based on random renders from the label NeRF. The few remaining, if any, ambiguities can be corrected via sparse human annotation.

Neural Label Field Our label field represents the scene as a function $\Phi(x, d)$ that maps each point $\mathbf{x} \in \mathbb{R}^3$ and view-direction directions $\mathbf{d} \in \mathbb{S}^2$ to a corresponding density $\sigma(\mathbf{x})$ and other modalities $f(\mathbf{x}, \mathbf{d})$. These neural fields are queried via volumetric rendering by sampling per pixel rays across images, sampling points on rays, and estimating the volume rendering integral along those sampled points (see [7] for more details). We base our NeRF on the off-the-shelf Nerfacto architecture [18]. We add a multiresolution hashgrid [19] for instance labels in addition to one for color and geometry as well as an instance prediction MLP to render instance segmentation logits, which we supervise with cross entropy using the pseudolabel masks from InstanceMap. As in [5], [14], we block gradient flow from the label grid to the geometry grid. We train our model to learn volume density, color, and instance labels simultaneously.

Label Merging: Once the label NeRF is trained, we build a *redundancy graph* where the nodes, v denote labels and the number of edges, $|e|$ between nodes correlate with how likely the labels refer to the same object. Specifically, we randomly render K views and compare the rendered instance mask with some mask generated on the fly by FastSAM. For each region of FastSAM, we add pairwise edges between all labels that occupy greater than τ_{area} fraction of the region.

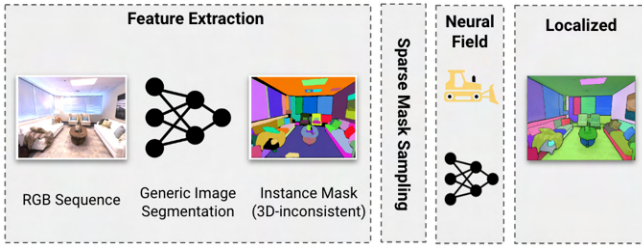


Fig. 3: InstanceLoc enables 3D-consistent instance segmentation for novel views of the scene unobserved by the InstanceMap pipeline. We leverage off-the-shelf instance segmentation models to first produce 3D-inconsistent instance labels for a new input image. We then query the label field over a sparse set of points on the image and use this to localize each 2D instance mask i.e., assign a 3D-consistent label to each mask.

Post graph construction, we merge labels a, b if

$$\frac{2|e_{ab}|}{\deg(v_a) + \deg(v_b)} > \tau_{merge}$$

Since we only need coarse information i.e. instance mask noise, we render images downsampled by a factor of 2.

C. Fast Instance Localization and Rendering

Training a label field enables us to predict 3D-consistent instance labels for novel viewpoints without rerunning *3DIML*. However, rendering every pixel is slow, and rendering from a novel viewpoint is often noisy. We propose a *fast localization* approach that instead precomputes instance masks for the input image using an instance segmentation model (here FastSAM [8]). Given this instance mask, for each instance region, we sample the corresponding pixelwise 3D object labels from the label NeRF and take the majority label. Another benefit is that the input instance masks can be constructed using prompts and edited before localization.

IV. EXPERIMENTS

We benchmark our method against Panoptic and Contrastive Lifting using the same Mask2Former frontend. For fairness, we render semantics as in [5], [6] using the same multiresolution hashgrid for semantics and instances. For other experiments, we utilize GroundedSAM as our frontend and FastSAM for runtime performance critical tasks such as label merging and InstanceLoc.

A. Datasets

We evaluate our methods on a challenging subset of scans from Replica and ScanNet, which provide ground truth annotations. Since our methods are based on structure from motion techniques, we utilize the Replica-vMap [20] sequences, which is more indicative of real-world collected image sequences. For Replica and ScanNet, we avoid scans with various incompatibilities with our method (multi-room, low visibility, Nerfacto doesn’t converge) as well as those containing many close-up views of identical objects that easily confuse NetVLAD and LoFTR.

