3D-Aware Vision-Language Models Fine-Tuning with Geometric Distillation

Seonho Lee*, Jiho Choi*, Inha Kang, Jiwook Kim, Junsung Park, Hyunjung Shim[†] Graduate School of Artificial Intelligence, KAIST, Republic of Korea

{glanceyes, jihochoi, rkswlsj13, tom919, jshackist, kateshim}@kaist.ac.kr

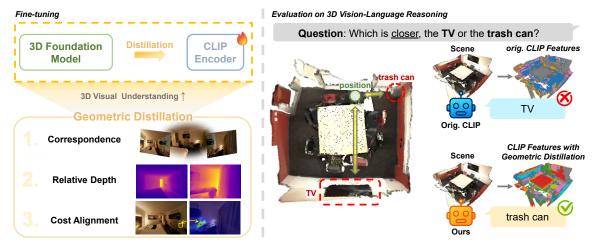


Figure 1: **Geometric Distillation enhances 3D spatial reasoning in vision-language models.** By distilling geometric cues such as correspondences, relative depth, and cost alignment from 3D foundation models, our method improves 3D visual understanding and enables accurate reasoning in tasks like answering which object is closer.

Abstract

Vision-Language Models (VLMs) have shown remarkable performance on diverse visual and linguistic tasks, yet they remain fundamentally limited in their understanding of 3D spatial structures. We propose Geometric Distillation, a lightweight, annotation-free finetuning framework that injects human-inspired geometric cues into pretrained VLMs without modifying their architecture. By distilling (1) sparse correspondences, (2) relative depth relations, and (3) dense cost volumes from off-the-shelf 3D foundation models (e.g., MASt3R, VGGT), our method shapes representations to be geometry-aware while remaining compatible with natural image-text inputs. Through extensive evaluations on 3D visionlanguage reasoning and 3D perception benchmarks, our method consistently outperforms prior approaches, achieving improved 3D spatial reasoning with significantly lower computational cost. Our work demonstrates a scalable and efficient path to bridge 2D-trained VLMs with 3D understanding, opening up wider use in spatially grounded multimodal tasks.

1 Introduction

Vision-Language Models (VLMs) (e.g., CLIP (Radford et al., 2021), ALIGN (Jia et al., 2021), and BLIP (Li et al., 2022, 2023)), trained on large-scale image-text datasets, have demonstrated competitive performance on diverse multimodal tasks (Li et al., 2021; Gao et al., 2024; Lee et al., 2022). Despite their progress, these models struggle with understanding 3D spatial structures (El Banani et al., 2024; Man et al., 2024; Chen et al., 2024; Danier et al., 2024; Li et al., 2024; Kamath et al., 2023; Qiu et al., 2025). Specifically, VLMs remain limited in grounded spatial reasoning tasks such as depth ordering, occlusion, or object layout in a scene (El Banani et al., 2024; Chen et al., 2024; Kamath et al., 2023). This limitation stems from their reliance on 2D projections, which lack depth cues and multiview supervision (Eigen et al., 2014; Tulsiani et al., 2017; Qin et al., 2019). It is illustrated in Figure 1, where features of standard VLMs like CLIP incorrectly predict relative depth due to their limited 3D awareness. These shortcomings greatly hinder applications requiring spatial reasoning, including navigation, scene understanding, and robotic planning (Peng et al., 2023; Shridhar et al., 2022; Hong et al., 2023).

^{*}Equal contribution

†Corresponding author

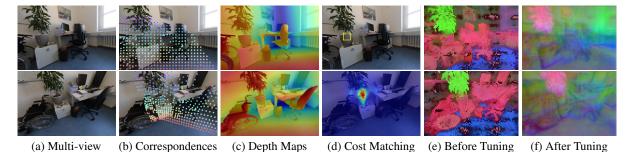


Figure 2: Geometric cues and PCA visualization of feature transformation through geometric distillation.

To address this, recent work has explored injecting 3D priors into VLMs. FiT3D (Yue et al., 2024) reconstructs 3D scenes from multi-view images using Gaussian splatting (Kerbl et al., 2023), and then aligns VLM's features with those rendered 3D views. Multiview Equivariant Finetuning (MEF) (You et al., 2024) improves 3D equivariance by reinforcing feature consistency across rendered views of the same object. SpatialVLM (Chen et al., 2024) improves its spatial reasoning abilities by generating billions of synthetic spatial question—answer pairs to train VLMs.

Despite these advancements, existing methods suffer from notable drawbacks. FiT3D incurs a high computational cost and suffers from semantic degradation due to its reliance on explicit 3D reconstruction. MEF depends on 3D object-centric datasets, which restricts its generalizability to real-world scenes. SpatialVLM requires extensive synthetic data generation and task-specific tuning, making it resource-intensive and less flexible. These limitations motivate the need for more efficient and generalizable approaches to endow VLMs with robust 3D awareness.

propose Geometric Distillation, lightweight and glanannotation-free fine-tuning framework that enriches 3D spatial understanding in VLMs. Our approach introduces supervision signals aligned with human perceptual strategies, derived from pretrained 3D foundation models such as MASt3R (Leroy et al., 2024) and VGGT (Wang et al., 2025) as in Figure 2 (a) - (d). First, we supervise the VLM to align features at sparse correspondences that are visually stable and semantically meaningful regions, such as object corners or room boundaries, derived from pretrained 3D foundation models without any explicit 3D annotations. These locations provide strong geometric anchors across views, and featurelevel matching at these points encourages the model to learn consistent and viewpoint-invariant representations. Second, we supervise relative

depth reasoning through ordinal comparisons both within and across views. This reflects the human tendency to reason in relative terms and aligns with the way spatial relationships are expressed in language (Zhang et al., 2022b; Auty and Mikolajczyk, 2023). Lastly, we incorporate dense cost volume alignment, which captures soft correspondences across views by fully exploiting the geometric priors and warping relationships (Weinzaepfel et al., 2022; An et al., 2024) provided by 3D foundation models, thereby enabling the model to learn fine-grained geometric consistency. These signals collectively reshape the visual representations into a geometry-aware space that better supports grounded spatial reasoning and improves VLM performance on 3D-aware tasks as shown in Figure 2 (e), (f). Additionally, since our approach operates without modifying the VLM's architecture and retains compatibility with natural image-text inputs, it preserves the strong generalization capabilities of the original model.

To overcome these limitations, we draw inspiration from human 3D spatial perception. Humans infer depth and structure from sparse relational cues such as occlusions, relative size, and perspective, rather than absolute measurements (Todd and Norman, 2003; Howard and Rogers, 1995; Landy et al., 1995). In addition, spatial relationships are often expressed in language using relative terms (e.g., "next to the table", "behind the sofa") rather than absolute metric units, suggesting that the reasoning is both perceptually and linguistically grounded. These observations suggest that incorporating human-inspired geometric cues into VLM can enhance their spatial reasoning abilities.

Our approach enhances the model's ability to infer spatial relations, such as object proximity, without explicit 3D labels or costly reconstruction. We demonstrate consistent improvements across a range of 3D-aware tasks, including semantic correspondence, depth estimation, and 3D visual question answering. Our method outper-

forms strong baselines on benchmarks such as PF-PASCAL, TAP-Vid, and ScanQA, illustrating both the effectiveness and scalability of our approach.

2 Related Work

Fine-tuning VFMs and VLMs. Various attempts have been made to integrate 3D information into Visual Foundation Models (VFMs) or Vision-Language Models (VLMs) (Yue et al., 2024; You et al., 2024). FiT3D (Yue et al., 2024) lifts 2D visual features into a 3D Gaussian representation and re-renders them from multiple viewpoints. By fine-tuning, this approach guides the original 2D features to align with re-rendered features, which enhances 3D perception in VFMs. However, its dense L1 optimization introduces noise, which potentially leads to semantic information loss and significant computational overhead. Multiview Equivariance Finetuning (MEF) (You et al., 2024) enhances 3D correspondence understanding by maximizing cross-view feature equivariance within pretrained foundation models. This allows them to improve on tasks such as camera pose estimation, object tracking, and semantic transfer. Nevertheless, MEF requires explicit 3D annotations and does not provide direct supervision for depth understanding. SpatialVLM (Chen et al., 2024) generates extensive 3D spatial QA corpora using pretrained detectors, depth estimators, and segmentation models. Training on this large-scale data strengthens the spatial question-answering capabilities of VLMs. However, the reliance on massive synthetic datasets limits their practicality. In our work, we address these limitations by introducing a lightweight and annotation-free fine-tuning method that efficiently enhances 3D spatial reasoning in VLMs.

