

SPATIALBOOST: ENHANCING VISUAL REPRESENTATION THROUGH LANGUAGE-GUIDED REASONING

Anonymous authors

Paper under double-blind review

ABSTRACT

Despite the remarkable success of large-scale pre-trained image representation models (i.e., vision encoders) across various vision tasks, they often fail to learn 3D spatial relationships between objects and backgrounds in the real world, constraining their effectiveness in various downstream applications. We attribute this to the limited availability of large-scale 3D training data, which makes it difficult for current image representation learning approaches to learn spatial relationships. This motivates the need for learning paradigms that rely on strong supervision while requiring less data. To address this, we propose a novel learning framework that enhances the spatial awareness of existing pre-trained vision encoders by injecting dense 3D spatial knowledge expressed in linguistic forms. To be specific, the core idea involves converting dense 3D spatial information from 2D images into linguistic expressions, which is then used to inject such spatial knowledge into vision encoders through a Large Language Model (LLM). To this end, we adopt a multi-turn Chain-of-Thought (CoT) reasoning process that progressively incorporates dense spatial knowledge and builds hierarchical spatial understanding. To validate effectiveness, we adapt SpatialBoost to state-of-the-art vision encoders such as DINOv3, and evaluate its performance gains on a wide range of benchmarks requiring both 3D perception and general vision abilities.

1 INTRODUCTION

Pre-trained image representation models (He et al., 2020; Donahue & Simonyan, 2019; Chen et al., 2020b; Dosovitskiy et al., 2021; Li et al., 2023b; Assran et al., 2023) have shown remarkable success in various downstream tasks, such as image classification (Krizhevsky et al., 2009; Cui et al., 2018), semantic segmentation (Lin et al., 2014; Zhou et al., 2019), monocular depth prediction (Silberman et al., 2012; Geiger et al., 2012), and vision-language understanding (Antol et al., 2015; Hudson & Manning, 2019). The core idea behind these successes is extracting transferrable representation from large-scale image datasets such as ImageNet (Deng et al., 2009), enabling the model to understand semantic information within images that is significantly useful for various downstream tasks.

Despite their success, these models are predominantly trained on 2D images and hence face a fundamental challenge in acquiring 3D spatial awareness capabilities. Consequently, large vision language models struggle to discern 3D spatial relationships between objects in images (Liu et al., 2023a; Fu et al., 2024b; Wang et al., 2025b; Cheng et al., 2024), and demonstrate sub-optimal performance in vision-based robotic control tasks compared to approaches that directly utilize 3D information (Ze et al., 2024; Ke et al., 2024; Zhen et al., 2024). To address these limitations, several works train vision models on multi-view images that naturally encode spatial information (Zhang et al., 2024; Wang et al., 2024b; Charatan et al., 2024). While these approaches have shown promise in robot control tasks (Seo et al., 2023; Sermanet et al., 2018), their broader applicability remains constrained by the need to use carefully curated data (Yu et al., 2023) or obtain multi-view datasets from simulation environments (Savva et al., 2019), creating significant limitations for scaling up these approaches. These challenges highlight the need for a novel framework that enables effective learning of 3D information with substantially less data.

However, we note that vision models specialized for individual tasks are able to infer object positions and point depths from standard 2D images. These extracted cues make it possible to extend spatial information by modeling geometric relationships between objects in a scene. We hypothesize that

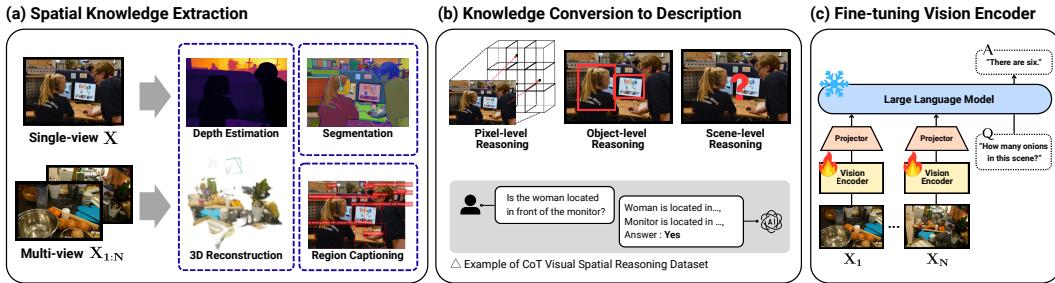


Figure 1: **Overview of SpatialBoost.** We enhance spatial and geometric understanding of pre-trained vision encoders by leveraging language-guided spatial reasoning. SpatialBoost consists of (a) spatial knowledge extraction through depth estimation, 3D reconstruction, segmentation, and region captioning, (b) converting spatial knowledge into multi-turn spatial reasoning from pixel to scene levels, and (c) building a spatial-aware vision encoder with LLM using generated data in (b).

such spatial information can be systematically converted into explicit representations by leveraging language. Moreover, since language naturally composes information in a sequential and structured form, this property allows the construction of labels that capture dense spatial relationships within a scene.

Based on these insights, we introduce SpatialBoost, a training framework that enhances the spatial understanding of pre-trained vision encoders by leveraging language-guided reasoning (see Figure 1). We inject linguistically described spatial knowledge through decoder-based fine-tuning with Large Language Models (LLM), where the model takes single or multi-view images as input and generates descriptions. In particular, to leverage this knowledge without forgetting the existing knowledge, we incorporate additional learnable parameters (*i.e.*, dual-channel attention module) into the vision encoder and train only them while freezing the existing parameters. Furthermore, to incorporate dense spatial information in a structured manner, we present a multi-turn visual spatial reasoning approach that builds hierarchical spatial understanding through pixel-level, object-level, and scene-level sub-questions and answers.

To validate the effectiveness of our method, we apply SpatialBoost to state-of-the-art image encoders, including DINOv3 (Siméoni et al., 2025) and SigLIPv2 (Tschanne et al., 2025), and evaluate them across a diverse set of vision tasks: monocular depth estimation, semantic segmentation, 3D scene understanding, vision-based robotic control, image classification, image retrieval, spatial reasoning, and general VQA.¹ Our experiment first shows that SpatialBoost consistently improves performance on tasks requiring 3D spatial knowledge. For example, on the 3D scene understanding task, SpatialBoost improves DINOv3 by 3.5% (51.4% \rightarrow 54.9%) on the SQA3D task from Lexicon3D Benchmark (Man et al., 2024). In addition, on depth estimation tasks, SpatialBoost improves SigLIPv2 from an RMSE score of 0.51 to 0.39 on NYUd linear probing. Moreover, we show that SpatialBoost even improves the performance of the vision encoders across all benchmarks, notably in image classification: SpatialBoost improves ImageNet linear probing performance of DINOv3 from 88.4% to 90.2%.

2 RELATED WORK

Self-supervised Learning for Image Representation. In earlier years, most approaches relied on supervised learning with large-scale labeled datasets to train models (Deng et al., 2009; Simonyan & Zisserman, 2014; Szegedy et al., 2014; He et al., 2016). However, the dependence on annotated data introduced scalability challenges due to label expense. To address this, self-supervised learning (SSL) has emerged as a dominant paradigm, leveraging unlabeled data to learn image representations. Contrastive learning methods, including SimCLRv2 (Chen et al., 2020c), MoCov3 (Chen et al., 2021), DINOv2 (Oquab et al., 2023), and iBOT (Zhou et al., 2021), are trained to distinguish between representations of augmented views of the same image and those of different images. Concurrently, mask prediction approaches such as BEiT (Bao et al., 2021) and MAE (He et al., 2022),

¹Due to space constraints, results on spatial reasoning and general VQA tasks are provided in the appendix.

108 learn representations by reconstructing masked portions of input images. While these methods excel
 109 at capturing rich semantic features within 2D images, they lack mechanisms to effectively encode
 110 3D spatial knowledge. On the other hand, we overcome this limitation by enhancing image repre-
 111 sentations through a novel method that injects 3D spatial knowledge by utilizing language decoding.
 112

113 **Multi-modal Learning for Image Representation.** The increasing prominence of multi-modal
 114 tasks has catalyzed the development of vision-language models that jointly represent visual and
 115 textual information. These models typically employ weakly supervised learning by leveraging text
 116 caption. Contrastive learning schemes, *e.g.*, CLIP (Radford et al., 2021), SigLIP (Zhai et al., 2023)
 117 and OpenCLIP (Cherti et al., 2023), consist of vision and text encoders and are trained to align their
 118 representations in a shared embedding space. Alternative methodologies like M3AE (Geng et al.,
 119 2022), jointly encode image patches and text tokens, employing masked prediction objectives to
 120 reconstruct both modalities. More recently, autoregressive formulations such as iGPT (Chen et al.,
 121 2020b), have emerged, treating image patches and text tokens as sequential elements for predictive
 122 modeling. These approaches successfully enrich visual representations with semantic context de-
 123 rived from natural language descriptions. However, existing models necessitate joint pre-training of
 124 both modalities from scratch, imposing significant computational demands and preventing efficient
 125 adaptation of existing pre-trained models. Our method eliminates the need for joint text-image rep-
 126 resentation learning by using LLM, thereby enhancing pre-trained models with relevant linguistic
 127 information efficiently.

128 **Multi-View Learning for Image Representation.** Recent advances in vision tasks that require 3D
 129 spatial understanding and generation have increased the demand for effective 3D spatial represen-
 130 tations (Chen et al., 2024b; Wu et al., 2024; Goyal et al., 2023; Shridhar et al., 2023). Multi-view
 131 images from different camera viewpoints or video sequences serve as input for these tasks. Our
 132 focus is specifically on augmenting image representations with useful 3D information. Typically,
 133 following approaches similar to single-view image representation learning, multi-view data has been
 134 processed by converting images into patches for masked prediction such as MV-MWM (Seo et al.,
 135 2023) or through contrastive learning methods (Sermanet et al., 2018). Additionally, to learn 3D-
 136 related information more explicitly, approaches that predict 3D features from image representation
 137 (Ke et al., 2024; Gervet et al., 2023; Ze et al., 2024) have been proposed. These approaches have led
 138 to significant performance improvements in vision-based robot control. However, such methods are
 139 limited by multi-view data, making it difficult to develop them into pre-trained models for general
 140 3D understanding. Our approach proposes a method to learn 3D spatial representations from both
 141 single-view and multi-view images, avoiding these limitations.

3 METHOD

144 In this section, we introduce SpatialBoost, a visual representation learning framework designed to
 145 improve vision encoders by injecting 3D spatial information expressed in natural language. We first
 146 present a multi-modal architecture that incorporates linguistically expressed visual information into
 147 the vision encoder through a dual-channel attention layer, ensuring that original visual features are
 148 preserved while 3D spatial information is fully exploited (see Section 3.1). On top of this archi-
 149 tecture, we design a Visual-Question-Answering (VQA) dataset that hierarchically disentangles 3D
 150 spatial relations from both single/multi-view images, enabling the vision encoder to learn spatial
 151 information more effectively (see Figure 1).

