

VirPro: Visual-referred Probabilistic Prompt Learning for Weakly-Supervised Monocular 3D Detection

Chupeng Liu^{1*}, Jiyong Rao^{2,*}, Shangquan Sun³, Runkai Zhao^{1†}, Weidong Cai^{1†}
¹The University of Sydney ²Tongji University ³Nanyang Technological University

runkai.zhao@sydney.edu.au, tom.cai@sydney.edu.au

Abstract

Monocular 3D object detection leverages deterministic linguistic cues as effective auxiliary weak supervision, providing complementary semantic context. However, hand-crafted textual descriptions struggle to capture the inherent visual diversity of individuals across scenes, limiting the model’s ability to learn scene-aware representations. To address this challenge, we propose **Visual-referred Probabilistic Prompt Learning (VirPro)**, an adaptive multi-modal pretraining paradigm that can be seamlessly integrated into diverse weakly supervised monocular 3D detection (WS-M3D) frameworks. Specifically, we generate a diverse set of learnable, instance-conditioned prompts across scenes and store them in an Adaptive Prompt Bank (APB). Subsequently, we introduce Multi-Gaussian Prompt Modeling (MGPM), which incorporates scene-based visual features into the corresponding textual embeddings, allowing the text prompts to express visual uncertainties. From the fused vision–language embeddings, we further decode a prompt-targeted Gaussian distribution and derive a unified object-level prompt embedding for each instance. RoI-level contrastive matching is employed to enforce modality alignment, bringing embeddings of co-occurring objects within the same scene closer in the latent space, thus enhancing semantic coherence. Extensive experiments on the KITTI benchmark demonstrate that integrating our pre-training paradigm consistently yields substantial performance gains, achieving up to a 4.8% average precision improvement than the baseline. Code is available at [VirPro](#).

1. Introduction

Accurately perceiving 3D objects by a monocular detector is challenging due to the absence of explicit depth information, making existing approaches reliant on costly and labor-intensive annotations. To address this, recent

*These authors contributed equally to this work.

†Co-corresponding author.

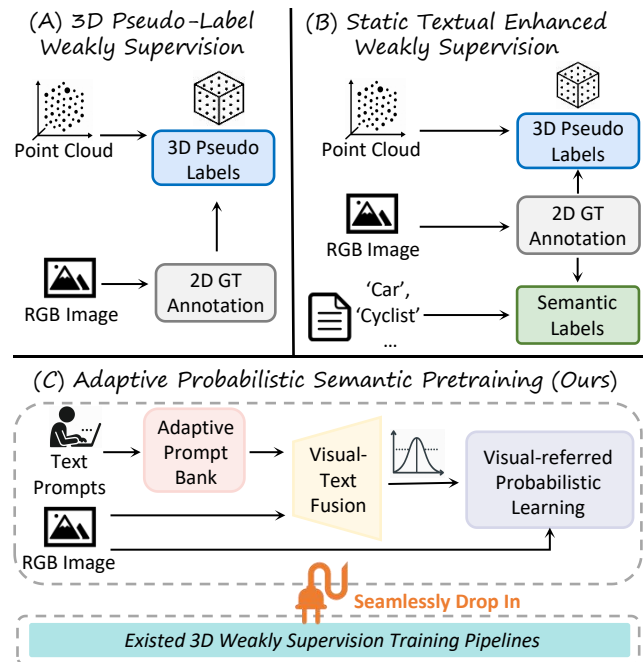


Figure 1. **Comparison of weak supervision labels for monocular 3D detection.** To mitigate label scarcity, we propose an adaptive multi-modal pretraining paradigm that leverages visually-referred probabilistic prompts as auxiliary labels and can be seamlessly integrated into existing WS-M3D pipelines.

label-efficient strategies employ pseudo-3D label generation [33, 39, 43–46, 51, 53, 59, 60], 3D knowledge distillation [16, 21, 25, 52, 61] and geometry constraint-based supervision [15, 31, 56], complementing the missing 3D depth information in 2D images. Furthermore, with the advancement of text-visual alignment [8, 11, 13, 17, 18, 23, 29, 38, 41, 49, 57], deterministic linguistic cues such as plain text have emerged as effective auxiliary weak supervision signals for context learning [1, 6, 24, 26, 47, 48, 50, 55]. Drawing inspiration from CLIP [36], aligning visual and textual embeddings by projecting semantically related concepts onto “latent” proximal regions, CAW3D [24] employs hand-crafted prompts for weak supervision to facil-

itate the detector in capturing scene-specific contextual semantics. However, relying solely on deterministic textual descriptions is insufficient to capture the intricate visual nuances, including variations in object appearance and spatial localization across different scenes, thus constraining the model’s ability to learn effective scene-aware representations. Therefore, a crucial question arises: *How can we craft prompt-based supervision that embraces cross-scene visual diversity, thereby achieving robust scene-aware representations without additional manual annotations?*