3D Foundation Models. Recently, geometry-based models have emerged as foundation models for 3D vision. CroCo (Weinzaepfel et al., 2022, 2023) performs self-supervised cross-view completion by reconstructing one view from another, which allows the model to acquire multiview consistent features. Based on CroCo pretraining, DUSt3R (Wang et al., 2024) introduces a unified approach to directly estimate scene point maps from two or more images taken from different viewpoints. DUSt3R effectively simplifies the Structure-from-Motion (SfM) pipeline. MASt3R (Leroy et al., 2024) further extends these approaches by incorporating a global matching head that aligns partial reconstructions and predicts dense 3D cor-

respondences. These models inherently provide 3D perceptual priors by learning scene geometry without explicit supervision or accurate dense reconstructions from limited views. Additionally, VGGT (Wang et al., 2025) introduces a large transformer-based model to jointly estimate camera poses, depth maps, and point clouds from a few images. Training VGGT on large-scale 3D datasets enables accurate depth prediction even from a single image, which significantly improves 3D downstream tasks. Consequently, these models embed critical 3D knowledge that is beneficial for robust 3D understanding. In our work, we propose a method to effectively inject these rich 3D priors into VLMs.

Bridging VLMs and 3D Understanding. Recent studies have explored analyzing or improving vision-language representations to better understand 3D scenes. Lexicon3D (Man et al., 2024) evaluates various vision foundation encoders across vision-language reasoning tasks and identifies their strengths and limitations. Notably, image-text alignment supervised models (Qiu et al., 2025; Auty and Mikolajczyk, 2023; Radford et al., 2021; Jia et al., 2021) still exhibit substantial weaknesses in complex 3D spatial reasoning and languagedriven question answering tasks. This suggests that vision-language pretraining alone may not sufficiently capture comprehensive 3D concepts. These observations underscore the necessity of incorporating explicit 3D signals or specialized training strategies into VLMs. To address these limitations, various approaches have been proposed. Some studies (Hegde et al., 2023) extend CLIP via prompt tuning by prepending learnable tokens to the vision encoder and training it contrastively on rendered 3D object images paired with textual labels. Other notable efforts include PointCLIP (Zhang et al., 2022a; Zhu et al., 2023), which aligns 3D point clouds with CLIP's textual embedding space, and methods designed to enhance text-image alignment in 3D contexts (Kim et al., 2023; Zeng et al., 2021). Collectively, these studies introduce additional representations or strategies to enrich 3D understanding within VLMs. In contrast, our work directly injects robust 3D knowledge into 2D VLMs using multi-view images. This enables leveraging their inherent rich 2D vision-language priors without relying on explicit supervision from other 3D data modalities such as point clouds or 3D Gaussians.

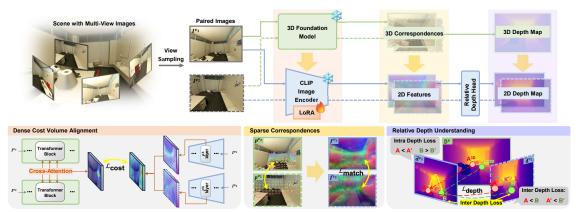


Figure 3: Overview of Geometric Distillation Architecture. A 3D foundation model extracts geometric cues including (1) sparse correspondences, (2) depth maps, and (3) dense cost volumes from multi-view inputs. These cues supervise a frozen CLIP image encoder with a lightweight adapter (LoRA) via three loss branches: \mathcal{L}_{match} , \mathcal{L}_{depth} , and \mathcal{L}_{cost} . The distillation enables the VLM to acquire 3D spatial awareness without explicit 3D annotations.

3 Proposed Method

We propose a geometric knowledge distillation framework that transfers 3D spatial understanding from high-performance 3D foundation models such as MASt3R (Leroy et al., 2024) and VGGT (Wang et al., 2025) into a pretrained visionlanguage model (VLM) (Radford et al., 2021; Jia et al., 2021) without requiring any ground truth 3D annotations. Inspired by human perception, which infers spatial structure by integrating visual cues from multiple viewpoints, our method uses paired images, $\{I^{v_1}, I^{v_2}\}\$, of the same scene captured from different perspectives v_1 and v_2 . From these image pairs, we extract geometric signals including sparse correspondences, ordinal depth relations, and viewpoint-induced disparities, which guide the VLM to learn geometry-aware representations. An overview of our framework is illustrated in Figure 3.

Our framework obtains these geometric cues using a teacher model that generates pseudo-3D supervision from image pairs. Specifically, we utilize the following information provided by 3D foundation models: (i) sparse correspondences $\mathbb{P}^{v_1,v_2} = \{(p_i^{v_1},p_i^{v_2})\}_{i=0}^{|\mathbb{P}^{v_1,v_2}|}$ for matching 3D points across views, (ii) estimated depth maps $\tilde{\mathbb{D}}^{v_1}, \tilde{\mathbb{D}}^{v_2}$ for each viewpoint, and (iii) a dense cost volume, $\mathbb{C}^{v_1 \to v_2}$, representing patch-level features similarity between two viewpoints. These heterogeneous signals serve as supervision for three complementary objectives: sparse correspondence matching, relative depth learning using both intraview and inter-view comparisons, and alignment of dense feature similarity. Combined, they enrich the model's multimodal representations and facilitate 3D-aware reasoning in complex scenes.

3.1 Sparse Correspondences

Background. Humans often rely on sparse but stable visual features, such as corners or edges, to estimate spatial layout. In a similar way, sparse correspondences across views serve as geometric anchors that help enforce cross-view consistency and identify matching 3D points. These signals are essential for enforcing consistency across viewpoints (Leroy et al., 2024; Wang et al., 2025) and have been widely adopted in multi-view geometry (Weinzaepfel et al., 2022, 2023; An et al., 2024) as well as recent representation learning methods such as MEF (You et al., 2024). To exploit these correspondences, we adopt a feature-matching objective that promotes accurate feature-level alignment between image pairs. Given a set of pseudo correspondence pairs $\mathbb{P}^{v_1,v_2} = \{(p_i^{v_1},p_i^{v_2})\}_{i=1}^{|\mathbb{P}^{v_1,v_2}|}$ generated by a geometric teacher, we extract local image features $\{(f_i^{v_1}, f_i^{v_2})\}_{i=1}^{|\mathbb{P}^{v_1, v_2}|}$ and intermediate ate patch features $\{h^{v_*}\}$ from each viewpoint. We adopted a matching-based loss (Brown et al., 2020; You et al., 2024) that encourages high retrieval performance by maximizing the Smooth Average Precision (SmoothAP) (Brown et al., 2020), computed within a spatial neighborhood. For a query feature f_i , the SmoothAP is calculated using positive matches \mathbb{P}^{v_1,v_2} and negative matches (nonmatches), $\mathcal{N}(i)$, of point p_i as:

$$SmoothAP_{v_1 \to v_2} = \frac{1}{|\mathbb{P}^{v_1, v_2}|} \sum_{i \in \mathbb{P}^{v_1, v_2}} \frac{1 + \sigma(D_{ii})}{1 + \sigma(D_{ii}) + \sum_{j \in \mathcal{N}(i)} \sigma(D_{ij})}, (1)$$

where $D_{ij} = f_j^{v_2} \cdot f_i^{v_1} - f_i^{v_1} \cdot f_i^{v_1}$ measures the difference in similarity between features, and $\sigma(x)$ denotes the sigmoid function. This objective promotes higher similarity for true matches than for

non-matches, thereby incorporating relative similarity into the training. To ensure symmetry across views, we apply the objective in both matching directions and define the final loss as:

$$\mathcal{L}_{\text{match}} = 1 - \frac{1}{2} \{ \text{SmoothAP}_{v_1 \to v_2} + \text{SmoothAP}_{v_2 \to v_1} \}.$$
(2)

3.2 Relative Depth Understanding

To complement sparse correspondences, we enhance the VLM's geometric reasoning by supervising its understanding of relative depth. Unlike absolute depth estimation, which is fundamentally ambiguous in monocular settings due to scale uncertainty, relative depth reasoning (i.e., determining which of two points is closer) is intuitive and practically robust across domains (Todd and Norman, 2003; Howard and Rogers, 1995; Landy et al., 1995). Numerous studies (Fu et al., 2018; Chen et al., 2016; Xian et al., 2020; Zoran et al., 2015) show that models trained with ordinal depth constraints generalize better to diverse scenes and produce sharper depth maps with preserved structure.

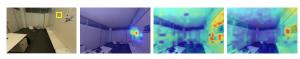
Inspired by this, we leverage the outputs of high-capacity 3D foundation models (e.g., MASt3R (Leroy et al., 2024), VGGT (Wang et al., 2025)) to construct pseudo ground-truth relative depth labels. This approach allows us to inject 3D awareness into VLMs without explicit 3D supervision or reconstruction. The learning proceeds on two levels: intra-view and inter-view, capturing both local monocular cues and multi-view disparities, akin to human depth perception mechanisms. **Intra-view Relative Depth.** Given an image I^v , we sample point pairs $(x, y) \in \mathcal{P}^v$ and determine their ordinal pseudo ground-truth relation using the depth map \mathbb{D}^v provided by a 3D foundation model (e.g., MASt3R, VGGT). The relative depth ordering is defined as:

$$s_{xy} = sign(\tilde{d}_x - \tilde{d}_y) \in \{-1, +1\},$$
 (3)

where \tilde{d}_x and \tilde{d}_y denote the estimated depths of points x and y from viewpoint v, respectively. The VLM predicts a scalar depth ranking score $\hat{\mathbf{s}}_{xy}$ for each pair based on its encoded features, and is trained with a logistic ranking loss (Chen et al., 2009; Fu et al., 2018):

$$\mathcal{L}_{\text{intra_depth}} = \frac{1}{|\mathcal{P}^v|^2} \sum_{(x,y) \in \mathcal{P}^v} \log \left(1 + \exp \left[-\mathbf{s}_{xy} \cdot \hat{\mathbf{s}}_{xy} \right] \right). \tag{4}$$

This loss encourages correct ordinal predictions without relying on metric depth values, allowing



(a) Anchor (b) MASt3R (c) V

(c) Vanilla

(d) Ours

Figure 4: **Visualization of cost volume**. (a) Anchor view with query location (yellow box). Cost volume heatmaps from (b) the teacher (MASt3R), (c) the vanilla CLIP, and (d) after geometric distillation. The proposed method better captures localized geometric similarity, closely aligning with the teacher's output.

the model to learn scale-invariant depth cues from local monocular structure.