3.1 TRAINING PIPELINE

152 To train a vision encoder from rich spatial information encoded in large-scale linguistic expressions,
 153 our key idea is to utilize Large-Language Models (LLM) by constructing a multi-modal architecture
 154 composed of a vision encoder f_V , a trainable projection module g_P , and the LLM f_L . However,
 155 without proper alignment between visual and textual representations, the training signals from the
 156 LLM cannot effectively propagate back to the vision encoder, making the learning process ineffec-
 157 tive. To fully exploit language supervision, we begin by aligning the visual encoder with the textual
 158 embedding space of the LLM. Specifically, we adopt LLaVA (Liu et al., 2023b), a two-stage train-
 159 ing for the alignment: feature alignment (Stage 1) and visual instruction tuning (Stage 2). After
 160 the alignment, we introduce a training framework that uses a language-guided reasoning dataset to
 161

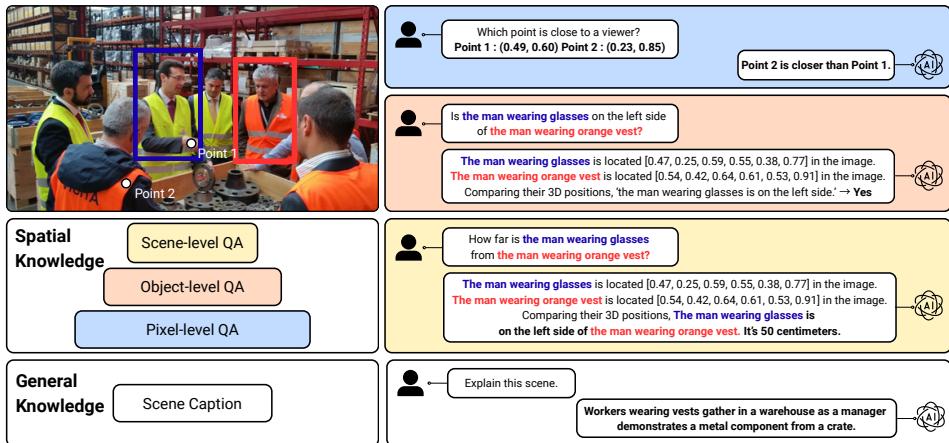


Figure 2: **Illustration of multi-turn visual spatial reasoning dataset**, exhibiting pixel-level, object-level, and scene-level reasoning QAs. At the pixel-level, the QA task queries the 3D positions of points (e.g., via depth estimation). At the object-level, it extracts spatial properties of objects (e.g., by predicting bounding cubes or relative positions). At the scene-level, it determines the exact distances between multiple objects that require the rationales of the previous steps. At last, we add 2-turn for general scene caption. These are listed in order and constitute 12 multi-turn visual spatial reasoning conservation.

fine-tune the vision encoder (Stage 3). Notably, direct full fine-tuning in this final stage would lead to catastrophic forgetting of the pre-trained knowledge embedded in the vision encoder. To address this challenge, we introduce *dual-channel attention* layers that enable the model to acquire spatial understanding while preserving its original representational capabilities.

Formally, given an input image \mathbf{x} and multi-turn conversation data $(\mathbf{x}_q^1, \mathbf{x}_a^1, \dots, \mathbf{x}_q^T, \mathbf{x}_a^T)$ from question-answering (QA) pairs (Q_x, A_x) , we first encode \mathbf{x} to obtain visual features $\mathbf{z}_v = f_V(\mathbf{x})$, which are mapped into the token embedding space via $g_P(\mathbf{z}_v)$. These visual tokens are then concatenated with text tokens and fed into the LLM. Given the multi-turn conversation data and input image, we optimize the model through autoregressive loss. Our training pipeline consists of three stages and all stages are trained with supervised fine-tuning (SFT) loss. We describe each stage in the following paragraphs.

Stage 1: Feature alignment. In this stage, we train a projector g_P that maps image features into the textual embedding space of the LLM. This projector pre-training contributes to the stable vision-language alignment. Following the training setup in multi-modal large language models (Liu et al., 2023a; 2024a), we freeze the parameters of both the visual encoder f_V and the language model f_L , and optimize only the projector g_P .

Stage 2: Visual instruction tuning. Following the projector alignment in Stage 1, this stage extends the alignment to the LLM. We freeze the visual encoder f_V and fine-tune the projector g_P and the language model f_L using our multi-view VQA data, combined with the single-view visual instruction data from LLaVA (Liu et al., 2023a). This step enables f_L and g_P to handle multi-view visual questions. We provide details of proposed multi-view VQA data in Section 3.2.

Stage 3: Vision encoder fine-tuning with dual-channel attention. Finally, we fine-tune the vision encoder f_V to have the capability of spatial understanding. To effectively inject dense spatial knowledge into the vision encoder, we use multi-turn visual spatial reasoning dataset (see Section 3.2), which is carefully designed for hierarchical spatial reasoning. We train the vision encoder f_V and

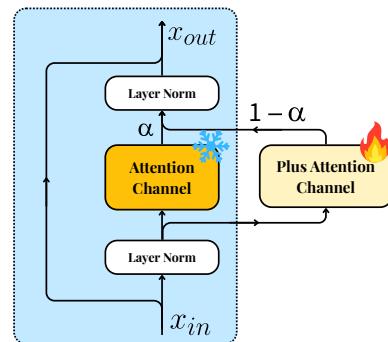


Figure 3: **Illustration of the dual-channel attention layer** (Hong et al., 2023a), where an additional attention block is introduced alongside the original attention block and merged via a learnable mixture factor α .

216 **Table 1: Results on monocular depth estima-
217 tion** from NYUd (Silberman et al., 2012) and
218 KITTI (Geiger et al., 2013) benchmarks. We re-
219 port the RMSE score between ground truth and
220 re-predicted depth values. Lower is better. For all
221 results, we freeze the encoder backbone and train
222 a linear head (lin.) or DPT head (Ranftl et al.,
223 2021) on top of the image features of the last
224 layer.

Method	NYUd		KITTI	
	lin.	DPT	lin.	DPT
OpenCLIP	0.53	0.41	3.54	2.70
+SpatialBoost (Ours)	0.40	0.38	2.79	2.54
SigLIPv2	0.51	0.40	3.32	2.64
+SpatialBoost (Ours)	0.39	0.34	2.71	2.50
DINOv2	0.37	0.29	2.60	2.11
+SpatialBoost (Ours)	0.30	0.25	2.53	2.07
DINOv3	0.31	0.25	2.33	2.02
+SpatialBoost (Ours)	0.25	0.21	2.20	1.84

Table 2: Results on semantic segmentation from ADE20K (Zhou et al., 2017) and Pascal VOC (Everingham et al., 2010) benchmarks. We report mIoU score. Higher is better. For all results, we freeze the encoder backbone and report results of linear probing (lin.) or multi-scale evaluation (+ms), where the multi-scale approach uses features from the last four layers of the visual encoder to perform segmentation.

Method	ADE20K		Pascal VOC	
	lin.	+ms	lin.	+ms
OpenCLIP	39.5	46.0	71.7	79.3
+SpatialBoost (Ours)	40.5	47.3	75.1	80.9
SigLIPv2	42.8	48.7	72.6	79.1
+SpatialBoost (Ours)	45.1	50.8	79.0	82.2
DINOv2	49.3	53.0	83.0	86.2
+SpatialBoost (Ours)	52.0	54.9	84.5	87.6
DINOv3	55.9	60.3	86.6	89.8
+SpatialBoost (Ours)	59.7	63.1	88.5	90.9

the projection module g_P while keeping the parameters of the LLM f_L frozen, allowing only the vision encoder to benefit from language-driven spatial information. We employ SFT loss, and through this training process, the vision encoder learns to extract meaningful representations necessary for producing answers. However, direct full fine-tuning risks forgetting of the pre-trained knowledge embedded in the vision encoder. To address this challenge, we introduce a dual-channel attention mechanism (see Figure 3). Specifically, for each attention layer $\text{Attn}(\cdot)$ in the visual encoder f_V , we introduce an additional attention layer $\text{Attn}^+(\cdot)$, whose weight parameters are initialized to the same values as those of $\text{Attn}(\cdot)$. Given an input \mathbf{x} to each attention layer, we merge the outputs of $\text{Attn}(\cdot)$ and $\text{Attn}^+(\cdot)$ by introducing a trainable mixture factor $\alpha = \text{sigmoid}(\mathbf{a}) \in (0, 1)^d$ with zero-initialized parameter $\mathbf{a} \in \mathbb{R}^d$, where d is the hidden dimension of \mathbf{x} , as follows:

$$\text{Attn}^{\text{final}}(\mathbf{x}) = \alpha \cdot \text{Attn}(\mathbf{x}) + (1 - \alpha) \cdot \text{Attn}^+(\mathbf{x}). \quad (1)$$

During fine-tuning, we only update the parameters of Attn^+ and α while keeping all other parameters frozen. This approach allows the vision encoder to initially rely on pre-trained attention weights and gradually incorporate new attention weights, smoothly enhancing spatial awareness without discarding existing knowledge (see classification result in Figure 6).

3.2 ENHANCING VISION ENCODER WITH SPATIAL COT

To effectively inject dense spatial information into vision encoders, we address the fundamental limitations of existing spatial datasets. Current spatial VQA data consist of simple single-turn QA pairs with limited information content, insufficient for transferring comprehensive 3D understanding. To overcome this limitation, We introduce Multi-view VQA, which helps align the vision encoder with the LLM to effectively handle multi-view data and a multi-turn Chain-of-Thought (CoT) framework (Wei et al., 2022) for both single-view and multi-view images that enables the injection of substantially richer spatial information in a single training instance.

Multi-view VQA Dataset. To enhance multi-view VQA capabilities during the visual instruction tuning (Stage 2), we construct multi-view VQA dataset. We first apply LPIPS (Zhang et al., 2018) metric to the 3D or video dataset to obtain a pair of images. Given the pair of images, we employ GPT-4o (Achiam et al., 2023) to generate visual questions targeting general multi-view knowledge. We provide more details in Section C.

Multi-Turn Visual Spatial Reasoning Dataset. To enhance spatial reasoning capabilities of the vision encoder (Stage 3), we construct multi-turn visual spatial reasoning dataset for single-view and multi-view. Additionally, to enhance general knowledge of the vision encoder, we append GPT-generated scene captions after spatial reasoning turn. For single-view image, we first extract a 3D point cloud from given an image \mathbf{x} by applying diverse vision models (e.g., depth estimation

270 **Table 3: Results on 3D-centric tasks.** We evaluate unified probing on diverse 3D-related tasks
 271 from ScanNet (Dai et al., 2017) scenes. We report BLEU-1 score for Vision-Language Reasoning
 272 (VLR) on ScanQA (Azuma et al., 2022) and SQA3D (Ma et al., 2023). For Visual Grounding (VG),
 273 we report accuracy on overall category of ScanRefer (Chen et al., 2020a) dataset. For Geometric
 274 Understanding (GU), we report Registration Recall (RR) at 0.05m RMSE threshold and Relative
 275 Translation Error (RTE). For 3D Semantic Understanding (3D SU), we report accuracy and mIoU.
 276 Lower is better for RTE and higher is better for all other metrics.

Method	VLR		VG	GU		3D SU	
	ScanQA \uparrow	SQA3D \uparrow		ScanRefer-Overall \uparrow	RR@0.05m (%) \uparrow	RTE (m) \downarrow	Acc \uparrow
OpenCLIP	36.9	48.0	50.1	22.6	0.40	39.8	6.9
+SpatialBoost (Ours)	39.2	49.9	56.6	78.8	0.17	76.9	54.9
SigLIPv2	38.1	48.5	51.4	47.8	0.28	47.7	9.2
+SpatialBoost (Ours)	40.8	50.1	56.8	86.4	0.15	81.0	55.5
DINOv2	39.5	49.8	52.7	82.4	0.15	83.0	64.1
+SpatialBoost (Ours)	40.3	50.4	57.0	92.4	0.13	89.8	68.3
DINOv3	40.6	51.4	56.2	86.9	0.10	91.1	69.1
+SpatialBoost (Ours)	43.3	54.9	61.1	97.5	0.06	91.9	70.6

286 model (Bochkovskii et al., 2024) and image segmentation model (Ravi et al., 2024)). For multi-
 287 view images $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, we use 3D reconstruction model (Wang et al., 2025a) to extract a 3D
 288 point cloud from given images. Using the point cloud, we synthesize QA pairs specialized in spatial
 289 reasoning about \mathbf{x} or $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$.

290 We then design spatial reasoning QA pairs at three hierarchical levels: pixel, object, and scene,
 291 enabling LLM to perform CoT reasoning from narrow to broad view. Specifically, at the pixel-
 292 level, the QA task is designed to capture the overall geometry in the image by querying the absolute
 293 or relative 3D position of a point, e.g., “What is the depth value at coordinate (x, y) ?” At the
 294 object-level, the QA task tackles the semantic spatial information of objects inside the image using
 295 a bounding cube of the object in 3D space, e.g., “Is [A] on the left side of [B]?", where [A] and
 296 [B] is the descriptions about the object in image. We note that this level uses the pixel-level spatial
 297 information as a rationale, enabling LLM to reason about the geometry of objects in 3D space.
 298 Lastly, at the scene-level, the QA task is designed to predict the exact distance between multiple
 299 objects that requires coherent 3D spatial understanding, e.g., “How far is [A] from [B]?”.