Recent advancements in Probabilistic Prompt Distribution Learning [7, 28, 37] have introduced a novel paradigm that dynamically generates diverse probabilistic prompts for multi-modal tasks [7, 22, 28], achieving superior scalability and adaptability compared to traditional static prompts [19, 62]. These probabilistically generated prompts offer varying semantic perspectives, capturing variations in object appearance and spatial context across scenes. By associating each object with such prompts, the model effectively learns high-level semantic relationships with scene-aware understanding and accurately localizes object instances without explicit human annotation. Inspired by this progress, we hypothesize that enriching textual prompts with visual cues could further encode subtle appearance variations, thus facilitating the learning of more robust and semantically meaningful scene-aware representations.

Building on this insight, we propose a novel pretraining paradigm that introduces rich semantic weak-supervision signals and can be seamlessly integrated into diverse WS-M3D frameworks. As shown in Fig. 1, prior approaches primarily rely on 3D pseudo-labels derived from 2D bounding boxes and LiDAR alignment, combined with deterministic textual descriptions for semantic guidance. Moreover, we present **Visual-referred Probabilistic Prompt Learning (VirPro)**, an adaptive text-image aligned pretraining strategy that leverages probabilistic prompts enriched with visual context to learn expressive, scene-aware representations without requiring manual annotations. Specifically, an **Adaptive Prompt Bank** assigns multiple learnable prompts to each object instance by embedding class names into natural-language templates, enabling robust contextual representation learning. To further model scene-specific variability, a **Multi-Gaussian Prompt Modeling** module injects visual cues into the prompt embeddings and parameterizes them as multivariate Gaussian distributions, where means capture canonical semantics and variances represent visual uncertainty. Randomly sampled prompts are then normalized as object-level textual embeddings for RoI contrastive matching, ensuring cross-modal semantic alignment and contextual consistency among objects within the same scene. Our contributions are summarized as follows:

- We introduce **VirPro**, an adaptive multi-modal pretraining paradigm that enriches weak supervision through vi-

ually referred probabilistic prompts and is compatible with diverse WS-M3D pipelines.

- We design an **Adaptive Prompt Bank (APB)** that generates diverse, learnable prompts for each object instance, and a **Multi-Gaussian Prompt Modeling (MGPM)** module that injects visual features and parameterizes prompts as multivariate Gaussians.
- Extensive experiments on the KITTI benchmark show up to a 4.8% AP gain over the baseline, confirming the effectiveness of visually enriched probabilistic prompts as a weak supervisory signal.

2. Related Works

2.1. Label-Efficient Monocular 3D Detection

Monocular 3D object detection typically relies on costly, labor-intensive 3D annotations, motivating recent advances in weakly supervised learning to reduce annotation dependency. One prevalent research direction leverages 2D ground-truth annotations in conjunction with LiDAR point clouds to design supervision signals with minimal spatial inconsistency [16, 40, 55], thereby aligning 3D predictions with geometric structures and providing inherent advantages in true positive localization. Additionally, to mitigate reliance on ground-truth 2D annotations, several studies adopt off-the-shelf 2D detectors to generate 2D bounding boxes to replace the 2D ground-truth labels [31, 35, 54]. Extending the pure 3D pseudo-label supervision, a novel direction exploits deep semantic cues as supervisory signals for contextual learning within text-image alignment frameworks [6, 26, 36]. For example, CAW3D [24] and GGA [55] leverage static textual prompts to facilitate semantic learning with minimal labeling effort. However, static prompts lack expressiveness and fail to capture cross-scene visual diversity, highlighting the need for adaptive, visually grounded supervision.

2.2. Probabilistic Prompt Distribution Learning

Probabilistic prompt learning [4, 7, 22, 28, 37] decouples prompt embeddings from fixed class labels, enabling more flexible adaptation to unseen categories. APP [7] introduces a Bayesian framework that models uncertainty in prompt embeddings within the input space. However, due to the sparsity and semantic variability of natural language descriptions, modeling a coherent prompt distribution in this space remains challenging. To mitigate this, several studies have shifted the focus to the output space. ProDA [28] is the first to model text prompts as multi-variable Gaussian distributions, with a regularization term that promotes diversity and improves zero-shot generalization. PPL [22] further constructs a mixture of Gaussians over attribute prompts, effectively capturing shared semantics across categories for dense prediction tasks. Departing from prior global modeling strategies, we prioritize probabilistic modeling for each