Inter-view Relative Depth. To further infuse geometric awareness, we supervise relative depth relationships across multiple views, as absolute depth values may differ due to scale variations between viewpoints. Unlike intra-view supervision, which assumes a consistent scale within a single image, inter-view supervision requires the model to reason about depth differences under potential scale shifts.

Given a correspondence pair $(p_i^{v_1}, p_i^{v_2}) \in \mathbb{P}^{v_1, v_2}$ that observes the same 3D point from views v_1 and v_2 , we extract the pseudo ground-truth depths $\tilde{d}_i^{v_1}$ and $\tilde{d}_i^{v_2}$ from the teacher model's depth maps $\tilde{\mathbb{D}}^{v_1}$ and $\tilde{\mathbb{D}}^{v_2}$, respectively. To mitigate the effect of absolute scale mismatch, we define a bounded signed depth difference using the \tanh function as $\delta_i^* = \tanh(\tilde{d}_i^{v_1} - \tilde{d}_i^{v_2})$. The model is trained to regress this value using a lightweight MLP head, which is applied to the feature representations of each view. The loss is defined as:

$$\mathcal{L}_{\text{inter_depth}}^{v_1, v_2} = \frac{1}{\left| \mathbb{P}^{v_1, v_2} \right|} \sum_{i \in \mathbb{P}^{v_1, v_2}} \left| \hat{\delta}_i - \delta_i^* \right|. \quad (5)$$

This supervision encourages the model to be sensitive to viewpoint-induced disparities and relative geometry, even in the absence of explicit camera calibration or metric consistency. To jointly capture both local (intra-view) and cross-view (inter-view) depth relationships, we define the final relative depth loss as a combination of the two components: $\mathcal{L}_{\mathsf{depth}} = \sum_{p} \{\mathcal{L}_{\mathsf{intra_depth}}^{v_p} + \sum_{q} \mathcal{L}_{\mathsf{inter_depth}}^{v_p, v_q} \}$. By unifying intra-view ordinal supervision with interview relative regression, the model learns to infer consistent and structurally-aware depth relationships. This multi-scale depth reasoning framework fosters a more human-like, scale-invariant understanding of 3D geometry, enhancing the generalization ability of vision-language models across diverse visual domains.

Table 1: Comparison of zero-shot semantic correspondence on PF-PASCAL.

Method	Dataset]	Different View	s	Same Views			
		PCK@0.05	PCK@0.10	PCK@0.15	PCK@0.05	PCK@0.10	PCK@0.15	
(Vanilla) CLIP FiT3D (Yue et al., 2024) MEF (You et al., 2024)	ScanNet++ Objaverse	16.61 15.90 21.18	26.96 23.40 33.54	37.64 30.34 43.58	18.23 14.93 25.94	32.27 26.52 43.33	43.01 34.56 53.87	
Ours Ours	Objaverse ScanNet++	25.87 28.48 (+11.87)	39.85 43.07 (+16.11)	50.21 53.55 (+15.91)	36.77 42.16 (+23.93)	56.61 61.57 (+29.30)	67.93 72.16 (+29.15)	

¹ The best score is **bold** and the second-best score is <u>underlined</u>. These are the same for all experiments.

Table 2: Comparison of video tracking on TAP-Vid and pose estimation on OnePose-LowTexture.

Method	Dataset	Video	Tracking	Pose Estimation				
		Avg. Jaccard Index	Avg. Position Accuracy	1cm-1deg	3cm-3deg	5cm-5deg		
(Vanilla) CLIP	=	27.73	42.59	2.50	19.32	33.11		
FiT3D (Yue et al., 2024)	ne et al., 2024) ScanNet++ 28.45		43.51	2.86	20.14	34.75		
MEF (You et al., 2024)	Objaverse	34.61	50.58	6.32	36.00	52.33		
Ours	Objaverse	35.60	54.65	8.50	39.30	57.68		
Ours	ScanNet++	40.09 (+12.36)	57.75 (+15.16)	10.96 (+8.46)	44.93 (+25.61)	63.65 (+30.54)		

3.3 Dense Cost Volume Alignment

Beyond sparse matching and relative depth supervision, we introduce a dense cost volume alignment method to extract richer geometric cues from intermediate features of 3D foundation models. This alignment is further enhanced by leveraging geometric priors from cross-view completion models such as CroCo (Weinzaepfel et al., 2022, 2023), and transformer-based models using crossattention mechanisms across multiple views like VGGT (Wang et al., 2025). Recent findings from ZeroCo (An et al., 2024) show that cross-attention maps learned through cross-view completion pretext tasks encode high-quality dense correspondences, effectively acting as self-supervised cost volumes. These maps inherently learn to warp source features to reconstruct masked target views by estimating correspondences across views. By treating these attention-derived correspondences as pseudo ground-truth warping functions, we can supervise the VLM's dense feature similarity to better reflect geometric consistency, thereby enhancing its capacity for dense 3D-aware reasoning.

To enforce dense geometric consistency across entire feature maps, we align the feature similarities produced by a vision-language model with geometrically grounded predictions from a 3D foundation model as in Figure 4. Given two views v_1 and v_2 , we construct a 4D cost volume that encodes normalized feature similarity between all spatial positions (patch index) across the views:

$$\mathbb{C}_{v_1 \to v_2}(i,j) = \frac{h_i^{v_1} \cdot h_j^{v_2}}{\|h_i^{v_1}\| \|h_j^{v_2}\|},\tag{6}$$

where $h_i^{v_*} \in \mathbb{R}$ denotes the intermediate feature vector at patch index i in view v_1 , and j is a corre-

sponding patch index in view v_2 . This similarity matrix captures the VLM's inherent geometric understanding between all patch pairs across views. We convert this cost volume into a probability distribution using temperature-scaled softmax as:

$$P_{v_1 \to v_2}(j \mid i) = \operatorname{softmax}_j(\mathbb{C}_{v_1 \to v_2}(i, j) / \tau), (7)$$

where temperature τ controls the sharpness of the matching distribution. The geometric teacher provides target distributions $\tilde{P}_{v_1 \to v_2}$ derived from its robust 3D understanding. Our alignment loss minimizes the Jensen-Shannon Divergence (Menéndez et al., 1997) as:

$$\mathcal{L}_{\text{cost}} = \frac{1}{2} \{ D_{\text{KL}} (\tilde{P}_{v_{1} \rightarrow v_{2}} \| P_{v_{1} \rightarrow v_{2}}) + D_{\text{KL}} (\tilde{P}_{v_{2} \rightarrow v_{1}} \| P_{v_{2} \rightarrow v_{1}}) \}.$$
(8)

This dense supervision compels the VLM's feature similarities to mirror the teacher's geometrically grounded predictions, enforcing subpixel-level geometric awareness.

3.4 Overall Objective

To jointly train the vision-language model with rich geometric supervision, we combine the proposed loss components into a single objective function. Given a pair of images (I^{v_1}, I^{v_2}) from the same scene, the total loss is defined as:

$$\mathcal{L}_{\text{total}} = \lambda_{\text{match}} \mathcal{L}_{\text{match}} + \lambda_{\text{depth}} \mathcal{L}_{\text{depth}} + \lambda_{\text{cost}} \mathcal{L}_{\text{cost}}.$$
(9)

where λ_{match} , λ_{depth} , and λ_{cost} are hyperparameters for balancing each loss term.