B. Metrics

For lifting panoptic segmentation, we utilize Scene Level Panoptic Quality [5], defined as the Panoptic Quality for the concatenated sequence of images. For Grounded SAM, especially with instance masks for smaller objects, there is a divergence in alignment with ground truth annotations. As such, we report mIoU for predicted, reference masks that have IoU > 0.5 over all frames (TP in Scene Level Panoptic Quality) as well as the number of such matched masks and the total number of reference masks.

C. Implementation Details

For comparison with Panoptic and Contrastive Lifting, we follow their evaluation framework and map the Mask2Former outputs and ground truth semantic labels to 21 reduced ScanNet classes, considering only the “thing” subset of those classes. We train Panoptic and Contrastive Lifting using the code provided by [6]. We use the default parameterizations as their Replica experiment, which scales the number of iterations (128000) with the number of images. We build *3DIML* on top of the popular framework Nerfstudio [18] with `tinycudann` enabled [19]. We train until convergence for 5000 iterations for Replica-vMap (both Mask2Former and GroundedSAM) and 10000 iterations for ScanNet. For our panoptic experiments, we use $(\tau_{global}, \tau_{local}, \tau_{community}, \tau_{area}, \tau_{merge}) = (0.25, 0.8, 2, 0.15, 0.75)$. For GroundedSAM, we use $(0.25, 0.75, 3, 0.3, 0.6)$ for Replica-vMap and $(0.4, 0.75, 3, 0.25, 0.75)$ for ScanNet.

D. Results

Comparison with Panoptic and Contrastive Lifting: Table II shows the Scene Level Panoptic Quality for *3DIML* and other methods on Replica-vMap sequences subsampled by 10 (200 frames). We observe *3DIML* approaches Panoptic and Contrastive Lifting in performance despite not relying on updating or querying explicit/implicit world models for much of the pipeline, which makes it much faster. Figure 4 shows a comparison between the instances identified by Panoptic Lifting compared to *3DIML*.

We benchmark all runtimes using a single RTX 3090 post-mask generation. Specifically, Panoptic Lifting requires on average 5.7 hours of training over all scans, with min and max of 3.6 and 6.6 hours, respectively, since its runtime is dependent on the number of objects. Contrastive Lifting takes around 3.5 hours on average while *3DIML* runs under 20 minutes (14.5 minutes on average) for all scans. Note several components of *3DIML* can be easily parallelized, such as dense descriptors extraction using LoFTR and label merging. The runtime of our method is dependent on the number of correspondences produced by LoFTR, which doesn’t change for different frontend segmentation models, and we observe similar runtimes for other experiments.

Grounded SAM Table III shows our results for lifting GroundedSAM masks for Replica-vMap. From Figure 5 we see that InstanceLift is effective at interpolating labels missed by InstanceMap and resolving ambiguities produced



Fig. 4: Comparison between Panoptic Lifting and *3D DIML* for room0 from Replica-vMap

Scene	office0	office2	office3	office4	room0	room1	Average
Panoptic Lifting [5]	39.1	52.8	60.0	65.6	63.7	94.8	62.6
Contrastive Lift [6]	11.4	19.4	17.2	20.4	37.2	94.8	33.4
InstanceMap (ours)	34.9	26.1	41.3	27.3	63.1	52.5	40.8
InstanceLift (ours)	39.7	28.1	55.2	46.4	68.8	87.4	54.2

TABLE I: Quantitative comparison between panoptic lifting [5], contrastive lift [6], and our framework components, InstanceMap and InstanceLift. We measure the scene level panoptic quality metric (higher value indicates better performance). Our approach offers competitive performance while being far more efficient to train. Best performing numbers for each scene are in bold, while the second-best numbers are shaded yellow.