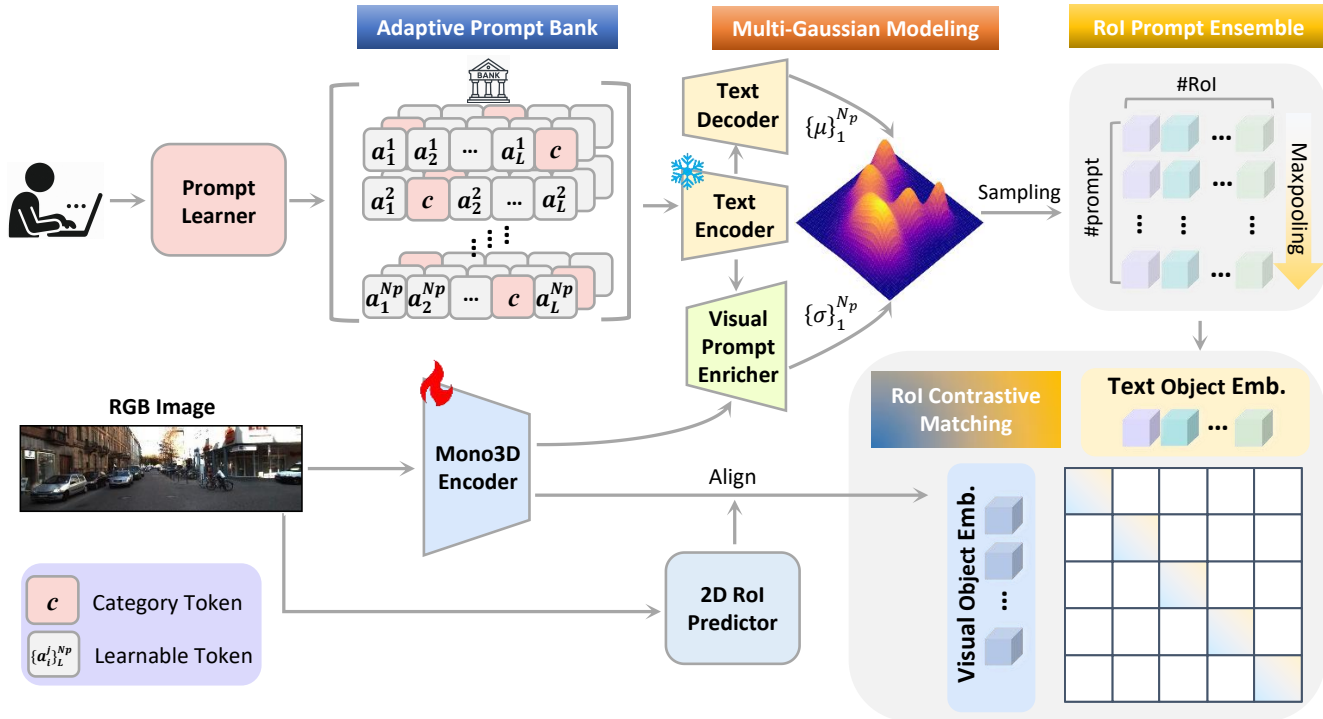


Figure 2. **Overview of the VirPro paradigm.** We propose an adaptive pretraining paradigm that generates scene-aware probabilistic prompts enriched with visual context, which can be seamlessly integrated into diverse WS-M3D frameworks. An **Adaptive Prompt Bank** includes diverse learnable prompts for each object, while **Multi-Gaussian Prompt Modeling** injects scene-specific visual features into textual embeddings and encodes prompts as a multivariate Gaussian distribution. The sampled probabilistic prompts are max-pooled for RoI-level contrastive learning to align semantics across modalities.

individual prompt as the RoI level. This design provides richer contextual cues, enabling the model to learn more robust and semantically meaningful scene-aware representations, thereby equipping it with contextual awareness.

3. Methodology

Our weakly supervised paradigm adopts a two-stage training pipeline. In the first stage, as shown in Fig. 2, we introduce an **Adaptive Prompt Bank (APB)** (Sec. 3.1) to generate diverse, instance-specific prompts. We further propose **Multi-Gaussian Prompt Modeling (MGPM)** (Sec. 3.2), which injects visual cues into textual embeddings and represents each prompt as a multivariate Gaussian distribution. A unified prompt embedding is then sampled and normalized for each instance, followed by **RoI-level Contrastive Matching** (Sec. 3.3) to align monocular 3D object embeddings with their corresponding textual prompts embeddings. In the second stage, we adopt the Dual-to-One Distillation (D2OD) strategy from CAW3D [24] to transfer the learned scene-aware priors into the monocular encoder. The overall loss formulation is provided in Sec. 3.4.