4 Experiments

4.1 Experimental Setups

Datasets. We evaluate our method in two main sets of downstream tasks to examine the effectiveness

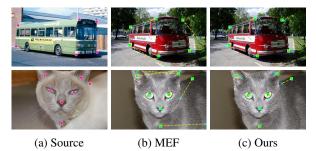


Figure 5: **Semantic Transfer.** (a) Source image with annotated keypoints. Transfer results using (b) MEF (You et al., 2024) and (c) our approach. Our method produces more accurate and spatially consistent transfers.

of our 3D-aware VLM representations: 3D visual understanding and vision-language understanding tasks. Specifically, to measure the 3D correspondence understanding, we conduct experiments on three downstream benchmarks introduced by (You et al., 2024): (1) semantic correspondence on PF-PASCAL (Ham et al., 2016), (2) video tracking on TAP-Vid (Doersch et al., 2022), and (3) object pose estimation on the OnePose-LowTexture dataset (He et al., 2022). Additionally, we perform experiments on downstream tasks for dense scene understanding via linear probing as in FiT3D (Yue et al., 2024), including semantic segmentation on ADE20K (Zhou et al., 2019) and VOC2012 (Everingham et al., 2015), and monocular depth estimation on ScanNet++ (Yeshwanth et al., 2023) and KITTI (Geiger et al., 2013). Furthermore, we assess improvements in 3D vision-language understanding by evaluating our method on the 3D visual question-answering benchmarks SQA3D (Ma et al., 2022) and ScanQA (Azuma et al., 2022).

Implementation Details. We fine-tune the ViT-based CLIP model for up to 500 epochs on either Objaverse (Deitke et al., 2023) or ScanNet++. We perform parameter-efficient fine-tuning through LoRA (Hu et al., 2022), adopting settings similar to those used in MEF (You et al., 2024). Our method primarily leverages MASt3R (Leroy et al., 2024) as a pretrained 3D foundation teacher during geometric distillation. Further implementation details, including experiments with VGGT (Wang et al., 2025), are provided in the appendix.

4.2 Experimental Results

4.2.1 3D Visual Understanding

3D Correspondence Understanding. We evaluate how effectively our distilled 3D-aware VLM representations capture robust multi-view correspondences, following established protocols from You

et al. (2024). As summarized in Tables 1 and 2, the baseline CLIP and FiT3D (Yue et al., 2024) exhibit limited performance. Specifically, FiT3D slightly degrades the ability of semantics matching, corroborating findings by (You et al., 2024). MEF (You et al., 2024) significantly improves performance as it leverages explicit 3D annotations. Nevertheless, our approach consistently outperforms MEF even without such annotations. On the Objaverse dataset, our geometric distillation yields notable improvements over the vanilla CLIP. Moreover, training on the real-world ScanNet++ dataset results in further substantial gains of +11.87% in PCK@0.05, +12.36% in average Jaccard index, and +8.46% accuracy at the 1cm-1deg threshold. This demonstrates the practical value and strong generalization power of our method. Unlike MEF, which indiscriminately uses 3D annotations, our distillation naturally selects semantically meaningful key regions, leading to more effective correspondence learning. These observations confirm that our approach effectively transfers strong geometric priors into VLM representations by improving cross-view consistency without explicit ground-truth 3D supervision. Further qualitative comparisons provided in Figure 5 support these quantitative results.

Depth Estimation and Semantic Segmentation.

We demonstrate the transferability of our distilled VLM features via linear probing on monocular depth estimation and semantic segmentation tasks after fine-tuning on ScanNet++. Although traditionally 2D-oriented, performance on these tasks heavily relies on robust 3D geometric understanding (Yue et al., 2024). We measure depth prediction accuracy with RMSE and absolute relative error (Rel.), and semantic segmentation using mIoU and mAcc. As shown in Table 3, FiT3D significantly improves both tasks but requires approximately three days of training on four NVIDIA A6000 GPUs due to costly 3D Gaussian optimization across training scenes. MEF shows marginal improvements over baseline CLIP, indicating limited effectiveness for dense predictions. Our approach achieves the best depth estimation performance, reducing RMSE from **0.432** to **0.367** on ScanNet++, and obtains competitive semantic segmentation results while requiring up to 54 times less computation than FiT3D on a single GPU. Without explicit dense 3D optimization, our method effectively injects robust depth priors into VLMs, enhancing semantic scene understanding.

Table 3: Quantitative comparison with linear probing on depth estimation and semantic segmentation.

Method	Train Time (↓)	ScanN	et++	KIT	TI	ADI	E20K	VOC2012	
	***	RMSE (↓)	Rel. (↓)	RMSE (↓)	Rel. (↓)	mIoU (†)	mAcc (†)	mIoU (†)	mAcc (†)
(Vanilla) CLIP	_	0.432	0.317	3.946	0.150	40.11	55.75	76.44	89.42
FiT3D (Yue et al., 2024)	\sim 3 d	0.394	0.278	3.542	0.125	42.53	56.61	79.21	90.25
MEF (You et al., 2024)	$\sim 1\mathrm{h}$	$\overline{0.429}$	0.312	3.891	0.145	40.16	55.93	76.47	89.46
Ours	\sim 1 h 20 m	0.367 (-0.065)	0.260 (-0.057)	3.529 (-0.417)	0.117 (-0.033)	41.86 (+1.75)	57.01 (+1.26)	78.74 (+2.30)	90.41 (+0.99)

Table 4: Comparison of 3D vision-language reasoning on SQA3D and ScanQA.

Method			SQA3D		ScanQA					
	EM-1	BLEU-1	METEOR	ROUGE	CIDEr	EM-1	BLEU-1	BLEU-4	METEOR	ROUGE
(Vanilla) CLIP	48.1	47.3	34.6	48.6	124.5	19.6	36.4	10.7	14.4	36.0
MEF (You et al., 2024)	48.2	<u>47.4</u>	<u>34.6</u>	<u>48.7</u>	124.7	19.0	36.1	10.4	14.3	35.1
Ours	48.6 (+0.5)	47.7 (+0.4)	35.0 (+0.4)	49.0 (+0.4)	125.5 (+1.0)	20.7 (+1.1)	36.6 (+0.2)	11.6 (+0.9)	14.5 (+0.1)	36.3 (+0.3)

Table 5: Ablation study of loss components on 3D correspondence understanding after finetuning on Objaverse.

Loss	Compon	ents	Semantic Co			rrespond	lence		Video T	Tracking	Pose Estimation			
$\mathcal{L}_{ exttt{match}}$	\mathcal{L}_{depth}	\mathcal{L}_{cost}	Dif	ferent Vi	ews	Same Views			Jaccard	Avg. Pts	Accuracy within Thresholds			
	~ ucptii		0.05	0.10	0.15	0.05	0.10	0.15		8	1cm-1deg	3cm-3deg	5cm-5deg	
~	X	×	21.18	33.54	43.58	25.94	43.33	53.87	34.61	50.58	6.32	32.00	48.33	
V	V	X	24.89	38.32	49.00	31.92	<u>52.05</u>	62.88	<u>35.36</u>	53.43	8.38	42.01	60.26	
~	~	~	25.87	39.85	50.21	36.77	56.61	67.93	35.60	54.65	8.50	<u>39.30</u>	<u>57.68</u>	

4.2.2 3D Vision-Language Understanding

To evaluate whether our distilled VLM features effectively enhance 3D vision-language understanding, we conduct experiments on two representative 3D VOA benchmarks with fine-tuned CLIP features, following the evaluation protocol from Lexicon3D (Man et al., 2024). We measure performance using EM-1, BLEU, METEOR, ROUGE, and CIDEr. Among these metrics, EM-1 is particularly crucial as it directly measures the model's exact answer prediction accuracy. For fair comparisons, we fine-tune all baselines on the Objaverse dataset. As shown in Table 4, MEF does not show significant improvements over the vanilla CLIP on SQA3D and even lower performance on ScanQA. In contrast, our method consistently outperforms both CLIP and MEF across all metrics and datasets. Specifically, our approach increases EM-1 on SQA3D to 48.6%, and notably improves EM-1 on ScanQA from 19.6% to 20.7%. These results demonstrate that our fine-tuning approach provides better 3D visual understanding which effectively leads to improvement of 3D spatial knowledge for vision-language reasoning.

4.3 Ablation Study

We conduct an ablation study to analyze the effectiveness of each loss component for 3D correspondence understanding as in Section 4.2.1 after fine-tuning on Objaverse. Compared to fine-tuning solely with \mathcal{L}_{match} equivalent to MEF, adding \mathcal{L}_{depth} consistently improves performance across

all metrics. Incorporating $\mathcal{L}_{\text{cost}}$ further boosts PCK@0.05 by +4.69% and video tracking position accuracy by +4.07%. Although pose estimation accuracy slightly decreases at some thresholds, it maintains improved performance with a gain of +2.18% at the challenging 1cm-1deg threshold. These results demonstrate that $\mathcal{L}_{\text{depth}}$ significantly enhances semantic matching and precise localization, while cost $\mathcal{L}_{\text{cost}}$ further strengthens cross-view feature consistency. Additional ablation analyses are provided in the appendix.

5 Conclusion

We present Geometric Distillation, a lightweight and annotation-free framework that enhances 3D spatial awareness and reasoning in VLMs. By distilling rich geometric signals such as multiview correspondences, relative depth relations, and dense cost volumes from high-capacity 3D foundation models like MASt3R and VGGT, our method equips pretrained 2D VLMs with robust 3D perception. Without requiring architectural modifications or explicit 3D annotations, our approach improves state-of-the-art results across diverse spatial reasoning tasks, including semantic correspondence, depth estimation, and 3D visual question answering. Extensive experiments demonstrate that our method consistently outperforms prior approaches while offering greater scalability and generalization to real-world scenes. Our work highlights an effective pathway to bridge the gap between 2D vision-language understanding and 3D perception.