Method	Time
Panoptic lifting [5]	5.7 hours
Contrastive lift [6]	3.6 hours
InstanceMap (w/ LoFTR)	6.2 minutes
InstanceMap	2.1 minutes
InstanceLift (train)	5.6 minutes
InstanceLift (refine)	33.2 seconds
<i>3D DIML</i> (total)	14.5 minutes

TABLE II: Runtime in minutes benchmarked on a single RTX 3090 of Panoptic Lifting, Contrastive Lifting, and *3D DIML*




Fig. 5: InstanceLift is able to fill in labels missed by InstanceMap as well as correct ambiguities. Here we show comparisons between them for office0 and room0 from Replica-vMap.

by GroundedSAM¹. Figure 7 shows that InstanceMap and *3D DIML* are robust to large viewpoint changes and as well as duplicate objects, assuming nice scans, that is enough context for NetVLAD and LoFTR to somewhat distinguish between them. This is achieved by efficiently relying on implicit scene representation methods only at critical junctions, which gives *3D DIML* a large edge in runtime compared to Panoptic and Contrastive Lifting. Table IV and Fig. 6 illustrate our performance on ScanNet [21].

¹GroundedSAM produces lower quality frontend masks than SAM due to prompting using bounding boxes instead of a point grid.



Fig. 6: Some results for scans 0144_01, 0050_02, and 0300_01 from ScanNet [21] (one scene per row, top to bottom), showcasing how *3D DIML* accurately and consistently delineates instances in 3D.

Novel View Rendering and InstanceLoc Table V shows the performance if *3D DIML* on the second track provided in Replica-vMap. We observe that InstanceLift is effective at rendering new views, and therefore InstanceLoc performs well. For Replica-vMap and FastSAM, InstanceLoc takes on average 0.16s per localized image (6.2 frames per second). In addition, InstanceLoc can be applied as a post-processing step to the renders of the input sequence as a denoising operation.

E. Limitations and Future Work

In extreme viewpoint changes, our method sometimes produced discontinuous 3D instance labels. For example, on the worst performing scene, office2, we see that since the scan images only the front of a chair facing the back of the room and the back of a chair facing the front of the room for many frames, InstanceMap is not able to conclude these labels refer to the same object, and InstanceLift was unable to fix it, as NeRF’s correction performance rapidly degrades with increasing label inconsistency [12]. However, there are very few of these left per scene post *3D DIML*, and they can be easily fixed via sparse human annotation.

V. CONCLUSION

In this paper, we present *3D DIML*, which addresses the problem of 3D instance segmentation in a class-agnostic and computationally efficient manner. By employing a novel approach that utilizes InstanceMap and InstanceLift for generating and refining view-consistent pseudo instance masks from a sequence of posed RGB images, we circumvent the complexities associated with previous methods that only optimize a neural field. Furthermore, the introduction of InstanceLoc allows rapid localization of instances in unseen views by combining fast segmentation models and a refined neural label field. Our evaluations across Replica and ScanNet and different frontend segmentation models showcase *3D DIML*’s speed and effectiveness. It offers a promising avenue for real-world applications requiring efficient and accurate scene analysis.