4 EXPERIMENTS

303 Through extensive experiments, we validate the performance of SpatialBoost and ablate its key
 304 components, focusing on following questions:

- 306 • Can SpatialBoost improve spatial knowledge of the vision encoder? (Tables 1 to 4)
- 307 • Isn’t SpatialBoost overfitted to spatial knowledge? (Table 5)
- 308 • Which components contribute to SpatialBoost performance? (Table 6 and Figure 6)

4.1 EXPERIMENTAL SETUP

313 **VQA Dataset Construction.** For single-view image, we use randomly sampled 100K images from
 314 the SA1B dataset (Kirillov et al., 2023) to construct the single-view VQA dataset specialized in
 315 chain-of-thought spatial reasoning. For multi-view images, we use filtered 200K samples from the
 316 ego-centric video dataset (Grauman et al., 2022) and 3D dataset (Jensen et al., 2014; Dai et al.,
 317 2017; Mildenhall et al., 2021; Barron et al., 2022) to construct multi-view VQA dataset niche in
 318 multi-view reasoning or alignment. More details in Section D.

320 **Baselines.** For all experiments, we compare our methods with the recent widely-used pre-trained
 321 image representation models. To be specific, we first consider OpenCLIP (Cherti et al., 2023) ViT-
 322 G/14 and SigLIPv2 (Tschannen et al., 2025) ViT-g/16, known for language-aligned vision encoder.
 323 We also consider DINOv2 (Oquab et al., 2023) ViT-g/14 and DINOv3 (Siméoni et al., 2025) ViT-
 7B/16, which is a recent state-of-the-art vision encoder.

324
 325 **Table 4: Results on vision-based robot learning.** We report the
 326 performance of imitation learning agents on 4 domains from Cor-
 327 texBench (Majumdar et al., 2023), which are trained upon the im-
 328 age representations. In particular, we report the normalized score
 329 for DMControl and success rates (%) for other tasks.

Method	Adroit	MetaWorld	DMControl	Trifinger	Avg.
OpenCLIP	52.6 ± 4.9	83.0 ± 2.7	58.5 ± 1.9	67.7 ± 0.5	65.5
+SpatialBoost (Ours)	61.1 ± 3.4	87.0 ± 3.3	61.0 ± 1.6	72.9 ± 0.3	70.5
SigLIPv2	56.5 ± 3.0	84.7 ± 2.9	69.4 ± 2.1	68.3 ± 0.8	69.7
+SpatialBoost (Ours)	66.5 ± 1.9	89.1 ± 0.9	73.5 ± 1.8	73.9 ± 0.7	75.8
DINOv2	55.4 ± 2.7	82.4 ± 4.0	67.9 ± 1.0	66.8 ± 0.2	68.1
+SpatialBoost (Ours)	68.1 ± 2.9	88.5 ± 3.1	75.0 ± 1.1	71.4 ± 0.8	75.8
DINOv3	63.9 ± 1.5	83.8 ± 1.6	70.8 ± 1.8	72.8 ± 0.5	72.8
+SpatialBoost (Ours)	71.8 ± 3.4	92.0 ± 1.9	80.4 ± 2.4	79.0 ± 0.6	80.8

330
 331
 332
 333
 334
 335
 336
 337
 338
 339 **Implementation Details.** We choose Qwen-2.0-7B (Yang et al., 2024) as the LLM backbone and
 340 2-layer MLP as the projector, following the architecture of LLaVA-1.5 (Liu et al., 2024a). Further
 341 details are provided in Section A.

343 4.2 DENSE PREDICTION TASKS

345
 346 **Setup.** We evaluate SpatialBoost on dense prediction tasks requiring geometric and semantic spatial
 347 understanding. For geometric understanding, we perform monocular depth estimation on NYUD
 348 (Silberman et al., 2012) and KITTI (Geiger et al., 2013) using linear or DPT (Ranftl et al., 2021)
 349 heads. For semantic understanding, we evaluate on ADE20K (Zhou et al., 2017) and Pascal VOC
 350 (Everingham et al., 2010) segmentation benchmarks using linear or multi-scale heads. All experi-
 351 ments freeze the visual backbone during training (see Section A for details).

352 **Results.** As shown in Table 1 and 2, SpatialBoost consistently improves both geometric and se-
 353 mantic spatial understanding across various encoders. For instance, OpenCLIP’s RMSE on NYUD
 354 decreases from 0.53 to 0.40 with a linear head, while DINOv3’s mIoU on ADE20K increases from
 355 55.9% to 59.7%. These consistent gains demonstrate that language-based spatial knowledge transfer
 356 effectively enhances visual encoders’ spatial understanding capabilities.

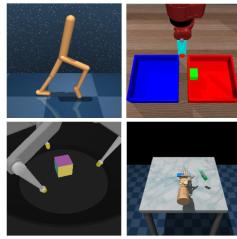
357 4.3 COMPLEX 3D-CENTRIC TASKS

360
 361 **Setup.** We evaluate SpatialBoost on Lexicon3D (Man et al., 2024), a unified benchmark for 3D
 362 scene understanding covering vision-language reasoning, visual grounding, semantic understand-
 363 ing, and geometric understanding. Following Lexicon3D protocols, we freeze visual backbones and train
 364 task-specific heads (see Section A for details).

365 **Results.** As shown in Table 3, SpatialBoost shows comprehensive improvements across diverse 3D
 366 tasks. OpenCLIP’s BLEU-1 improves from 36.9 to 39.2 on ScanQA (Azuma et al., 2022), while
 367 DINOv3 increases from 51.4 to 54.9 on SQA3D (Ma et al., 2023), demonstrating that SpatialBoost
 368 improves spatial understanding without compromising language capabilities. Notably, SigLIPv2’s
 369 3D semantic segmentation dramatically improves from 6.9 to 54.9 mIoU, highlighting SpatialBoost
 370 can inject robust spatial knowledge into encoders with initially limited spatial awareness.

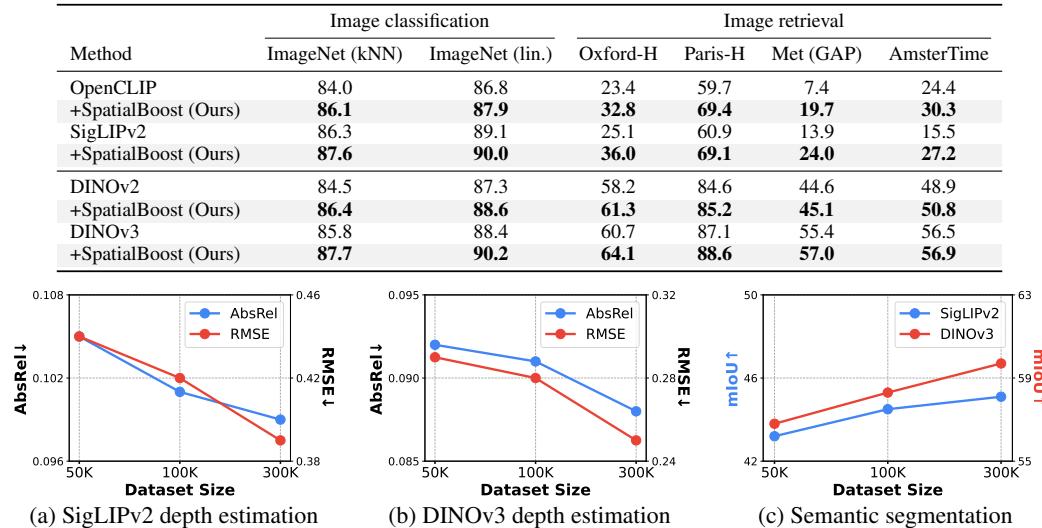
371 4.4 VISION-BASED ROBOT LEARNING

372
 373 **Setup.** We evaluate SpatialBoost on vision-based robot control using 4 domains from CortexBench
 374 (Majumdar et al., 2023) spanning locomotion and manipulation tasks (Rajeswaran et al., 2017; Yu
 375 et al., 2020; Tassa et al., 2018; Wüthrich et al., 2020). Following CortexBench protocols, we train
 376 behavior cloning agents using [CLS] representations to predict expert actions from visual observa-
 377 tions. We report the mean of best performance across 5 evaluation runs (see Section A for details).



378
 379 **Figure 4: Examples of vi-
 380 sual observations from Cor-
 381 texBench.** We train imitation
 382 learning agents to learn a map-
 383 ping from these visual observa-
 384 tions to expert actions.

378
379
380
381
382
383
Table 5: **Results on image classification and retrieval tasks.** We report Top-1 accuracy of kNN
384 performance and linear probing (lin.) for image classification on validation set of ImageNet-1K
385 (Russakovsky et al., 2015). For image retrieval, we report global average precision (GAP) on Met
386 (Ypsilantis et al., 2021) and mean average precision (mAP) on Oxford-Hard (Oxford-H) (Radenović
387 et al., 2018), Paris-Hard (Paris-H) (Radenović et al., 2018), and AmsterTime dataset (Yıldız et al.,
388 2022). For all results, we freeze the encoder backbone.
389



400
401
402
403
404
405
Figure 5: **Effect of dataset scalability.** We investigate the effect of the size of analysis of data scalability effects on (a) depth estimation results (AbsRel, RMSE) on NYUd benchmark for SigLIPv2,
406 (b) depth estimation results (AbsRel, RMSE) on NYUd benchmark for DINOv3, and (c) semantic
407 segmentation results (mIoU) on ADE20K benchmark for SigLIPv2 and DINOv3. The results show
408 scalable performance improvements with increased data size.
409

410
411
412
413
414
415
416
Results. As shown in Table 4, SpatialBoost significantly improves robot task performance across all
417 vision encoders. For example, DINOv2 + SpatialBoost achieves 68.1% on Adroit versus 55.4% for
418 DINOv2 alone, demonstrating that enhanced spatial representations directly benefit robot control.
419

4.5 IMAGE CLASSIFICATION AND RETRIEVAL TASKS

420
421
422
423
424
425
426
427
428
429
430
431
Setup. We evaluate SpatialBoost’s impact on instance recognition using ImageNet-1K (Russakovsky et al., 2015) classification and retrieval benchmarks (Oxford, Paris (Radenović et al., 2018), Met (Ypsilantis et al., 2021), AmsterTime (Yıldız et al., 2022)). Following DINOv3 protocols, we use linear probing on [CLS] representations for classification and similarity-based ranking for retrieval (see Section A for details).

432
433
434
435
436
437
438
439
440
441
442
443
Results. As shown in Table 5, SpatialBoost improves both classification and retrieval despite these
444 tasks not explicitly requiring spatial understanding. DINOv3’s ImageNet accuracy increases from
445 88.4% to 90.2%, while Oxford-Hard mAP improves from 60.7 to 64.1. These results demonstrate
446 that SpatialBoost enhances general vision capabilities without overfitting to spatial features, likely
447 due to our dual-channel attention preserving pre-trained knowledge and the inclusion of general
448 scene captions alongside spatial reasoning.

4.6 ABLATION STUDY AND ANALYSIS

449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
20

432
 433 **Table 6: Effect of LLM-based fine-tuning.** We fine-tune
 434 the vision encoder with different headers. We report ac-
 435 curacy (%) for classification (Cls) on ImageNet-1K, mIoU
 436 for segmentation (Seg) on ADE20K, RMSE for depth esti-
 437 mation on NYUd, and BLEU-1 score for vision-language
 438 reasoning (VLR) on ScanQA. We use ViT-L/14 as the
 439 backbone architecture of the encoder.