6 Limitations & Future Work

While our approach achieves notable improvements in 3D spatial reasoning for vision-language models without requiring explicit annotations or architectural changes, several limitations remain. First, the method assumes access to multi-view imagery during training, which may not always be feasible in practical applications. Second, the reliance on 3D foundation models as supervision sources introduces potential biases and limits the controllability over the distilled geometric signals. Additionally, our framework does not directly generalize to other 3D modalities such as point clouds or meshes.

Future work will focus on extending geometric distillation to monocular settings and exploring self-supervised alternatives to reduce dependence on external teacher models.

Acknowledgements

This research was supported by the Basic Science Research Program through the National Research Foundation of Korea (NRF), funded by the MSIP (RS-2025-00520207, RS-2023-00219019), KEIT grant funded by the Korean government (MOTIE) (No. 2022-0-00680, No. 2022-0-01045), Artificial Intelligence Graduate School Program (KAIST) (RS-2019-II190075), and SAMSUNG Research, Samsung Electronics Co., Ltd.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Honggyu An, Jinhyeon Kim, Seonghoon Park, Jaewoo Jung, Jisang Han, Sunghwan Hong, and Seungryong Kim. 2024. Cross-view completion models are zero-shot correspondence estimators. *arXiv preprint arXiv:2412.09072*.
- Dylan Auty and Krystian Mikolajczyk. 2023. Learning to prompt clip for monocular depth estimation: Exploring the limits of human language. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2039–2047.
- Daichi Azuma, Taiki Miyanishi, Shuhei Kurita, and Motoaki Kawanabe. 2022. Scanqa: 3d question answering for spatial scene understanding. In *proceedings* of the IEEE/CVF conference on computer vision and pattern recognition, pages 19129–19139.
- Andrew Brown, Weidi Xie, Vicky Kalogeiton, and Andrew Zisserman. 2020. Smooth-ap: Smoothing the

- path towards large-scale image retrieval. In *European conference on computer vision*, pages 677–694. Springer.
- Boyuan Chen, Zhuo Xu, Sean Kirmani, Brain Ichter, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. 2024. Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14455–14465.
- Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. 2022. Adapt-former: Adapting vision transformers for scalable visual recognition. *Advances in Neural Information Processing Systems*, 35:16664–16678.
- Wei Chen, Tie-Yan Liu, Yanyan Lan, Zhi-Ming Ma, and Hang Li. 2009. Ranking measures and loss functions in learning to rank. *Advances in Neural Information Processing Systems*, 22.
- Weifeng Chen, Zhao Fu, Dawei Yang, and Jia Deng. 2016. Single-image depth perception in the wild. Advances in neural information processing systems, 29.
- Duolikun Danier, Mehmet Aygün, Changjian Li, Hakan Bilen, and Oisin Mac Aodha. 2024. Depthcues: Evaluating monocular depth perception in large vision models. *arXiv preprint arXiv:2411.17385*.
- Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. 2023. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13142–13153.
- Carl Doersch, Ankush Gupta, Larisa Markeeva, Adria Recasens, Lucas Smaira, Yusuf Aytar, Joao Carreira, Andrew Zisserman, and Yi Yang. 2022. Tap-vid: A benchmark for tracking any point in a video. *Advances in Neural Information Processing Systems*, 35:13610–13626.
- David Eigen, Christian Puhrsch, and Rob Fergus. 2014. Depth map prediction from a single image using a multi-scale deep network. *Advances in neural information processing systems*, 27.
- Mohamed El Banani, Amit Raj, Kevis-Kokitsi Maninis, Abhishek Kar, Yuanzhen Li, Michael Rubinstein, Deqing Sun, Leonidas Guibas, Justin Johnson, and Varun Jampani. 2024. Probing the 3d awareness of visual foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21795–21806.
- Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. 2015. The pascal visual object classes challenge: A retrospective. *International journal of computer vision*, 111:98–136.

- Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. 2018. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2002–2011.
- Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. 2024. Clip-adapter: Better vision-language models with feature adapters. *International Journal of Computer Vision*, 132(2):581–595.
- Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. 2013. Vision meets robotics: The kitti dataset. *The international journal of robotics research*, 32(11):1231–1237.
- Bumsub Ham, Minsu Cho, Cordelia Schmid, and Jean Ponce. 2016. Proposal flow. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3475–3484.
- Xingyi He, Jiaming Sun, Yuang Wang, Di Huang, Hujun Bao, and Xiaowei Zhou. 2022. Onepose++: Keypoint-free one-shot object pose estimation without cad models. *Advances in Neural Information Processing Systems*, 35:35103–35115.
- Deepti Hegde, Jeya Maria Jose Valanarasu, and Vishal Patel. 2023. Clip goes 3d: Leveraging prompt tuning for language grounded 3d recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2028–2038.
- Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 2023. 3d-llm: Injecting the 3d world into large language models. Advances in Neural Information Processing Systems, 36:20482–20494.
- Ian P Howard and Brian J Rogers. 1995. *Binocular vision and stereopsis*. Oxford University Press, USA.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR.
- Amita Kamath, Jack Hessel, and Kai-Wei Chang. 2023. What's" up" with vision-language models? investigating their struggle with spatial reasoning. *arXiv* preprint arXiv:2310.19785.
- Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1.

- Seoyeon Kim, Minguk Kang, Dongwon Kim, Jaesik Park, and Suha Kwak. 2023. Extending clip's image-text alignment to referring image segmentation. *arXiv preprint arXiv:2306.08498*.
- Michael S Landy, Laurence T Maloney, Elizabeth B Johnston, and Mark Young. 1995. Measurement and modeling of depth cue combination: in defense of weak fusion. *Vision research*, 35(3):389–412.
- Janghyeon Lee, Jongsuk Kim, Hyounguk Shon, Bumsoo Kim, Seung Hwan Kim, Honglak Lee, and Junmo Kim. 2022. Uniclip: Unified framework for contrastive language-image pre-training. Advances in Neural Information Processing Systems, 35:1008–1019
- Vincent Leroy, Yohann Cabon, and Jérôme Revaud. 2024. Grounding image matching in 3d with mast3r. In *European Conference on Computer Vision*, pages 71–91. Springer.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. Blip-2: Bootstrapping language-image pretraining with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pretraining for unified vision-language understanding and generation. In *International conference on machine learning*, pages 12888–12900. PMLR.
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705.
- Siting Li, Pang Wei Koh, and Simon Shaolei Du. 2024. On erroneous agreements of clip image embeddings. *arXiv preprint arXiv:2411.05195*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101.
- Xiaojian Ma, Silong Yong, Zilong Zheng, Qing Li, Yitao Liang, Song-Chun Zhu, and Siyuan Huang. 2022. Sqa3d: Situated question answering in 3d scenes. *arXiv preprint arXiv:2210.07474*.
- Yunze Man, Shuhong Zheng, Zhipeng Bao, Martial Hebert, Liangyan Gui, and Yu-Xiong Wang. 2024. Lexicon3d: Probing visual foundation models for complex 3d scene understanding. *Advances in Neural Information Processing Systems*, 37:76819–76847.
- María Luisa Menéndez, Julio Angel Pardo, Leandro Pardo, and María del C Pardo. 1997. The jensenshannon divergence. *Journal of the Franklin Institute*, 334(2):307–318.

- Songyou Peng, Kyle Genova, Chiyu Jiang, Andrea Tagliasacchi, Marc Pollefeys, Thomas Funkhouser, and 1 others. 2023. Openscene: 3d scene understanding with open vocabularies. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 815–824.
- Zengyi Qin, Jinglu Wang, and Yan Lu. 2019. Monogrnet: A geometric reasoning network for monocular 3d object localization. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 8851–8858.
- Congpei Qiu, Yanhao Wu, Wei Ke, Xiuxiu Bai, and Tong Zhang. 2025. Refining clip's spatial awareness: A visual-centric perspective. *arXiv preprint arXiv:2504.02328*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, and 1 others. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR.
- Mohit Shridhar, Lucas Manuelli, and Dieter Fox. 2022. Cliport: What and where pathways for robotic manipulation. In *Conference on robot learning*, pages 894–906. PMLR.
- James T Todd and J Farley Norman. 2003. The visual perception of 3-d shape from multiple cues: Are observers capable of perceiving metric structure? *Perception & psychophysics*, 65(1):31–47.
- Shubham Tulsiani, Tinghui Zhou, Alexei A Efros, and Jitendra Malik. 2017. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2626–2634.
- Narek Tumanyan, Assaf Singer, Shai Bagon, and Tali Dekel. 2024. Dino-tracker: Taming dino for self-supervised point tracking in a single video. In *European Conference on Computer Vision*, pages 367–385. Springer.
- Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. 2025. Vggt: Visual geometry grounded transformer. *arXiv preprint arXiv:2503.11651*.
- Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. 2024. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709.
- Philippe Weinzaepfel, Vincent Leroy, Thomas Lucas, Romain Brégier, Yohann Cabon, Vaibhav Arora, Leonid Antsfeld, Boris Chidlovskii, Gabriela Csurka, and Jérôme Revaud. 2022. Croco: Self-supervised pre-training for 3d vision tasks by cross-view completion. *Advances in Neural Information Processing Systems*, 35:3502–3516.