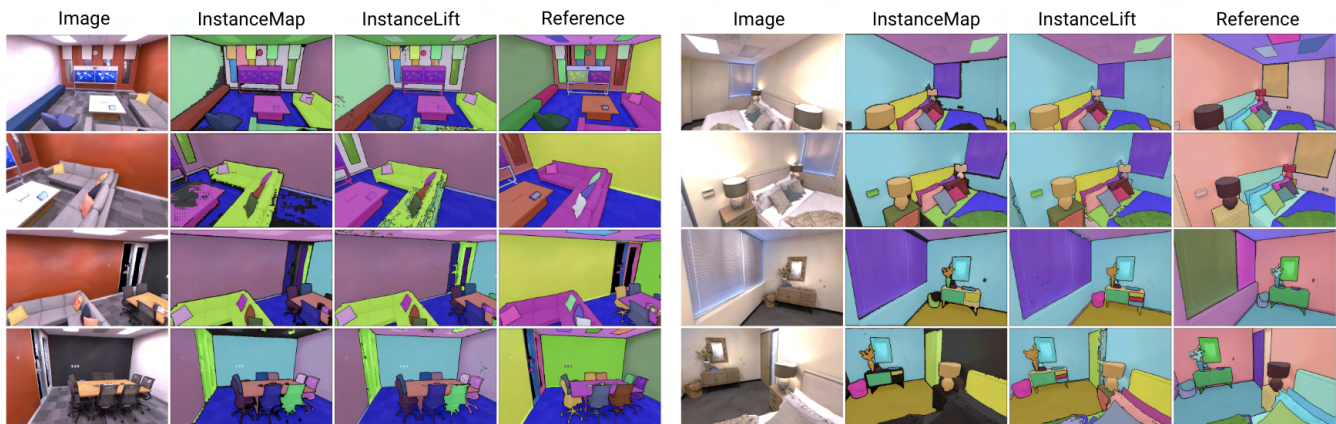


Fig. 7: **Qualitative results** on *office3* and *room1* from the Replica Vmap split [20]. Both InstanceMap and InstanceLift are able to maintain quality and consistency over the image sequence despite duplicate objects due to sufficient image context overlap across the sequence.

Scene	office0	office2	office3	office4	room0	room1	Average
InstanceMap	0.740 / 14	0.753 / 27	0.769 / 30	0.758 / 22	0.825 / 34	0.729 / 25	0.762
InstanceLift	0.785 / 18	0.780 / 30	0.799 / 36	0.794 / 26	0.851 / 38	0.769 / 25	0.796

TABLE III: Quantitative (mIoU, TP) results for GroundedSAM frontend on Replica-vMap. The average number of reference instances for all Replica scenes we evaluated on is 67.

Scene	0050_02	0144_01	0221_01	0300_01	0389_00	0494_00	0693_00	Average
InstanceMap	0.675 / 12	0.646 / 14	0.667 / 12	0.726 / 13	0.704 / 13	0.716 / 6	0.700 / 14	0.691
InstanceLift	0.657 / 14	0.664 / 14	0.621 / 10	0.754 / 14	0.745 / 12	0.712 / 8	0.782 / 16	0.705

TABLE IV: Quantitative (mIoU, TP) results for GroundedSAM frontend on ScanNet. The average number of reference instances for all ScanNet scenes we evaluated on is 32.

Scene	office0	office2	office3	office4	room0	room1	Average
InstanceLift	0.822 / 16	0.787 / 29	0.762 / 35	0.769 / 30	0.849 / 39	0.768 / 28	0.793
InstanceLoc	0.845 / 15	0.842 / 30	0.799 / 37	0.812 / 30	0.877 / 41	0.800 / 28	0.829

TABLE V: Quantitative (mIoU, TP) results for InstanceLift and InstanceLoc on novel views over the Replica-vMap split [20].




Fig. 8: InstanceLoc is able to correct for noise rendered by InstanceLift.



Fig. 9: Our methods do not perform well in cases where the scan sequence contains only images of the different sides of an object (chair) or surface (floor) from differing directions without any smooth transitions in between, which occurs for *office2* from Replica (vMap split [20]).