Method	Cls ↑	Seg ↑	Depth ↓	VLR ↑
DINOv2	86.3	47.7	0.38	39.2
+Linear (depth)	85.7 (-1.39%)	47.9 (+0.42%)	0.35 (-7.89%)	36.9 (-5.87%)
+Linear (seg.)	86.6 (+0.35%)	48.8 (+2.31%)	0.45 (+18.42%)	37.1 (-5.36%)
+SAM decoder	86.3 (+0.0%)	50.1 (+5.03%)	0.42 (+10.53%)	37.6 (-4.08%)
+VGGT decoder	84.8 (-1.74%)	45.6 (-4.40%)	0.35 (-7.89%)	37.3 (-4.85%)
+LLM (Ours)	88.3 (+2.32%)	51.5 (+7.97%)	0.32 (-15.79%)	40.0 (+2.04%)

444
 445 **Table 7: Component-wise analysis.** We investigate the effect of multi-turn spatial reasoning data
 446 and the effect of single-view and multi-view data. Multi-turn order means the order of three levels
 447 (*i.e.*, pixel, object, and scene) in our visual spatial reasoning data.

Method	Multi-turn order	Single-view data	Multi-view data	Cls ↑	Seg ↑	Depth ↓
DINOv2	✗	-	-	86.3	47.7	0.38
+SpatialBoost	Reverse	+100K	-	87.4	48.4	0.35
	Random	+100K	-	87.4	48.5	0.36
	Forward	+100K	-	87.6	48.9	0.34
	Forward	-	+100K	87.6	48.2	0.36
	Forward	+50K	+50K	87.6	49.2	0.32

455
 456 **Effect of Multi-turn Visual Reasoning.** In Table 7, we investigate how the hierarchical structure of
 457 reasoning affects representation learning. We compare dataset construction strategies: (a) shuffled
 458 multi-turn, (b) reversed order (scene→object→pixel), and (c) forward order (pixel→object→scene).
 459 The forward hierarchical ordering shows optimal performance, demonstrating that reasoning order
 460 significantly impacts the quality of representation.

461 **Effect of Single-view and Multi-view Data.** In Table 7, we investigate the effect of single-view and
 462 multi-view reasoning data. With fixed total samples, we compare single-view only, multi-view only,
 463 and combined training. While both data types independently improve performance, the combination
 464 achieves the highest results, confirming their complementary nature.

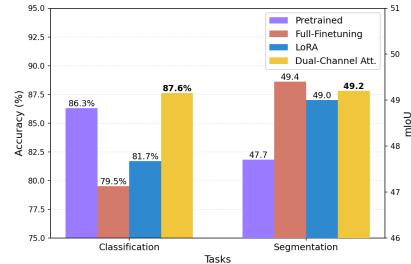
465
 466 **Comparison with Naive Post-training.** In Table 8, we investigate the effect of post-training. With
 467 fixed total samples (*i.e.*, 300K data in multi-turn reasoning data), we compare the naive post-training
 468 scheme and SpatialBoost. We evaluate the performance of the vision encoder across five tasks: depth
 469 estimation, segmentation, vision-language reasoning, robot learning, and classification. The results
 470 show that naive post-training does not yield effective representations for downstream tasks.

471 **Effect of Dual-channel Attention Layer.** In Figure 6, we investigate whether our dual-channel
 472 attention mechanism preserves pre-trained knowledge during fine-tuning. We evaluate several
 473 approaches for fine-tuning the vision encoder including full fine-tuning, LoRA (Hu et al., 2021), and
 474 dual-channel (Hong et al., 2023a) on ImageNet (Russakovsky et al., 2015) and ADE20K (Zhou
 475 et al., 2017). Dual-channel attention uniquely preserves and even enhances pre-trained knowledge,
 476 while other approaches cause degradation.

477
 478 **Dataset Scalability.** We analyze the impact of dataset sizes on depth estimation results from NYUd
 479 (Silberman et al., 2012) benchmark and semantic segmentation results from ADE20K (Zhou et al.,
 480 2017) benchmark. With matched training iterations (*i.e.*, one epoch for 300K data), larger datasets
 481 yield consistent improvements, indicating robust scalability potential.

482 5 CONCLUSION

483
 484 In this paper, we have presented SpatialBoost, a framework to enhance the vision encoders by lever-
 485 aging linguistic expressions of geometric and semantic information within images. SpatialBoost



486
 487 **Figure 6: Effect of dual-channel attention layer.** We report the linear evalua-
 488 tion performance of DINOv2-ViT-L/14 across different fine-tuning strategies.

486
 487 Table 8: **Effect of post-training.** We fine-tune vision encoders with their original pre-training objec-
 488 tives (simple FT). We report RMSE for monocular depth estimation on NYUD, mIoU for semantic
 489 segmentation on ADE20K, BLEU-1 score for vision-language reasoning on ScanQA, average score
 for robot learning on CortexBench, and Top-1 accuracy (%) for classification on ImageNet-1K.

Method	Depth Estimation ↓	Segmentation ↑	Vision-Language Reasoning ↑	Robot Learning ↑	Classification ↑
OpenCLIP	0.53	39.5	36.9	65.5	84.0
+Simple FT	0.56	39.6	37.7	63.7	84.3
+SpatialBoost (Ours)	0.40	40.5	39.2	72.9	86.1
SigLIPv2	0.51	42.8	38.1	69.7	86.3
+Simple FT	0.53	43.0	38.4	67.9	86.4
+SpatialBoost (Ours)	0.39	45.1	40.8	75.8	87.6
DINOv2	0.37	49.3	39.5	68.1	84.5
+Simple FT	0.36	49.6	39.4	69.4	84.7
+SpatialBoost (Ours)	0.30	52.0	40.3	75.8	86.4
DINOv3	0.31	55.9	40.6	72.8	85.8
+Simple FT	0.31	56.4	40.2	75.5	86.1
+SpatialBoost (Ours)	0.25	59.7	43.3	80.8	87.7

500 uses LLM and dual-channel attention layers to exploit linguistic information into image representations,
 501 generates a multi-turn visual spatial reasoning dataset, and leverages them to improve the image representations.
 502 Our experiments show that SpatialBoost consistently enhances the vision encoders on various downstream tasks that require a spatial understanding of images. We hope that our work further facilitates future research on designing and enhancing vision encoders.
 503
 504

505 REFERENCES

506 Saminda Abeyruwan, Joshua Ainslie, Jean-Baptiste Alayrac, Montserrat Gonzalez Arenas, Travis Armstrong, Ashwin Balakrishna, Robert Baruch, Maria Bauza, Michiel Blokzijl, et al. Gemini robotics: Bringing ai into the physical world. *arXiv preprint arXiv:2503.20020*, 2025.

507 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

508 Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *ICCV*, 2015.

509 Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *CVPR*, 2023.

510 Daichi Azuma, Taiki Miyanishi, Shuhei Kurita, and Motoaki Kawanabe. Scanqa: 3d question answering for spatial scene understanding. In *CVPR*, pp. 19129–19139, 2022.

511 Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021.

512 Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *CVPR*, pp. 5470–5479, 2022.

513 Aleksei Bochkovskii, Amaël Delaunoy, Hugo Germain, Marcel Santos, Yichao Zhou, Stephan R Richter, and Vladlen Koltun. Depth pro: Sharp monocular metric depth in less than a second. *arXiv preprint arXiv:2410.02073*, 2024.

514 David Charatan, Sizhe Lester Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *CVPR*, 2024.

515 Boyuan Chen, Zhuo Xu, Sean Kirmani, Brain Ichter, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. In *CVPR*, 2024a.

516 Dave Zhenyu Chen, Angel X Chang, and Matthias Nießner. Scanrefer: 3d object localization in
 517 rgb-d scans using natural language. In *ECCV*, pp. 202–221. Springer, 2020a.

540 Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever.
 541 Generative pretraining from pixels. In *ICML*, 2020b.

542

543 Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big
 544 self-supervised models are strong semi-supervised learners. In *NeurIPS*, 2020c.

545

546 Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision
 547 transformers. In *ICCV*, 2021.

548

549 Yuedong Chen, Haofei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-
 550 Jen Cham, and Jianfei Cai. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view
 551 images. In *ECCV*, 2024b.

552

553 An-Chieh Cheng, Hongxu Yin, Yang Fu, Qiushan Guo, Ruihan Yang, Jan Kautz, Xiaolong Wang,
 554 and Sifei Liu. Spatialrgpt: Grounded spatial reasoning in vision-language models. In *NeurIPS*,
 555 2024.

556

557 Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gor-
 558 don, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for
 559 contrastive language-image learning. In *CVPR*, 2023.

560

561 Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li,
 562 Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned lan-
 563 guage models. *Journal of Machine Learning Research*, 25(70):1–53, 2024.

564

565 Yin Cui, Yang Song, Chen Sun, Andrew Howard, and Serge Belongie. Large scale fine-grained
 566 categorization and domain-specific transfer learning. In *CVPR*, 2018.

567

568 Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias
 569 Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, pp. 5828–
 570 5839, 2017.

571

572 Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need
 573 registers. In *ICLR*, 2023.

574

575 Google DeepMind. Gemini 2.0 model updates: 2.0 flash, flash-lite, pro experimental. February
 576 2025.

577

578 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale
 579 hierarchical image database. In *CVPR*, 2009.

580

581 Jeff Donahue and Karen Simonyan. Large scale adversarial representation learning. In *NeurIPS*,
 582 2019.

583

584 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
 585 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An
 586 image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.

587

588 M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object
 589 classes (voc) challenge. *IJCV*, 2010.

590

591 Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu
 592 Zheng, Ke Li, Xing Sun, Yunsheng Wu, and Rongrong Ji. Mme: A comprehensive evaluation
 593 benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*, 2024a.

594

595 Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A
 596 Smith, Wei-Chiu Ma, and Ranjay Krishna. Blink: Multimodal large language models can see but
 597 not perceive. In *ECCV*, 2024b.

598

599 Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti
 600 vision benchmark suite. In *CVPR*, 2012.

601

602 Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The
 603 kitti dataset. *The international journal of robotics research*, 2013.

594 Xinyang Geng, Hao Liu, Lisa Lee, Dale Schuurmans, Sergey Levine, and Pieter Abbeel. Multimodal
 595 masked autoencoders learn transferable representations. *arXiv preprint arXiv:2205.14204*, 2022.
 596

597 Theophile Gervet, Zhou Xian, Nikolaos Gkanatsios, and Katerina Fragkiadaki. Act3d: 3d feature
 598 field transformers for multi-task robotic manipulation. *arXiv preprint arXiv:2306.17817*, 2023.
 599

600 Ankit Goyal, Jie Xu, Yijie Guo, Valts Blukis, Yu-Wei Chao, and Dieter Fox. Rvt: Robotic view
 601 transformer for 3d object manipulation. In *CoRL*, 2023.
 602

603 Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in
 604 vqa matter: Elevating the role of image understanding in visual question answering. In *CVPR*,
 605 pp. 6904–6913, 2017.
 606

607 Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Gird-
 608 har, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in
 609 3,000 hours of egocentric video. In *CVPR*, pp. 18995–19012, 2022.
 610

611 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-
 612 nition. In *CVPR*, 2016.
 613

614 Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for
 615 unsupervised visual representation learning. In *CVPR*, 2020.
 616

617 Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked
 618 autoencoders are scalable vision learners. In *CVPR*, 2022.
 619

620 Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pre-
 621 training for text-to-video generation via transformers. In *ICLR*, 2023a.
 622

623 Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang
 624 Gan. 3d-llm: Injecting the 3d world into large language models. In *NeurIPS*, 2023b.
 625

626 Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
 627 and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint
 628 arXiv:2106.09685*, 2021.
 629

630 Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning
 631 and compositional question answering. In *CVPR*, 2019.
 632

633 Stephen James and Andrew J Davison. Q-attention: Enabling efficient learning for vision-based
 634 robotic manipulation. *IEEE Robotics and Automation Letters*, 2022.
 635

636 Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-
 637 view stereopsis evaluation. In *CVPR*, pp. 406–413, 2014.
 638

639 Wolfgang Kabsch. A solution for the best rotation to relate two sets of vectors. *Foundations of
 640 Crystallography*, 32(5):922–923, 1976.
 641

642 Tsung-Wei Ke, Nikolaos Gkanatsios, and Katerina Fragkiadaki. 3d diffuser actor: Policy diffusion
 643 with 3d scene representations. *arXiv preprint arXiv:2402.10885*, 2024.
 644

645 Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete
 646 Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *ICCV*,
 647 2023.
 648

649 Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images.
 650 2009.
 651

652 Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image
 653 pre-training with frozen image encoders and large language models. In *International conference
 654 on machine learning*, pp. 19730–19742. PMLR, 2023a.
 655

656 Tianhong Li, Huiwen Chang, Shlok Mishra, Han Zhang, Dina Katabi, and Dilip Krishnan. Mage:
 657 Masked generative encoder to unify representation learning and image synthesis. In *CVPR*, 2023b.
 658

648 Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr
 649 Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.
 650

651 Fangyu Liu, Guy Emerson, and Nigel Collier. Visual spatial reasoning. *Transactions of the Association
 652 for Computational Linguistics*, 2023a.