- Philippe Weinzaepfel, Thomas Lucas, Vincent Leroy, Yohann Cabon, Vaibhav Arora, Romain Brégier, Gabriela Csurka, Leonid Antsfeld, Boris Chidlovskii, and Jérôme Revaud. 2023. Croco v2: Improved cross-view completion pre-training for stereo matching and optical flow. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17969–17980.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and 1 others. 2019. Huggingface's transformers: State-of-the-art natural language processing. *arXiv* preprint arXiv:1910.03771.
- Ke Xian, Jianming Zhang, Oliver Wang, Long Mai, Zhe Lin, and Zhiguo Cao. 2020. Structure-guided ranking loss for single image depth prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 611–620.
- Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. 2023. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12–22.
- Yang You, Yixin Li, Congyue Deng, Yue Wang, and Leonidas Guibas. 2024. Multiview equivariance improves 3d correspondence understanding with minimal feature finetuning. *arXiv preprint arXiv:2411.19458*.
- Yuanwen Yue, Anurag Das, Francis Engelmann, Siyu Tang, and Jan Eric Lenssen. 2024. Improving 2d feature representations by 3d-aware fine-tuning. In *European Conference on Computer Vision*, pages 57–74. Springer.
- Yan Zeng, Xinsong Zhang, and Hang Li. 2021. Multi-grained vision language pre-training: Aligning texts with visual concepts. arXiv preprint arXiv:2111.08276.
- Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. 2022a. Pointclip: Point cloud understanding by clip. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8552–8562.
- Renrui Zhang, Ziyao Zeng, Ziyu Guo, and Yafeng Li. 2022b. Can language understand depth? In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 6868–6874.
- Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. 2019. Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision*, 127:302–321.
- Xiangyang Zhu, Renrui Zhang, Bowei He, Ziyu Guo, Ziyao Zeng, Zipeng Qin, Shanghang Zhang, and Peng Gao. 2023. Pointclip v2: Prompting clip and

gpt for powerful 3d open-world learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2639–2650.

Daniel Zoran, Phillip Isola, Dilip Krishnan, and William T Freeman. 2015. Learning ordinal relationships for mid-level vision. In *Proceedings of the IEEE international conference on computer vision*, pages 388–396.

Appendix Contents

· A. Potential Risks

• B. Use or Create Scientific Artifacts

- B.1 Discuss The License For Artifacts
- B.2 Documentation of Artifacts
- B.3 Statistics for Dataset

• C. Computational Experiments

- C.1 Model Size and Budget
- C.2 Experimental Setup and Hyperparameters
- C.3 Descriptive Statistics
- C.4 Parameters for Packages

• D. Use of AI Assistants

- Information About Use Of AI Assistants

• E. Additional Quantitative Evaluation

- Feature Visualization
- More Qualitative Results
- Example Result of 3D VQA

• F. Additional Ablation Study

- Comparison of Absolute and Relative Depth Understanding
- Ablation on Loss Components with Different Training Dataset
- Comparison of MASt3R and VGGT as a Teacher Model

• G. Failure Cases

A Potential Risks

Our proposed method, Geometric Distillation, enhances vision-language models (VLMs) with 3D spatial understanding by leveraging supervision signals from pretrained 3D foundation models. While our approach is annotation-free and lightweight, there are potential risks associated with its deployment. First, since the 3D models used as teachers may contain biases learned from their own training data, such biases could be inadvertently transferred to the VLMs. Second, because our method relies on pseudo-supervision (e.g., depth maps and correspondences), inaccuracies in the geometric signals could result in incorrect spatial reasoning or degraded model performance. Finally, although our work is intended for academic and constructive use, enhanced spatial reasoning capabilities could potentially be misused in surveillance, military applications, or other ethically sensitive scenarios.

B Use or Create Scientific Artifacts

Our study builds entirely on existing resources, including publicly available pretrained models and benchmark datasets. In the following, we briefly describe the licensing status of the artifacts used and provide key statistics for the datasets involved in our experiments.

B.1 Discuss The License for Artifacts

In this work, we do not introduce new datasets, but instead make use of publicly available pretrained models and benchmarks. Specifically, we use MASt3R (Leroy et al., 2024) and VGGT (Wang et al., 2025) as geometric teacher models, which are distributed under research-friendly licenses: VGGT is released under the CC BY-NC 4.0 license, and MASt3R, DUSt3R is licensed under the CC BY-NC-SA 4.0 license. Additionally, we evaluate our method using several publicly available datasets: TAP-Vid-DAVIS (Doersch et al., 2022) (Apache 2.0), OnePose-LowTexture (He et al., 2022) (Apache 2.0), ADE20K (Zhou et al., 2019) (CC BSD 3), and Objaverse (Deitke et al., 2023) (Apache 2.0). All datasets are used strictly for non-commercial research purposes in accordance with their respective licenses or terms.

B.2 Documentation of Artifacts

All code, pretrained model checkpoints, and evaluation scripts used in this study will be publicly released upon publication. These artifacts will be hosted on a GitHub repository, accompanied by detailed documentation including installation instructions, dataset preparation scripts, and usage examples. A complete README file will be provided to ensure the reproducibility of our results. For datasets that cannot be redistributed due to licensing constraints, we include scripts and links to download them from their original sources. Our release is intended to support both reproduction and future research based on our approach.

B.3 Statistics for Dataset

We summarize the dataset statistics used in our experiments across different tasks in Table 6.

3D Correspondence Understanding. We evaluate on three benchmarks following the protocols from MEF (You et al., 2024). For semantic correspondence, we use PF-PASCAL that consists of 308 image pairs from 20 object classes, randomly shuffled in different viewpoint settings. For video

Table 6: Dataset statistics and split details for each downstream task.

Task / Dataset	Split information										
3D Correspondence	Understanding										
PF-PASCAL TAP-Vid (DAVIS) OnePose-LowTexture	20 object classes; 308 image pairs; pairs randomly shuffled (in different viewpoint settings) 30 object-centric videos; 34–104 frames per video 40 objects with two videos per object; evaluation every 10th frame										
Dense Scene Understanding											
ScanNet++ KITTI ADE20K VOC2012	Validation split — 50 scenes, 30,638 images Test split — 28 scenes, 697 images Validation split — 2,000 images Validation split — 1,449 images										
3D Vision-Language	e Understanding										
SQA3D ScanQA	over 33K question–answer pairs over 41K question–answer pairs										

tracking, we follow the protocols of (Doersch et al., 2022; Tumanyan et al., 2024) and use TAP-Vid-DAVIS, which contains 30 object-centric videos with 34–104 frames per video. For object pose estimation, we follow He et al. (2022) and evaluate on the OnePose-LowTexture dataset which comprises 40 objects, each with two videos, performing evaluations on every 10th frame.

Dense Scene Understanding. Following FiT3D (Yue et al., 2024), we perform linear probing evaluations to estimate monocular depth and semantic segmentation. For depth estimation, we use ScanNet++ (Yeshwanth et al., 2023), specifically utilizing its validation split of 50 scenes with 30,638 images. We also use KITTI (Geiger et al., 2013) to evaluate generalization performance on KITTI's test split consisting of 28 scenes and 697 images. For semantic segmentation, we follow standard protocols and evaluate on ADE20K (Zhou et al., 2019)'s validation split with 2,000 images and VOC2012 (Everingham et al., 2015)'s validation split with 1,449 images.

3D Vision-Language Understanding. We evaluate 3D visual question-answering capabilities on SQA3D (Ma et al., 2022) and ScanQA (Azuma et al., 2022), following Lexicon3D (Man et al., 2024). Both datasets contain diverse QA pairs designed to probe 3D spatial and semantic reasoning. Specifically, SQA3D comprises over 33K synthetic question—answer pairs, while ScanQA contains over 41K real-world question—answer pairs generated from ScanNet scenes.

C Computational Experiments

We conduct a series of computational experiments to evaluate the effectiveness and efficiency of our proposed method. This section outlines the scale and computational cost of our models, the training setup and hyperparameter choices, a summary of the reported evaluation metrics, and the software packages used for implementation and evaluation. Through careful design and efficient training strategies, we ensure that our method achieves strong performance while maintaining high computational efficiency.

C.1 Model Size and Budget

We utilize the CLIP (Radford et al., 2021) ViT-B/16 model as our vision-language backbone, which contains approximately 93 million parameters, closely comparable to the vanilla CLIP with about 87 million parameters. For parameter-efficient fine-tuning, we employ the Low-Rank Adaptation (LoRA) (Hu et al., 2022) technique, which takes up roughly 6 million parameters (about 6.5% of the total). All experiments are conducted on up to four NVIDIA A6000 GPUs, and our geometric distillation process takes approximately 1 hour and 20 minutes per model on a single NVIDIA A6000 GPU. Compared to prior methods such as FiT3D, which require up to three days of training on four A6000 GPUs due to costly optimizing 3D feature Gaussians for all training scene, our method significantly reduces computational cost while achieving superior performance.