REFERENCES

- [1] B. Cheng, I. Misra, A. G. Schwing, A. Kirillov, and R. Girdhar, “Masked-attention mask transformer for universal image segmentation,” *CoRR*, vol. abs/2112.01527, 2021. [Online]. Available: <https://arxiv.org/abs/2112.01527> 1, 2
- [2] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo *et al.*, “Segment anything,” *arXiv preprint arXiv:2304.02643*, 2023. 1, 2, 3
- [3] J. Schult, F. Engelmann, A. Hermans, O. Litany, S. Tang, and B. Leibe, “Mask3d: Mask transformer for 3d semantic instance segmentation,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 8216–8223. 1
- [4] A. Takmaz, E. Fedele, R. W. Sumner, M. Pollefeys, F. Tombari, and F. Engelmann, “Openmask3d: Open-vocabulary 3d instance segmentation,” *arXiv preprint arXiv:2306.13631*, 2023. 1
- [5] Y. Siddiqui, L. Porzi, S. R. Bulò, N. Müller, M. Nießner, A. Dai, and P. Kotschieder, “Panoptic lifting for 3d scene understanding with neural fields,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 9043–9052. 2, 3, 4, 5
- [6] Y. Bhargat, I. Laina, J. F. Henriques, A. Zisserman, and A. Vedaldi, “Contrastive lift: 3d object instance segmentation by slow-fast contrastive fusion,” 2023. 2, 4, 5
- [7] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “Nerf: Representing scenes as neural radiance fields

- for view synthesis,” *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021. 2, 3
- [8] X. Zhao, W. Ding, Y. An, Y. Du, T. Yu, M. Li, M. Tang, and J. Wang, “Fast segment anything,” *arXiv preprint arXiv:2306.12156*, 2023. 2, 4
- [9] L. Ke, M. Ye, M. Danelljan, Y. Liu, Y.-W. Tai, C.-K. Tang, and F. Yu, “Segment anything in high quality,” *arXiv preprint arXiv:2306.01567*, 2023. 2
- [10] Y. Zhang, X. Huang, J. Ma, Z. Li, Z. Luo, Y. Xie, Y. Qin, T. Luo, Y. Li, S. Liu, Y. Guo, and L. Zhang, “Recognize anything: A strong image tagging model,” 2023. 2
- [11] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, C. Li, J. Yang, H. Su, J. Zhu, and L. Zhang, “Grounding dino: Marrying dino with grounded pre-training for open-set object detection,” 2023. 2
- [12] S. Zhi, T. Laidlow, S. Leutenegger, and A. J. Davison, “In-place scene labelling and understanding with implicit scene representation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15 838–15 847. 2, 5
- [13] A. Chen, Z. Xu, A. Geiger, J. Yu, and H. Su, “Tensorf: Tensorial radiance fields,” in *European Conference on Computer Vision (ECCV)*, 2022. 2
- [14] B. Hu, J. Huang, Y. Liu, Y.-W. Tai, and C.-K. Tang, “Instance neural radiance field,” *arXiv preprint arXiv:2304.04395*, 2023. 2, 3
- [15] —, “Nerf-rpn: A general framework for object detection in nerfs,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 23 528–23 538. 2
- [16] P.-E. Sarlin, C. Cadena, R. Siegwart, and M. Dymczyk, “From coarse to fine: Robust hierarchical localization at large scale,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12 716–12 725. 2
- [17] J. Sun, Z. Shen, Y. Wang, H. Bao, and X. Zhou, “Loftr: Detector-free local feature matching with transformers,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 8922–8931. 2, 3
- [18] M. Tancik, E. Weber, E. Ng, R. Li, B. Yi, T. Wang, A. Kristoffersen, J. Austin, K. Salahi, A. Ahuja *et al.*, “Nerfstudio: A modular framework for neural radiance field development,” in *ACM SIGGRAPH 2023 Conference Proceedings*, 2023, pp. 1–12. 3, 4
- [19] T. Müller, A. Evans, C. Schied, and A. Keller, “Instant neural graphics primitives with a multiresolution hash encoding,” *ACM Transactions on Graphics*, vol. 41, no. 4, p. 1–15, Jul. 2022. [Online]. Available: <http://dx.doi.org/10.1145/3528223.3530127> 3, 4
- [20] X. Kong, S. Liu, M. Taher, and A. J. Davison, “vmap: Vectorised object mapping for neural field slam,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 952–961. 4, 6
- [21] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, “Scannet: Richly-annotated 3d reconstructions of indoor scenes,” in *Proc. Computer Vision and Pattern Recognition (CVPR)*, *IEEE*, 2017. 5