653 Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*,
 654 2023b.

655 Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction
 656 tuning. In *CVPR*, pp. 26296–26306, 2024a.

657 Yuliang Liu, Zhang Li, Mingxin Huang, Biao Yang, Wenwen Yu, Chunyuan Li, Xu-Cheng Yin,
 658 Cheng-Lin Liu, Lianwen Jin, and Xiang Bai. Ocrbench: on the hidden mystery of ocr in large
 659 multimodal models. *Science China Information Sciences*, 67(12):220102, 2024b.

660 Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord,
 661 Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for
 662 science question answering. In *NeurIPS*, volume 35, pp. 2507–2521, 2022.

663 Xiaojian Ma, Silong Yong, Zilong Zheng, Qing Li, Yitao Liang, Song-Chun Zhu, and Siyuan Huang.
 664 Sqa3d: Situated question answering in 3d scenes. In *ICLR*, 2023.

665 Arjun Majumdar, Karmesh Yadav, Sergio Arnaud, Jason Ma, Claire Chen, Sneha Silwal, Aryan Jain,
 666 Vincent-Pierre Berges, Tingfan Wu, Jay Vakil, et al. Where are we in the search for an artificial
 667 visual cortex for embodied intelligence? In *NeurIPS*, 2023.

668 Yunze Man, Shuhong Zheng, Zhipeng Bao, Martial Hebert, Liangyan Gui, and Yu-Xiong Wang.
 669 Lexicon3d: Probing visual foundation models for complex 3d scene understanding. In *NeurIPS*,
 670 2024.

671 Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. Docvqa: A dataset for vqa on document
 672 images. In *WACV*, pp. 2200–2209, 2021.

673 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and
 674 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications
 675 of the ACM*, 65(1):99–106, 2021.

676 Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov,
 677 Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning
 678 robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.

679 Filip Radenović, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondřej Chum. Revisiting oxford
 680 and paris: Large-scale image retrieval benchmarking. In *CVPR*, pp. 5706–5715, 2018.

681 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 682 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 683 models from natural language supervision. In *ICML*, 2021.

684 Aravind Rajeswaran, Vikash Kumar, Abhishek Gupta, Giulia Vezzani, John Schulman, Emanuel
 685 Todorov, and Sergey Levine. Learning complex dexterous manipulation with deep reinforcement
 686 learning and demonstrations. *arXiv preprint arXiv:1709.10087*, 2017.

687 René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction.
 688 In *ICCV*, 2021.

689 Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham
 690 Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Va-
 691 sudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Fe-
 692 ichtenhofer. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*,
 693 2024.

694 Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and
 695 David Novotny. Common objects in 3d: Large-scale learning and evaluation of real-life 3d cate-
 696 gory reconstruction. In *ICCV*, pp. 10901–10911, 2021.

702 Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng
 703 Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual
 704 recognition challenge. *International journal of computer vision*, 2015.

705

706 Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain,
 707 Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. Habitat: A platform for embodied ai
 708 research. In *ICCV*, 2019.

709 Younggyo Seo, Junsu Kim, Stephen James, Kimin Lee, Jinwoo Shin, and Pieter Abbeel. Multi-view
 710 masked world models for visual robotic manipulation. In *ICML*, 2023.

711

712 Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, Sergey
 713 Levine, and Google Brain. Time-contrastive networks: Self-supervised learning from video. In
 714 *ICRA*, 2018.

715 Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Perceiver-actor: A multi-task transformer for
 716 robotic manipulation. In *CoRL*, 2023.

717

718 Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and sup-
 719 port inference from rgbd images. In *ECCV*, 2012.

720

721 Oriane Siméoni, Huy V Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose,
 722 Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, et al. Dinov3. *arXiv*
 723 preprint *arXiv:2508.10104*, 2025.

724

725 Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image
 726 recognition. *arXiv preprint arXiv:1409.1556*, 2014.

727

728 Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. Eva-clip: Improved training
 729 techniques for clip at scale. *arXiv preprint arXiv:2303.15389*, 2023.

730

731 Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Du-
 732 mitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions
 733 (2014). *arXiv preprint arXiv:1409.4842*, 2014.

734

735 Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Bud-
 736 den, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, et al. Deepmind control suite. *arXiv*
 737 preprint *arXiv:1801.00690*, 2018.

738

739 Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdul-
 740 mohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, et al. Siglip 2:
 741 Multilingual vision-language encoders with improved semantic understanding, localization, and
 742 dense features. *arXiv preprint arXiv:2502.14786*, 2025.

743

744 Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns.
IEEE Transactions on pattern analysis and machine intelligence, 13(4):376–380, 2002.

745

746 Fei Wang, Xingyu Fu, James Y Huang, Zekun Li, Qin Liu, Xiaogeng Liu, Mingyu Derek Ma,
 747 Nan Xu, Wenxuan Zhou, Kai Zhang, et al. Muirbench: A comprehensive benchmark for robust
 748 multi-image understanding. *arXiv preprint arXiv:2406.09411*, 2024a.

749

750 Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David
 751 Novotny. Vggt: Visual geometry grounded transformer. In *CVPR*, pp. 5294–5306, 2025a.

752

753 Jiayu Wang, Yifei Ming, Zhenmei Shi, Vibhav Vineet, Xin Wang, Sharon Li, and Neel Joshi. Is a
 754 picture worth a thousand words? delving into spatial reasoning for vision language models. In
 755 *NeurIPS*, 2025b.

756

757 Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Ge-
 758 ometric 3d vision made easy. In *CVPR*, 2024b.

759

760 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 761 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. In *NeurIPS*,
 762 2022.

756 Rundi Wu, Ben Mildenhall, Philipp Henzler, Keunhong Park, Ruiqi Gao, Daniel Watson, Pratul P
 757 Srinivasan, Dor Verbin, Jonathan T Barron, Ben Poole, et al. Reconfusion: 3d reconstruction with
 758 diffusion priors. In *CVPR*, 2024.

759

760 Manuel Wüthrich, Felix Widmaier, Felix Grimminger, Joel Akpo, Shruti Joshi, Vaibhav Agrawal,
 761 Bilal Hammoud, Majid Khadiv, Miroslav Bogdanovic, Vincent Berenz, et al. Trifinger: An open-
 762 source robot for learning dexterity. *arXiv preprint arXiv:2008.03596*, 2020.

763 xAI org. Realworldqa: a dataset of real-world questions from the xai-bench suite. 2024.

764

765 An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li,
 766 Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang,
 767 Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren
 768 Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang,
 769 Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin,
 770 Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong
 771 Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu,
 772 Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru
 773 Zhang, Zhifang Guo, and Zhihao Fan. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*,
 2024. URL <https://arxiv.org/abs/2407.10671>.

774

775 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu,
 776 Chang Gao, Chengan Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint
 777 arXiv:2505.09388*, 2025.

778 Zi Jian Yew and Gim Hee Lee. Regtr: End-to-end point cloud correspondences with transformers.
 779 In *CVPR*, pp. 6677–6686, 2022.

780 Burak Yildiz, Seyran Khademi, Ronald Maria Siebes, and Jan Van Gemert. Amstertime: A visual
 781 place recognition benchmark dataset for severe domain shift. In *ICPR*, pp. 2749–2755. IEEE,
 782 2022.

783

784 Nikolaos-Antonios Ypsilantis, Noa Garcia, Guangxing Han, Sarah Ibrahimi, Nanne Van Noord, and
 785 Giorgos Tolias. The met dataset: Instance-level recognition for artworks. In *NeurIPS*, 2021.

786

787 Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey
 788 Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning.
 789 In *CoRL*, 2020.

790 Xianggang Yu, Mutian Xu, Yidan Zhang, Haolin Liu, Chongjie Ye, Yushuang Wu, Zizheng Yan,
 791 Chenming Zhu, Zhangyang Xiong, Tianyou Liang, et al. Mvimgnet: A large-scale dataset of
 792 multi-view images. In *CVPR*, 2023.

793 Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens,
 794 Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multimodal
 795 understanding and reasoning benchmark for expert agi. In *CVPR*, pp. 9556–9567, 2024.

796

797 Yanjie Ze, Gu Zhang, Kangning Zhang, Chenyuan Hu, Muhan Wang, and Huazhe Xu. 3d diffusion
 798 policy: Generalizable visuomotor policy learning via simple 3d representations. *arXiv preprint
 799 arXiv:2403.03954*, 2024.

800 Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language
 801 image pre-training. In *ICCV*, 2023.

802

803 Junyi Zhang, Charles Herrmann, Junhwa Hur, Varun Jampani, Trevor Darrell, Forrester Cole, De-
 804 Qing Sun, and Ming-Hsuan Yang. Monst3r: A simple approach for estimating geometry in the
 805 presence of motion. *arXiv preprint arXiv:2410.03825*, 2024.

806 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable
 807 effectiveness of deep features as a perceptual metric. In *CVPR*, pp. 586–595, 2018.

808

809 Yiming Zhang, ZeMing Gong, and Angel X Chang. Multi3drefrer: Grounding text description to
 810 multiple 3d objects. In *ICCV*, pp. 15225–15236, 2023.

810 Haoyu Zhen, Xiaowen Qiu, Peihao Chen, Jincheng Yang, Xin Yan, Yilun Du, Yining Hong, and
811 Chuang Gan. 3d-vla: A 3d vision-language-action generative world model. *arXiv preprint*
812 *arXiv:2403.09631*, 2024.

813

814 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang,
815 Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and
816 chatbot arena. In *NeurIPS*, 2023.

817 Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene
818 parsing through ade20k dataset. In *CVPR*, 2017.

819

820 Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba.
821 Semantic understanding of scenes through the ade20k dataset. *IJCV*, 2019.

822 Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot:
823 Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021.

824

825 Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen
826 Duan, Weijie Su, Jie Shao, et al. Internvl3: Exploring advanced training and test-time recipes for
827 open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025.