C.2 Experimental Setup and Hyperparameters

We use the AdamW optimizer (Loshchilov and Hutter, 2017) with a learning rate of 1×10^{-5} , and a train LoRA for up to 500 training epochs with early-stopping across all experiments. LoRA adapters

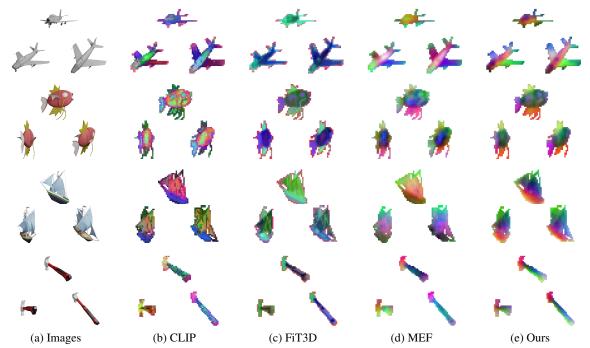


Figure 6: **Feature visualization.** PCA visualization of learned features on randomly selected 3D objects from Objaverse (Deitke et al., 2023). Compared to (b) CLIP (Radford et al., 2021), (c) FiT3D (Yue et al., 2024), and (d) MEF (You et al., 2024), our method (e) not only generates consistently smoother and more coherent features with reduced noise but also accurately preserves semantic correspondences across multiple viewpoints.

with rank r = 4 are applied to intermediate self attention layers in the CLIP model baseline. For the relative depth supervision, we add four LoRA layers to the 4th-7th attention layers, along with adapters following (Chen et al., 2022). The loss components are equally weighted: $\lambda_{\text{match}} = 1.0$, $\lambda_{\text{depth}} = 1.0$, and $\lambda_{\text{cost}} = 1.0$. Additionally, we apply temperature annealing to the cost volume alignment loss \mathcal{L}_{cost} as described in Equation (7), linearly decreasing τ from 1.0 to 0.5 during training. These hyperparameters were selected based on empirical tuning on ScanNet++ validation split and held consistent across all datasets to ensure fair comparison. We did not perform extensive hyperparameter search, and observed no significant sensitivity to small variations.

For view sampling during geometric distillation on ScanNet++, we randomly sample 10,000 views across 100 scenes, then subsequently select 100 random pairs of views that share overlapping 3D regions. This sampling results in a dataset size equivalent to the Objaverse view pairs used in MEF (You et al., 2024).

C.3 Descriptive Statistics

All results reported in the main paper and appendix represent the mean values over the full test set. For classification and tracking tasks, we use metrics such as PCK, Jaccard index, and positional accuracy at multiple thresholds. For depth estimation and semantic segmentation, we report RMSE, relative error, mIoU, and mAcc. We do not report error bars or variances, but all evaluations are deterministic and based on a single run unless otherwise specified. Our results are comparable to prior works under the same evaluation protocols and dataset splits.

C.4 Parameters for Packages

We rely on several well-established libraries and packages throughout our pipeline. For model implementation and training, we use PyTorch along with the HuggingFace Transformers (Wolf et al., 2019) and PEFT (Parameter Efficient Fine Tuning) libraries to incorporate LoRA into the CLIP backbone. For vision tasks such as depth estimation and segmentation, we use torchvision and mmsegmentation-based tools for data preprocessing and evaluation. NLP evaluation metrics including BLEU, ROUGE, METEOR, and CIDEr are computed using standard implementations from the NLTK and COCOEval toolkits. All packages are used with default parameters unless otherwise specified. No additional tuning or modification was made to external evaluation functions.

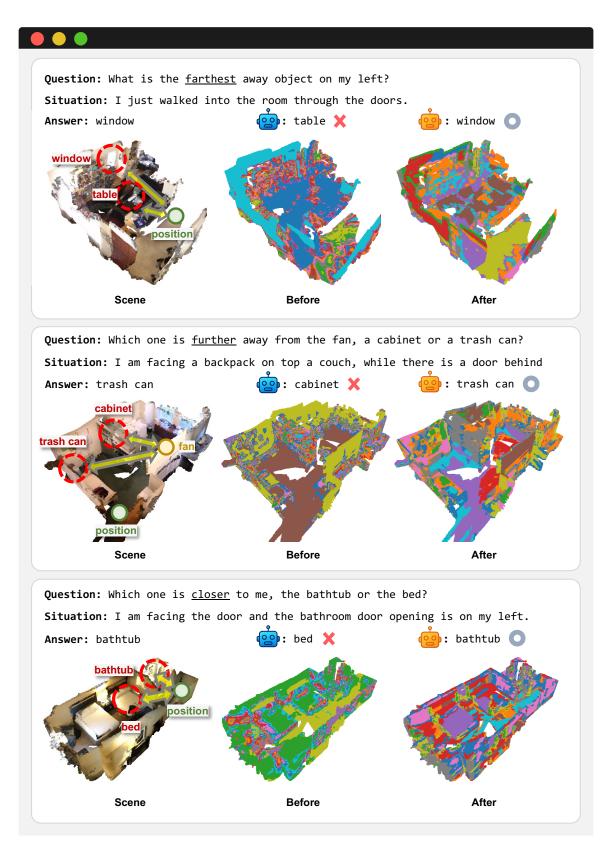


Figure 7: **Qualitative examples of 3D VQA on SQA3D.** Visualization of feature clustering for 3D scenes before and after our geometric distillation, following the protocol of Lexicon3D (Man et al., 2024). The 2D CLIP features and fine-tuned 2D CLIP features are lifted into 3D space and clustered using k-means. Each example presents a challenging VQA scenario, asking about relative object positions (e.g., "farthest," "further," "closer"). Compared to vanilla CLIP ("Before"), our distilled features ("After") offer clearer 3D spatial distinction and improved vision-language understanding for given 3D scenes.

D AI Assistants In Research Or Writing

D.1 Information About Use Of AI Assistants

We acknowledge the use of ChatGPT-4o (Achiam et al., 2023) for grammatical correction and style improvement during the writing of this paper. However, all technical content, experiment design, and conceptual development were performed solely by the authors. No AI-generated content was used for core research contributions or evaluations.

E Additional Qualitative Evaluation

E.1 Feature Visualization

To qualitatively analyze the effectiveness of our geometric distillation, we visualize PCA projections of features extracted from randomly sampled 3D objects in Objaverse. We compute a PCA between the patches of the images from the multi-view images of the same object and visualize their first 3 components. As illustrated in Figure 6, existing methods such as vanilla CLIP, FiT3D, and MEF produce noisy or inconsistent feature distributions across multiple views. In contrast, our method generates significantly smoother and more coherent feature maps that consistently preserve semantic correspondence across various viewpoints. This visualization confirms that our approach successfully injects robust multi-view geometric consistency into VLM features, which enables precise and noise-less representation of object parts and their spatial relationships.

E.2 More Qualitative Results

We provide additional qualitative comparisons for video tracking performance on the TAP-Vid-DAVIS dataset (described in Section 4.2.1) in Figure 8. Compared to MEF (You et al., 2024), our method produces notably cleaner and more accurate tracking results, which closely align with the ground-truth trajectories. Specifically, in the first row of Figure 8, MEF struggles to accurately track the trajectory of the rear wheel, confusing it with the front wheel of the car. In contrast, our approach clearly distinguishes and consistently tracks object parts. These results show that our method effectively enhances consistency to viewpoint changes and object motion.

E.3 Example Results of 3D VQA

As summarized in Figure 7. we provide example results of 3D visual question answering evaluation

on the SQA3D dataset following Section 4.2.2,. Specifically, we visualize features from vanilla CLIP and our fine-tuned CLIP obtained through geometric distillation. For visualization, we first lift the 2D CLIP features into their corresponding 3D scenes and apply k-means clustering. Our distilled features demonstrate clearer spatial coherence and improved geometric consistency compared to vanilla CLIP features. Consequently, our model exhibits superior spatial reasoning capabilities, which accurately identify relative object distances as required by challenging VQA questions, especially determining which object is farther or closer. For instance, while vanilla CLIP incorrectly identifies spatial relationships due to ambiguous feature representations, our method correctly interprets the precise spatial context, including spatially complex questions.

F Additional Ablation Study

F.1 Comparison of Absolute and Relative Depth Understanding

We perform an additional analysis comparing the effects of absolute and relative depth losses on 3D correspondence understanding. Specifically, we fine-tune models on ScanNet++ using either absolute depth loss or our proposed relative depth loss, and evaluate them across the 3D correspondence tasks described in Section 4.2.1. For absolute depth loss, we implement log-scale depth regression, which directly predicts depth values. Given predicted depth \hat{d}_p and ground-truth depth \tilde{d}_p at keypoint p for a single view, the absolute depth loss \mathcal{L}_{abs_depth} is computed as:

$$\mathcal{L}_{\text{abs_depth}} = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} |\hat{d}_p - s \cdot \tilde{d}_p|, \quad s = \frac{D_{\text{max}}^{\text{pred}}}{D_{\text{max}}^{\text{gt}}}$$

$$\tag{10}$$

where D_{\max}^{pred} and D_{\max}^{gt} denote the maximum depth from predictions and ground-truth, respectively, and s is the scale factor ensuring that predictions match the range of the ground-truth.