828

829

830

831

832

833

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864 **A IMPLEMENTATION DETAILS**
865866 **A.1 TRAINING DETAILS OF STAGE 1 & 2**
867868 We train our multi-modal architecture with 4x NVIDIA Tesla A100s. In multi-modal architecture,
869 we choose Qwen-2.0-7B (Yang et al., 2024) as the LLM backbone and 2-layer MLP as the projector.
870 In feature alignment pre-training (Stage 1), we train the projector on a BLIP-558K data (Liu et al.,
871 2024a) for one epoch with a learning rate of 2e-3 and a batch size of 256. In visual instruction
872 tuning (Stage 2), we fine-tune both the projector and the LLM backbone on the LLaVA-Instruct-
873 158K dataset (Liu et al., 2024a) and our multi-view VQA dataset (described in Section 3.2) for one
874 epoch with a batch size of 128.875 **A.2 TRAINING DETAILS OF STAGE 3**
876877 In this stage, we adapt dual-channel attention layers in training vision encoders by introducing additional
878 attention channels described in Section 3.1. By applying dual-channel attention, the number
879 of model parameters increased by 30% in OpenCLIP and SigLIPv2 and by 25% in DINOv2 and DI-
880 NOv3, respectively. We freeze the LLM decoder and fine-tune the vision encoder and projector on
881 a multi-turn visual spatial reasoning dataset (described in Section 3.2) for one epoch with a learning
882 rate of 2e-5 and a batch size of 128. We conduct hyperparameter search for the learning rate from
883 1e-6 to 1e-2.884 **A.3 DENSE PREDICTION TASKS**
885886 From the vision encoder obtained through SpatialBoost, we performed depth estimation and
887 semantic segmentation. We follow the same protocol as in DINOv2 (Oquab et al., 2023), defining
888 three primary hyperparameters for our linear probing setup: the learning rate, the number of output
889 layers, and whether we concatenate the average-pooled patch token features with the class token.
890 Concretely, we perform a grid search over learning rates in 1e-4 to 1e-1, choose the output layers
891 from {1, 4}, and optionally concatenate average-pooled representations. We train each linear layer
892 with SGD for 12500 iterations using random-resized-crop data augmentation. We then select the
893 best hyperparameter combination on validation accuracy.894 **A.4 3D SCENE UNDERSTANDING**
895896 We evaluate whether SpatialBoost enables complex 3D-centric reasoning using the Lexicon3D (Man
897 et al., 2024) benchmark. Lexicon3D provides a unified probing framework that freezes visual back-
898 bones and attaches task-specific heads to evaluate vision-language reasoning, visual grounding, 3D
899 semantic segmentation, and geometric correspondence. Following the Lexicon3D protocol, we ex-
900 tract features from 2D vision encoders and evaluate them on various 3D understanding tasks.901 **Vision-Language Reasoning.** To evaluate vision-language reasoning, we target the 3D visual
902 question-answering (3D-VQA) on ScanQA (Azuma et al., 2022) and SQA3D (Ma et al., 2023)
903 datasets. We follow the 3D-LLM (Hong et al., 2023b) architecture as our task head. Specifically, we
904 use a Q-Former module (Li et al., 2023a) to project multi-view visual features into the input space
905 of the language model. These projected features are then fed to the LLM (e.g., FlanT5 (Chung et al.,
906 2024)) for generating answers. We pre-train only the Q-Former projection module for 10 epochs us-
907 ing 3D-Language dataset (Hong et al., 2023b) and fine-tune the module for 35 epochs using training
908 split of ScanQA and SQA3D. We keep both the vision encoder and LLM frozen during training.909 **Visual Grounding.** To evaluate visual grounding with vision encoder, we target the object local-
910 ization task based on text descriptions on the ScanRefer (Chen et al., 2020a) dataset. We use an
911 attention-based fusion head following Multi3DRefer (Zhang et al., 2023). The task head consists of
912 a multi-layer attention module with 4 transformer layers that fuses visual and text embeddings. After
913 projecting multi-view features to 3D space and extracting object features via average pooling within
914 ground-truth bounding boxes, we apply cross-attention between object features and CLIP-encoded
915 text descriptions. The fusion module outputs confidence scores for each object. We train the header
916 for 30 epochs with cross-entropy loss.

918 **Geometric Understanding.** To evaluate geometric understanding, we target the geometric corre-
 919 spondence task. We adopt a REGTR-style (Yew & Lee, 2022) transformer cross-encoder as the task
 920 head. The head process features from two partial point clouds to establish correspondences. After
 921 obtaining point correspondences through the transformer, we apply the Kabsch-Umeyama (Kabsch,
 922 1976; Umeyama, 2002) algorithm for closed-form estimation of rotation and translation parameters.
 923 We train the transformer head using partial scene registration benchmark (Man et al., 2024) for 30
 924 epochs using a combination of correspondence loss and transformation loss.

925 **3D Semantic Understanding.** To evaluate 3D semantic understanding, we target the point-wise
 926 semantic classification task on ScanNet (Dai et al., 2017). We employ a linear probing head consist-
 927 ing of a single fully-connected layer followed by sigmoid activation: $\mathbf{y} = \text{Sigmoid}(\text{FC}(\mathbf{x}))$, where
 928 $\mathbf{x} \in \mathbb{R}^{N \times d}$ represents projected point features from multi-view images, $\mathbf{y} \in \mathbb{R}^{N \times C}$ represents class
 929 probabilities for $C = 20$ semantic classes and N is the number of points in each point cloud. The
 930 linear layer maps from feature dimension d to the number of classes. We train the linear layer using
 931 ScanNet segmentation dataset with cross-entropy loss at learning rate 1e-4 for 20 epochs.

933 A.5 VISION-BASED ROBOT LEARNING

935 We train the robot agents using 100 demos for each task. For training, we use keypoint augmen-
 936 tation (James & Davison, 2022) for each demonstration, and use the end-effector controller with
 937 path planning as an action mode. We use the front camera of 224×224 resolution without depth
 938 measurements. We evaluate the model 5 times by training with a pre-defined interval and report the
 939 mean of the best performance.

940 A.6 IMAGE CLASSIFICATION TASK

942 We train a linear classifier on top of the [CLS] token from the last feature of the vision encoder using
 943 the training split of ImageNet-1K (Deng et al., 2009) dataset. Following the evaluation protocol
 944 of DINOv3 (Siméoni et al., 2025), we employ SGD optimizer with momentum 0.9 and random-
 945 resized-crop data augmentation. We train the linear layer for 10 epochs with a batch size of 1024.
 946 We perform a grid search for the optimal learning rate, ranging from 1e-4 to 1e-1, selecting the best
 947 performing configuration.

949 A.7 IMAGE RETRIEVAL TASK

951 We evaluate the image retrieval performance of vision encoders using a non-parametric retrieval
 952 approach. Specifically, we compute cosine similarity between the output [CLS] tokens of query
 953 and target images to establish ranking. For Oxford (Radenović et al., 2018), Paris (Radenović et al.,
 954 2018), and AmsterTime (Yıldız et al., 2022) datasets, we resize images to 224×224 resolution,
 955 while for the Met (Ypsilantis et al., 2021) dataset, we resize to the nearest multiple of the patch size.
 956 All other setups follow evaluation protocols of each benchmark.

957
 958
 959
 960
 961
 962
 963
 964
 965
 966
 967
 968
 969
 970
 971

972 **B ADDITIONAL EXPERIMENTAL RESULTS**
973974 **B.1 VISUAL QUESTION-ANSWERING (VQA) TASKS**
975976
977 Table 9: Effect of the vision encoder on spatial reasoning and general VQA benchmarks.
978

980 Model	981 Vision encoder	982 Spatial Reasoning		983 General VQA			
		984 SpatialRGPT	985 BLINK-D	986 VQAv2	987 GQA	988 SQA-I	989 MME
GPT-4o	-	39.7	72.6	-	-	-	-
Gemini-2.5-Flash	-	42.5	77.4	-	-	-	-
990 Vicuna-1.5-7B	OpenCLIP	13.3	51.6	78.5	62.0	66.8	1510.7
	+SpatialBoost	52.0	84.9	79.0	65.6	67.1	1516.3
	SigLIPv2	21.1	52.3	79.4	62.5	66.8	1519.4
	+SpatialBoost	61.3	87.5	80.0	69.1	69.5	1527.6
	DINOv2	18.8	55.2	75.2	61.5	66.0	1509.2
	+SpatialBoost	54.2	87.2	76.8	62.5	66.8	1514.2
	DINOv3	17.6	53.9	78.7	61.9	65.8	1514.7
	+SpatialBoost	58.7	87.9	80.0	65.5	67.1	1520.6

992
993 **Setup.** To investigate whether SpatialBoost can enhance visual representations by capturing geometric and semantic information within images, we evaluate our framework on VQA tasks that
994 require (1) 3D geometric spatial reasoning and (2) general knowledge.
995

996 For spatial reasoning, we consider the VQA tasks from SpatialRGPT-bench (Cheng et al., 2024)
997 and BLINK’s Relative Depth Benchmark (*i.e.*, BLINK-D) (Fu et al., 2024b), where the goal is to
998 predict the relative or absolute positional relations between objects. For general VQA, we consider
999 widely-used benchmarks such as VQAv2 (Goyal et al., 2017), GQA (Hudson & Manning, 2019),
1000 SQA-I (Lu et al., 2022), and MME (Fu et al., 2024a). Given our SpatialBoost vision encoders, we
1001 follow the setup in LLaVA-1.5 (Liu et al., 2024a) that trains the LLM backbone (Vicuna-1.5-7B
1002 (Zheng et al., 2023)) and the 2-layer MLP projector in two stages while freezing our vision encoder.
1003

1004 **Details for Spatial Reasoning.** The SpatialRGPT-Benchmark is designed to assess 3D spatial
1005 understanding across a diverse range of scenes, incorporating both quantitative and qualitative QAs.
1006 We evaluate BLINK’s Relative Depth Benchmark for depth comparison between the coordinates of
1007 two objects. Given that these benchmarks allow for multiple correct answers, leveraging an LLM-
1008 based evaluation provides a reasonable and consistent approach to judging model responses. For
1009 this, we utilize GPT-4 (Achiam et al., 2023) to determine the accuracy of the responses. For qualitative
1010 questions, responses are assessed on 0 to 1 scoring scale. For quantitative questions, the LLM
1011 extracts numerical values from answers and model responses and standardizes them to a same unit
1012 for comparison. We use judging prompts following SpatialRGPT (Cheng et al., 2024).
1013

1014 In Table 9, we use the closed-source large vision language models (LVLMs), although they are not
1015 directly compared to our approach. We provide the versions of the closed-source LVLMs as follows:
1016

- 1017 • openai/gpt-4o-2024-11-20
- 1018 • Google/gemini-2.5-flash-preview-04-17

1019 **Results.** As shown in Table 9, we observe that SpatialBoost consistently and significantly enhances
1020 both the spatial reasoning capabilities and general knowledge of existing vision-language models,
1021 even though only the frozen vision encoder was changed. For instance, Vicuna-1.5-7B with
1022 SpatialBoost DINOv3 raises the score 17.6 to 58.7 on SpatialRGPT benchmark, surpassing frontier
1023 models like GPT-4o (Achiam et al., 2023) (39.7) and Gemini-2.5-Flash (DeepMind, 2025) (42.5).
1024 This demonstrates that our framework can indeed induce representations that are useful for solving
1025 complex QA tasks that require spatial understanding while preserving or even improving its general
knowledge.
1026

1026 **C MULTI-VIEW VQA DATASET**
1027

1028 We utilize multi-view data to inject rich 3D information into vision encoders. We found that proper
1029 instruction tuning is crucial for LLMs to stably transfer the 3D information to vision encoders.
1030 However, existing datasets are limited to enhance multi-view understanding, as most VQA datasets
1031 focus exclusively on single-view scenarios. We thereby construct a multi-view VQA dataset.

1032 We consider both 3D datasets and ego-centric video data for our multi-view VQA construction.
1033 Specifically, we utilize ScanNet (Dai et al., 2017), Mip-NeRF360 (Barron et al., 2022), and
1034 MVImgNet (Yu et al., 2023) for 3D data, and Ego4D (Grauman et al., 2022) for ego-centric video
1035 data. From these datasets, we extract pairs of images that satisfy the following LPIPS (Zhang et al.,
1036 2018) constraint:

1037
$$0.35 \leq \text{LPIPS}(\mathbf{x}_i, \mathbf{x}_j) \leq 0.65, \text{ where } \mathbf{x}_i, \mathbf{x}_j \in \{\mathbf{x}_1 \dots \mathbf{x}_N\}. \quad (2)$$
1038

1039 This constraint effectively filters out outlier samples for meaningful multi-view learning. Given
1040 the selected image pairs, we utilize GPT-4o (Achiam et al., 2023) to generate three types of visual
1041 questions: (1) common VQA, (2) adversarial VQA, and (3) multi-choice VQA. These question
1042 types are designed to probe general knowledge understanding from multi-view visual inputs, thereby
1043 guiding the model to accurately process and answer multi-view visual questions. We provide specific
1044 prompts used for generating multi-view VQA data in Table 10.

1045 **Table 10: Prompt examples for generating multi-view VQA data.**

1046
1047

```
system_prompt =[
1048     "You are a helpful multimodal assistant.
1049     Generate question-answer pairs for given two images.
1050     Both images are came from same scene.
1051     When referring to the image, please call it the first image or the second image."
1052 ]
1053 general_vqa.prompt =[
1054     "Please give me an exact question and answer by referring to the images.
1055     This is a common VQA.
1056     Create relevant question about these 2 images,
1057     referencing details that may only be visible if we consider both views.
1058     Then provide a concise, correct answer.
1059     The answer should be in length between 10 and 80 words."
1060 ]
1061 multi_choice_vqa.prompt =[
1062     "Please give me an exact question and answer by referring to the images.
1063     This is a multi-choice VQA.
1064     Create relevant question about these 2 images,
1065     referencing details that may only be visible if we consider both views.
1066     Then also generate 4 answer candidates,
1067     where only one candidate is correct and the others are very wrong.
1068     List candidates A to D or 1 to 4.
1069     The answer is the index of correct question.
1070     Each candidates should be in length between 5 and 20 words.
1071 ]
1072
1073
1074
1075
1076
1077
1078
1079
```

1080 D MULTI-TURN VISUAL SPATIAL REASONING DATASET 1081

1082 We here provide a detailed implementation of the data generation pipeline and examples of multi-
1083 turn visual spatial reasoning.
1084

1085 We construct a multi-turn visual spatial reasoning dataset by associating each single-view image \mathbf{x}
1086 or multi-view images $\{\mathbf{x}_1 \dots \mathbf{x}_N\}$ with 12 sequential QA turns. The first 5 turns focus on pixel-level
1087 view, prompting questions about point-wise depth or depth comparisons. The next 4 turns shift to
1088 object-level queries, referring to approximate bounding cubes (*i.e.*, 3D bounding boxes) for each
1089 object. The next one turn addresses scene-level understanding, requiring holistic 3D interpretation.
1090 The last 2 turns are GPT-generated scene captions for given image input. For instance, the entire
1091 sequence of question-answer pairs for image \mathbf{x} is described by
1092

$$\begin{aligned} \text{Pixel-level} : & (Q_{\mathbf{x}}^{(1)}, A_{\mathbf{x}}^{(1)}) \rightarrow \dots \rightarrow (Q_{\mathbf{x}}^{(5)}, A_{\mathbf{x}}^{(5)}) \rightarrow, \\ \text{Object-level} : & (Q_{\mathbf{x}}^{(6)}, A_{\mathbf{x}}^{(6)}) \rightarrow \dots \rightarrow (Q_{\mathbf{x}}^{(9)}, A_{\mathbf{x}}^{(9)}) \rightarrow, \\ \text{Scene-level} : & (Q_{\mathbf{x}}^{(10)}, A_{\mathbf{x}}^{(10)}) \rightarrow, \\ \text{Scene Caption} : & (Q_{\mathbf{x}}^{(11)}, A_{\mathbf{x}}^{(11)}) \rightarrow (Q_{\mathbf{x}}^{(12)}, A_{\mathbf{x}}^{(12)}). \end{aligned}$$

1093 Each turn builds on the previous answers, allowing the LLM to engage in CoT reasoning. To extract
1094 3D information for each image, we use the specialized vision models (*e.g.*, depth and segmentation
1095 networks) and synthesize QA pairs that reflect the relevant 3D information, ensuring that the
1096 final scene-level query can integrate pixel-level and object-level details into a coherent spatial
1097 understanding.
1098

1099 **Filtering for Single-view Image.** Generating visual spatial reasoning data requires multiple
1100 objects in an image. Therefore, selecting the appropriate images is necessary. Following SpatialVLM
1101 ([Chen et al., 2024a](#)) and SpatialRGPT ([Cheng et al., 2024](#)), we adopt a CLIP-based open-vocabulary
1102 classification model ([Sun et al., 2023](#)) to identify appropriate images with 100K samples from 314K
1103 samples of SA1B ([Kirillov et al., 2023](#)). We provide the labels to get filtered images in Table 11.
1104

1105 **Table 11: CLIP labels for filtering images.**

1106 Label type	1107 Labels
1108 Positive labels	“an iPhone photo of an indoor scene”
	“an iphone photo of an outdoor scene”
	“a DSLR photo of an indoor scene”
	“a DSLR of an outdoor scene”
1109 Negative labels	“a close up shot of a single object”
	“a product displayed in front of a white background”
	“an artwork”
	“a painting”
	“a screenshot of a graphical user interface”
	“a piece of text”
	“a sketch”

1110 **Filtering for Multi-view Images.** We apply LPIPS ([Zhang et al., 2018](#)) metric to 3D data (*e.g.*,
1111 ScanNet ([Dai et al., 2017](#)) trainset) and ego-centric video data (*e.g.*, Ego4D ([Grauman et al., 2022](#)))
1112 to obtain pairs of images that satisfy Equation (2). This constraint prevents sampling of image pairs
1113 that are either too dissimilar or overly redundant from the datasets.
1114

1115 **Point Cloud Processing.** We process two types of input: (1) single-view and (2) multi-view. For
1116 a single-view image, we use the results of the segmentation and depth estimation to generate a
1117 3D point cloud for objects in images. In particular, we use Depth-pro ([Bochkovskii et al., 2024](#))
1118 to perform metric depth estimation. For multi-view images, we obtain a 3D point cloud through
1119 VGGT ([Wang et al., 2025a](#)), which is a state-of-the-art 3D reconstruction model. For each image
1120 input $\{\mathbf{x}_1 \dots \mathbf{x}_N\}$, we first select an image \mathbf{x}_i , where $\mathbf{x}_i \in \{\mathbf{x}_1 \dots \mathbf{x}_N\}$, among the image input
1121 and generate pixel-level data by randomly selecting the 2D coordinates of bounding boxes in \mathbf{x}_i
1122 and then extract the depth information. We also generate object and scene-level data by randomly
1123 selecting the bounding cubes obtained by using 3D point cloud. We represent the bounding cubes in
1124 the canonical space, which is proposed by SpatialVLM ([Chen et al., 2024a](#)).
1125

Table 12: Template examples for pixel-level VQA.

```
1135     single_point_questions = [
1136         "What is the depth value at pixel point [A]?"
1137         "How far away is point [A]?"
1138         "Tell me the depth of point [A]."
1139     ]
1140     single_point_answers = [
1141         "[X] away."
1142         "It is [X]."
1143         "Depth value of point [A] is [X]."
1144     ]
1145     close_predicate_questions = [
1146         "Which point is close to a viewer? Point: [A], Point: [B]."
1147         "Is point [A] closer than [B]?"
1148         "Which point has a smaller depth value? Point [A] or Point [B]?"
1149         "Compare the depth of point [A] and point [B]."
1150     ]
1151     close_true_responses = [
1152         "Yes, point [A] is closer to the viewer than point [B]."
1153         "Indeed, point [A] has a smaller depth value than point [B]."
1154         "Correct, point [A] is closer than point [B]."
1155     ]
1156     close_false_responses = [
1157         "No, point [A] is not closer than point [B]."
1158         "In fact, point [B] is closer to the viewer than point [A]."
1159         "Incorrect, point [B] has a smaller depth value than point [A]."
1160     ]
1161 
```

Table 13: Template examples for object-level VQA.

```
1165 bounding_cube_questions = [
1166     "Identify [A] and [B]"
1167     "What is the center of the 3d bounding box coordinate for [A]?"
1168 ]
1169 bounding_cube_answers = [
1170     "[X]"
1171     "Center: [X]"
1172     "[A] in [X] and [B] in [Y]"
1173 ]
1174 left_predicate_questions = [
1175     "Is the [A] to the left of the [B] from the viewer's perspective?"
1176     "Does the [A] appear on the left side of the [B]?"
1177     "Can you confirm if the [A] is positioned to the left of the [B]?"
1178 ]
1179 left_true_responses = [
1180     "Yes, the [A] is to the left of the [B]."
1181     "Indeed, the [A] is positioned on the left side of the [B]."
1182     "Correct, you'll find the [A] to the left of the [B]."
1183 ]
1184 left_false_responses = [
1185     "No, the [A] is not to the left of the [B]."
1186     "In fact, the [A] is either to the right of or directly aligned with the [B]."
1187     "Incorrect, the [A] is not on the left side of the [B]."
1188 ]
```

Object-level VQA Data. Object-level dataset has two types of QAs: (1) predicting a bounding cube of an object from the bounding box of the object, and (2) predicting the relative positional relationship between two objects. We provide examples of templates for each type of QA of this level in Table 13.

Scene-level VQA Data. Scene-level dataset has single type of QA: predicting the 3D relative distance between two objects. We provide examples of templates for each type of QA of this level in Table 14.

1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
Table 14: Template examples for scene-level VQA.
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

```

distance_questions = [
    "What is the distance between the [A] and the [B]?"
    "How far is the [A] from the [B]?"
    "How distant is the [A] from the [B]?"
    "Measure the distance from the [A] to the [B]."
]
distance_answers = [
    "[X]"
    "the [A] and the [B] are [X] apart."
    "They are [X] apart."
    "The distance of the [A] from the [B] is [X]."
]

```

Expand Viewpoints in Multi-view Data. Through the aforementioned process, we obtain multi-view reasoning data for 2-view images. We denote these obtained views as anchor views. To extend beyond 2-view configurations, we additionally sample interpolated frames between the anchor views and validate whether the VQA pairs generated for the anchor views remain valid for these new viewpoints using GPT-4o. Specifically, if the existing VQA pairs are verified as correct for more than half of the interpolated views, we incorporate these interpolated views as additional viewpoints. This approach enables us to extend the 2-view input to arbitrary multi-view configurations. Among our 200K multi-view samples, we have 160K 2-view samples, 30K 4-view samples, and 10K 8-view samples.

E DETAILS OF ABLATION STUDY AND ANALYSIS

We here provide a detailed implementation of ablation study and analysis.

E.1 COMPARISON ON DIFFERENT HEADERS

Our key hypothesis is that language supervision, particularly through LLM-based supervised fine-tuning, can effectively distill rich 3D information into vision encoders. To validate this, we investigate whether LLM provides superior supervision compared to pixel-level alternatives. We align various headers with vision encoders following the SpatialBoost framework, then fine-tune the vision encoder with dual-channel attention. We evaluate each enhanced vision encoder on ImageNet-1K (Deng et al., 2009) image classification, ADE20K (Zhou et al., 2017) semantic segmentation, and NYUd (Silberman et al., 2012) monocular depth estimation. As shown in Table 6, pixel-level supervision leads to catastrophic forgetting, while language supervision preserves pre-trained knowledge. This validates our hypothesis that language serves as an effective modality for transferring dense and hierarchical spatial information.

For all experiments, we fine-tune the vision encoder with fixed 300K samples extracted from our multi-turn visual reasoning dataset, except for the VGGT experiment. We choose DINOv2-ViT-L/14 as a vision encoder architecture, with following evaluation protocols for each downstream task detailed in Section A. The specific implementation for each header-based fine-tuning approach is provided in following paragraphs:

SAM Decoder. We adopt the SAM decoder as a header and introduce an MLP layer to match dimensions with the vision encoder. Following the SpatialBoost training strategy, we first align only the MLP layer using 300K samples from SA1B (Kirillov et al., 2023) dataset. Subsequently, we apply dual-channel attention to the vision encoder and fine-tune it using 300K segmentation samples from our multi-turn visual reasoning dataset, which is also sampled from SA1B dataset.

VGGT Decoder. VGGT (Wang et al., 2025a) is a state-of-the-art 3D reconstruction model that employs DINOv2-ViT-L/14-reg (Darcret et al., 2023) as a feature extractor. Building upon this off-the-shelf pipeline, we apply dual-channel attention to the vision encoder and perform fine-tuning. We utilize 300K 3D data samples from Co3D (Reizenstein et al., 2021) for training.

1242 **Linear Layers.** We consider two different pixel-level modalities as input for linear layers: (1)
1243 depth and (2) segmentation. As linear layers are randomly initialized, we first train the linear layer
1244 while freezing the vision encoder. **We use 300K samples from SA1B to train the linear layer, then**
1245 **apply dual-channel attention to the vision encoder and fine-tune the vision encoder with 300K sam-**
1246 **ples from our reasoning data.** For depth data, we use depth maps obtained through Depth-Pro
1247 (Bochkovskii et al., 2024) on SA1B and a subset of our reasoning dataset. For segmentation data,
1248 we follow the same data configuration in SAM decoder experiment.

1249 **LLM (Ours).** We use Qwen-2.0-7B (Yang et al., 2024) as the LLM backbone. **We train the pro-**
1250 **jector with 300K SA1B and fine-tune the vision encoder and projector with 300K samples from our**
1251 **reasoning data.** We follow all other training setup described in Section A.