As shown in Table 7, the relative depth loss consistently outperforms absolute depth across all metrics. For semantic correspondence, it significantly improves PCK@0.05 from 27.04% to 28.48% (different views) and from 37.45% to 42.16% (same views). Similarly, relative depth supervision enhances video tracking, increasing the average Jaccard index from 39.27% to 40.09%, and boosts precise pose estimation accuracy at the 1cm–1deg threshold from 9.46% to 10.96%.

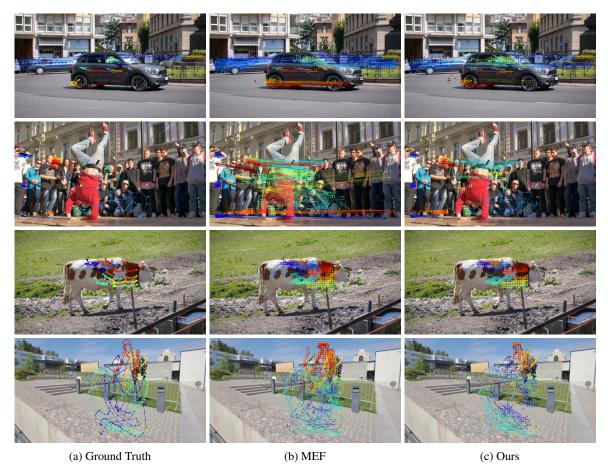


Figure 8: **Additional qualitative results on video tracking.** Visualization of predicted trajectories compared to (a) ground truth, (b) MEF (You et al., 2024), and (c) ours. Our method provides more accurate and coherent object tracking, which significantly reduces incorrect correspondences and aligns better with ground-truth trajectories.

Table 7: Absolute vs. relative depth loss in 3D correspondence understanding after fine-tuning on ScanNet++.

		Sem	antic Co	rrespond	lence		Video	Tracking	Pose Estimation			
Method	Different Views			S	Same Views			Avg. Pts	Thresholds			
	0.05	0.10	0.15	0.05	0.10	0.15	Jacc.	8	1cm-1deg	3cm-3deg	5cm-5deg	
Abs. Rel.	27.04 28.48	41.33 43.07	50.37 53.55	37.45 42.16	57.63 61.57	66.58 72.16	39.27 40.09	57.27 57.75	9.46 10.96	42.04 44.93	60.93 63.65	

These results indicate that explicitly modeling relative depth relationships, rather than absolute depth values, yields more generalizable geometric representations. Additionally, it reduces the risk of overfitting to the depth distribution of the training dataset.

F.2 Ablation on Loss Components with Different Training Dataset

To further investigate the generalization of each loss component in our geometric distillation, we conduct an additional ablation study by fine-tuning on the real-world ScanNet++ dataset, complementing our earlier analysis performed on Objaverse as in Section 4.3). Specifically, we evaluate the

effects of the matching loss \mathcal{L}_{match} , relative depth loss \mathcal{L}_{depth} , and cost volume alignment loss \mathcal{L}_{cost} across the downstream 3D correspondence tasks described in Section 4.2.1.

As shown in Table 8, adding the relative depth loss \mathcal{L}_{depth} significantly enhances semantic correspondence, increasing PCK@0.10 from 41.76% to 43.43% (different views), and improving pose estimation accuracy at the strict 1cm-1deg threshold from 9.61% to 10.80%. Incorporating the cost volume alignment loss \mathcal{L}_{cost} further strengthens performance, which yields substantial gains across most metrics. Specifically, semantic correspondence at PCK@0.05 notably increases from 26.32% to 28.48% (different views) and from

Table 8: Ablation study of loss components on 3D correspondence understanding after finetuning on ScanNet++.

Loss	Loss Components Semantic C					rrespond	lence		Video T	Tracking	Pose Estimation			
$\mathcal{L}_{ ext{match}}$ $\mathcal{L}_{ ext{depth}}$		\mathcal{L}_{cost}	Dif	ferent Vi	ews	S	ame Viev	vs	Jaccard	Avg. Pts	Accura	cy within Thr	esholds	
- materi	- асреп		0.05	0.10	0.15	0.05	0.10	0.15			1cm-1deg	3cm-3deg	5cm-5deg	
~	×	×	26.32	41.76	50.72	37.45	58.30	68.15	37.78	<u>57.45</u>	9.61	44.77	63.52	
V	~	×	27.25 28.48	43.43 43.07	52.18 53.55	38.82 42.16	60.20 61.57	69.64 72.16	38.26 40.09	56.43 57.75	10.80 10.96	47.40 44.93	64.93 63.65	

Table 9: Comparison of our VGGT and MASt3R-based methods on 3D correspondence understanding.

	Semantic Correspondence						Video	Tracking	Pose Estimation				
Method	Teacher	Dataset	Dif	Different Views			Same Views			Pos. Acc.	1cm-1deg	3cm-3deg	5cm-5deg
	1000101		0.05	0.10	0.15	0.05	0.10	0.15	Jaccard 			2 2 Daog	sem sacg
CLIP (Vanilla)	_	_	16.61	26.96	37.64	18.23	32.27	43.01	27.73	42.59	2.50	19.32	33.11
Ours (VGGT)	VGGT	Objaverse	19.84	32.79	44.24	25.44	42.48	55.18	36.77	52.68	6.94	34.37	51.83
ours (voor)	7001	ScanNet++	24.22	39.52	48.34	30.79	53.03	63.26	37.28	54.22	8.15	38.75	57.55
Ours (MASt3R)	MASt3R	Objaverse	<u>25.87</u>	<u>39.85</u>	50.21	<u>36.77</u>	56.61	<u>67.93</u>	35.60	<u>54.65</u>	8.50	39.30	<u>57.68</u>
		ScanNet++	28.48	43.07	53.55	42.16	61.57	72.16	40.09	57.75	10.96	44.93	63.65

37.45% to 42.16% (same views). Additionally, video tracking accuracy measured by the average Jaccard index improves from 37.78% to 40.09%, and pose estimation achieves the highest accuracy of 10.96% at 1cm–1deg threshold.

These results confirm that each loss component meaningfully contributes to enhancing cross-view consistency and spatial understanding. Particularly, the cost volume alignment loss $\mathcal{L}_{\text{cost}}$ improves the precision of representations, which significantly benefits performance on the most stringent evaluation metrics.

F.3 Comparison of MASt3R and VGGT as a Teacher Model

We conduct additional experiments to compare the effectiveness of different pretrained 3D foundation models, MASt3R and VGGT, used as teacher models in our geometric distillation method. Specifically, we evaluate their performance across multiple downstream 3D correspondence tasks as summarized in Table 9.

Both MASt3R and VGGT-based models substantially outperform the vanilla CLIP baseline, and this demonstrates the effectiveness of our geometric distillation approach. However, we observe consistent differences between the two teachers. Overall, MASt3R consistently generates superior results compared to VGGT, particularly when finetuned on real-world ScanNet++ data. For example, on ScanNet++, MASt3R achieves significantly better semantic correspondence accuracy (PCK@0.05 of 28.48% vs. 24.22% in different-view scenarios and 42.16% vs. 30.79% in same-view scenarios), enhanced video tracking performance (average Jaccard index 40.09% vs. 37.28%), and improved pose estimation accuracy (10.96% vs. 8.15% at

1cm-1deg threshold).

We attribute this difference in performance partly to the operational characteristics of each teacher model. Specifically, VGGT requires selecting an anchor viewpoint as user input to estimate dense correspondences across other views, so that it potentially introduces noise or inaccuracies. In contrast, MASt3R directly predicts dense and consistent semantic correspondences without requiring explicit selection of anchor points, which results in more reliable geometric guidance. Thus, while both models effectively enhance the geometric understanding of VLMs, MASt3R provides more precise and robust geometric priors in our experiments.

G Failure Cases

Although our geometric distillation method significantly enhances the VLM representations, we identify limitations under certain challenging scenarios, also shared by MEF (You et al., 2024). Specifically, our approach heavily relies on accurate geometric priors from pretrained 3D foundation models. Consequently, when input views have minimal or no overlapping 3D regions, these foundation models may fail to accurately infer or reconstruct the underlying geometry. Such failures can propagate erroneous geometric guidance into our distilled VLM features, which may degrade its performance on downstream tasks. This limitation might be alleviated through improved sampling strategies that explicitly consider shared viewing regions, as well as by enhancing the single-image 3D inference capability of the underlying 3D foundation models.

We believe that addressing these limitations is an important future direction. Potential improvements may include utilizing more powerful 3D foundation models trained on diverse, large-scale multi-view

datasets or integrating explicit uncertainty estimation to mitigate the impact of unreliable geometric guidance.