1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

1296 F ADDITIONAL ANALYSIS

1298 F.1 DETAILED ANALYSIS ON REASONING HIERARCHY

1300 In this section, we investigate which components of the multi-turn visual reasoning data contribute
 1301 most significantly to the performance of SpatialBoost. We provide a detailed analysis.

1302 Table 15: Effect of reasoning hierarchy.

1304 1305 1306 1307 1308 1309 1310 1311 1312 1313 1314 1315	Method	Depth ↓				Segmentation ↑				Classification ↑	
		OpenCLIP		DINOv2		OpenCLIP		DINOv2		OpenCLIP	DINOv2
		lin.	DPT	lin.	DPT	lin.	+ms	lin.	+ms		
Pre-trained		0.56	0.41	0.38	0.29	39.1	45.7	47.7	53.1	83.9	86.3
Pix		0.52	0.40	0.34	0.29	39.6	46.3	48.2	53.4	84.0	86.6
Obj		0.53	0.41	0.37	0.30	39.4	46.3	48.0	53.3	84.7	87.2
Scene		0.53	0.41	0.38	0.32	39.2	45.9	47.7	53.3	84.5	87.1
Pix + Obj		0.44	0.39	0.35	0.28	39.8	46.6	48.8	53.5	84.7	87.3
Pix + Scene		0.46	0.40	0.36	0.28	39.5	46.5	48.5	53.4	84.4	87.2
Obj + Scene		0.51	0.42	0.39	0.31	39.5	46.5	47.6	53.3	85.0	87.4
Pix + Obj + Scene		0.42	0.39	0.32	0.27	40.0	46.9	49.2	54.2	85.1	87.6

1316
 1317 **Setup.** We explore which levels of the reasoning hierarchy have an impact on the performance of
 1318 SpatialBoost by measuring the performance of vision encoders fine-tuned with different combi-
 1319 nations of reasoning levels. For all experiments, we fix the sample size at 100K and ensure identical
 1320 ratio for each combination. We evaluate monocular depth estimation on NYUd (RMSE), semantic
 1321 segmentation on ADE20K (mIoU), and classification on ImageNet-1K (Top-1 accuracy). We
 1322 use ViT-L/14 as a vision encoder architecture in all experiments. All other setups are the same as
 1323 described in Section A.

1324
 1325 **Results.** As shown in Table 15, we observe that pixel-level QA and its combinations remark superior
 1326 performance in dense prediction tasks, indicating pixel-level QA aids in higher-level understand-
 1327 ing. We also observe that object-level QA and its combinations achieve strong improvements in classi-
 1328 fication. The results highlights that the combination with all levels achieves the best performance
 1329 across all tasks, validating the effectiveness of our hierarchical reasoning.

1330 F.2 DETAILED ANALYSIS ON SINGLE-VIEW AND MULTI-VIEW DATA

1331 In this section, we investigate the effect of single-view and multi-view data on various downstream
 1332 tasks. We provide a detailed analysis.

1333 Table 16: Effect of single-view and multi-view data across diverse tasks.

1334 1335 1336	Model	SV	MV	Cls ↑	Seg ↑	Depth ↓	VLR		VG		GU		3D SU	
							ScanQA ↑	SQA3D ↑	ScanRef ↑	RR@0.05m ↑	RTE ↓	Acc ↑	mIoU ↑	
1337 1338 1339	SigLIPv2	-	-	89.1	42.8	0.51	38.1	48.5	51.4	47.8	0.28	47.7	9.2	
		+200K	+100K	90.2	44.7	0.41	40.5	50.0	56.6	84.1	0.18	77.7	51.8	
		+150K	+150K	90.0	44.9	0.39	40.6	50.1	56.6	84.9	0.16	80.2	52.4	
		+100K	+200K	90.0	45.1	0.39	40.8	50.1	56.8	86.4	0.15	81.0	55.5	
1340 1341	DINOv3	-	-	88.4	55.9	0.31	40.6	51.4	56.2	86.9	0.10	91.1	69.1	
		+200K	+100K	90.2	59.5	0.27	43.1	54.7	61.1	96.0	0.08	91.4	69.7	
		+150K	+150K	90.3	59.6	0.26	43.1	55.0	61.1	96.9	0.07	91.6	70.2	
		+100K	+200K	90.2	59.7	0.25	43.3	54.9	61.1	97.5	0.06	91.9	70.6	

1342
 1343 **Setup.** We explore the effect of single-view and multi-view data by fine-tuning the vision encoder
 1344 with different proportions of our reasoning data. With fixed total samples, *i.e.*, 300K from multi-
 1345 turn spatial reasoning data, we train the vision encoders and evaluate them on classification (Cls)
 1346 on ImageNet-1K, segmentation on ADE20K, depth estimation on NYUd, and 3D-centric tasks on
 1347 Lexicon3D (Man et al., 2024). We use ViT-g/16 and ViT-7B/16 as the architecture of SigLIPv2 and
 1348 DINOv3, respectively. All other setups are the same as described in Section A.

1349 **Results.** As shown in Table 16, we observe that multi-view reasoning data leads to improvements
 1350 in tasks which require spatial knowledge such as depth estimation, segmentation, geometric under-

standing (GU), and 3D semantic understanding (3D SU). Following the size of multi-view data, SigLIPv2’s GU registration recall improves from 84.1% to 86.4%, and 3D SU mIoU improves from 51.8% to 55.5%. These results demonstrate that multi-view reasoning data can effectively enhance 3D understanding of the vision encoder.

F.3 DETAILED ANALYSIS ON DUAL-CHANNEL ATTENTION

We provide quantitative and qualitative results for dual-channel attention (see Table 17 and Figure 7).

Table 17: Quantitative results of dual-channel attention.

Method	Classification \uparrow	Segmentation \uparrow	Depth estimation \downarrow
DINOv2 (Pre-trained)	86.3	47.7	0.38
Full Fine-tuning	79.5	49.4	0.31
LoRA	81.7	49.0	0.32
Dual-Channel Attn.	87.6	49.2	0.32

Setup. We evaluate different fine-tuning methodologies while fixing the reasoning data sample size at 100K. Performance is measured on ImageNet-1K classification (accuracy), ADE20K segmentation (mIoU), and NYUd depth estimation (RMSE). All experiments utilize DINOv2 with ViT-L/14 architecture.

Results. As shown in Table 17, we find that full fine-tuning and LoRA similarly exhibits performance drops in classification. In contrast, dual-channel attention shows consistent performance improvements across all tasks. This indicates that dual-channel attention effectively enhances spatial capabilities while preventing overfitting to spatial-specific features, maintaining the generalization ability. Partial results of Table 17 are visualized in Figure 6.



Figure 7: **Qualitative results for dual-channel attention.** We visualize attention heatmap from (a) DINOv3 and (b) SpatialBoost DINOv3. We rollout attention layers for cosine similarity between patches. Red cross denotes a query patch. We visualize pure attention heatmap (top) and RGB overlayed version (bottom).

F.4 DETAILED RESULTS ON DATA SCALABILITY

We provide more detailed results for data scalability. In Table 18, SpatialBoost improves SigLIPv2 and DINOv3 in all tasks.

1404 Table 18: Data scalability on classification, segmentation, and depth estimation.
1405

Model	Sample size	Classification \uparrow	Segmentation \uparrow	Depth estimation \downarrow
SigLIPv2	Pre-trained	89.1	42.8	0.51
+SpatialBoost	50K	89.5	43.2	0.44
	100K	89.7	44.5	0.42
	300K	90.0	45.1	0.39
DINOv3	Pre-trained	88.4	55.9	0.31
+SpatialBoost	50K	88.6	56.8	0.29
	100K	90.0	58.3	0.28
	300K	90.2	59.7	0.25

1416
1417 F.5 ANALYSIS ON BIAS PROPAGATION IN REASONING DATA
14181419 We provide an analysis of bias in vision foundation models used to generate spatial reasoning data.
1420

1421 Table 19: Comparison between VFM-based and GT-based reasoning data.

Method	Cls \uparrow	Seg \uparrow	Depth \downarrow	VLR \uparrow
DINOv2	86.3	47.7	0.38	39.2
+VFM-based	87.5	48.7	0.34	39.6
+GT-based	87.5	48.8	0.34	36.9
Δ (VFM – GT)	0.0	-0.1	0.0	0.0

1422 **Setup.** We explore the effect of bias propagation from vision foundation models (e.g., SAM, Depth-
1423 pro) used to generate spatial reasoning data. With fixed 100K ScanNet (Dai et al., 2017) single-
1424 view samples, we generate reasoning data based on 3D metadata extracted from vision foundation
1425 models (VFM-based) and ScanNet ground-truth annotation (GT-based). We then fine-tune the vision
1426 encoder and evaluate the performance on ImageNet-1K classification (Cls), ADE20K segmentation,
1427 NYUd depth estimation, and ScanQA vision-language reasoning (VLR).
14281429 **Results.** As shown in Table 19, we observe that the performance between VFM-based and GT-based
1430 is negligible. The results demonstrate that the effect of bias propagation is marginal in our reasoning
1431 data pipeline.
14321433 F.6 ADDITIONAL RESULTS ON MULTI-MODAL LARGE LANGUAGE MODELS
14341435 We provide results of application our framework on Multi-modal Large Language Models (MLLM).
1436

1437 Table 20: Effect of SpatialBoost on MLLM visual encoders.

Method	#Params	Cls \uparrow	Seg \uparrow	Depth \downarrow
InternViT-6B-v2.5	5.5B	86.6	39.4	0.46
+SpatialBoost (Ours)	6.0B	89.1	48.5	0.35
Qwen3-VL-VE	0.6B	87.9	40.8	0.44
+SpatialBoost (Ours)	0.7B	89.3	44.3	0.36

1437 **Setup.** We apply SpatialBoost on the vision encoders of InternVL-3 (Zhu et al., 2025) and Qwen3-
1438 VL (Yang et al., 2025). With fixed 300K samples from our reasoning data, we fine-tune the vision
1439 encoder and evaluate linear probing for ImageNet-1K classification, ADE20K segmentation, and
1440 NYUd depth estimation.
14411442 Additionally, we evaluate the performance of MLLM with SpatialBoost encoder on VQA tasks
1443 targeting multi-modal reasoning (MMMU (Yue et al., 2024)), real world comprehension (Real-
1444 WorldQA (xAI org., 2024)), OCR and document understanding (OCRBench (Liu et al., 2024b),
1445 DocVQA (Mathew et al., 2021)), multi-image comprehension (BLINK (Fu et al., 2024b), MUIR-
1446 Bench (Wang et al., 2024a)), and embodied reasoning (ERQA (Abeyruwan et al., 2025)).
1447

1458
1459 Table 21: Effect of SpatialBoost on MLLM VQA performance.
1460
1461

Method	MMMU	RealWorldQA	OCR Bench	DocVQA	BLINK	MUIRBench	ERQA
InternVL 3-38B	70.1	75.6	886	95.4	64.0	63.8	42.8
+SpatialBoost (Ours)	70.8	75.9	894	95.4	69.2	70.7	49.3
Qwen3-VL-32B-Instruct	76.0	79.0	895	96.9	67.3	72.8	48.8
+SpatialBoost (Ours)	76.4	79.6	909	97.1	70.8	76.4	51.5

1462
1463
1464
1465
1466 **Results.** As shown in Table 20, we observe that SpatialBoost produces notable performance gain
1467 in the vision encoders of Qwen3-VL and InternVL3. For example, InternViT-6B-v2.5 with Spatial-
1468 Boost raises the mIoU 39.4 to 48.5 on segmentation task. In Table 21, we observe that SpatialBoost
1469 yield consistent performance improvements on diverse VQA tasks. For instance, Qwen3-VL with
1470 SpatialBoost vision encoder rises the score from 72.8 to 76.4 on MUIRBench and from 48.8 to 51.5
1471 on ERQA.

1472
1473

G USE OF AI TOOLS

1474
1475 We acknowledge that a large language model (LLM) was used to refine the phrasing and grammar
1476 of the manuscript.

1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511