3.1. Adaptive Prompt Bank

Relying solely on visual features and a single object category prompt proves insufficient for modeling the diverse

scenario contexts in weakly-supervised monocular 3D detection. Prior studies [22, 28, 37] have shown that incorporating multiple diverse prompts offers complementary semantic cues, enhancing alignment between language and vision modalities and improving generalization on underrepresented object instances. To address the challenge of defining informative prompts for latent geometric reasoning, we propose a **prompt learner** that constructs the **Adaptive Prompt Bank**, including multiple learnable probabilistic scenario prompts for each object RoI, designed to guide the latent space structuring process. Specifically, for the i -th object query token o_i , we generate a set of N_p probabilistic prompt templates by composing learnable scenario descriptors, which serve as semantic anchors to organize the latent representation space in a geometrically aware manner:

$$p_i^t = \{a_1^t, a_2^t, \dots, a_L^t \mid o_i\}, \quad t = 1, \dots, N_p, \quad (1)$$

where $\{a_1^t, a_2^t, \dots, a_L^t\}$ comprises L learnable scenario descriptors, randomly initialized and jointly optimized during training. Furthermore, we follow the object token placement strategy [37], which enables flexible insertion of object-related tokens within prompt templates to enhances semantic grounding. Unlike prior methods such as ProDA [28], which fix the object token positions (*e.g.*, be-

ginning, middle, or end of the prompt), our approach allows randomized positioning across the template. This positional flexibility encourages the model to capture more robust contextual associations between language and visual features, which is especially critical under weak supervision. During training, both the scenario prompts and object tokens are jointly optimized with the monocular 3D detection task objective, improving both generalization and spatial reasoning in low-annotation regimes.

3.2. Multi-Gaussian Prompt Modeling

We propose a probabilistic reformulation of the prompt space to enable semantic diversity and structural disentanglement within prompt embeddings. To enable effective semantic disentanglement, the prompt loss in ‘‘Learning Objectives’’ section highlights the necessity of ensuring that scenario vectors exhibit low mutual correlation, ideally approaching orthogonality. To this end, we avoid using deterministic prompt embeddings and instead represent each scenario prompt as a distinct isotropic Gaussian distribution, parameterized by its own learnable mean and variance. Formally, for the i -th object and its N_p associated prompt scenarios, we define the distribution as:

$$\mathcal{P}(z_i^{(1:N_p)} | p_i) \sim \left\{ \mathcal{N} \left(\boldsymbol{\mu}_i^{(t)}, (\boldsymbol{\sigma}_i^{(t)})^2 \mathbf{I} \right) \right\}_{t=1}^{N_p}, \quad (2)$$

where each Gaussian corresponds to the t -th scenario-conditioned prompt. To estimate the parameters $\boldsymbol{\mu}_i^{(t)}$ and $\boldsymbol{\sigma}_i^{(t)}$, we utilize two decoders: A textual prompt decoder to produce the Gaussian mean $\boldsymbol{\mu}$. It employs a residual formulation consisting of an MLP projection and a self-attention module over the prompt set:

$$\boldsymbol{\mu}_i^t = \Phi_\mu(q_i^t) = \phi_\mu(q_i^t) + \text{SelfAttn}_\mu(q_i^t; P_i). \quad (3)$$

A cross-modal visual-text decoder to estimate the variance $\boldsymbol{\sigma}$ by attending to visual-language features F :

$$\boldsymbol{\sigma}_i^t = \Phi_\sigma(q_i^t) = \phi_\sigma(q_i^t) + \text{CrossAttn}_\sigma(q_i^t; F). \quad (4)$$

After estimating the statistical parameters for each scenario-conditioned prompt, we construct a Gaussian distribution to model the probabilistic representation space. Leveraging this distributional form, we perform stochastic sampling to produce multiple semantic variants of the original prompt, thereby capturing contextual diversity through different mean–variance combinations. Once the distribution parameters for each scenario are obtained, we instantiate a probabilistic embedding space by sampling from the corresponding Gaussians. Specifically, for each scenario t , we generate N_s stochastic samples from the learned distribution:

$$z_{i,j}^{(t)} \sim \mathcal{N} \left(\boldsymbol{\mu}_i^{(t)}, (\boldsymbol{\sigma}_i^{(t)})^2 \mathbf{I} \right), \quad j = 1, \dots, N_s, \quad (5)$$

where each sample $z_{i,j}^{(t)}$ captures a distinct contextual instantiation of the original scenario prompt. To enable end-to-end optimization, we apply the reparameterization trick:

$$\hat{z}_{i,j}^{(t)} = \boldsymbol{\mu}_i^{(t)} + \boldsymbol{\sigma}_i^{(t)} \odot \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (6)$$

where \odot denotes element-wise multiplication. This formulation facilitates efficient learning of prompt distributions while ensuring semantic diversity across samples.

3.3. RoI Contrastive Matching

We adopt an object-level matching paradigm based on image-text contrastive learning to ensure that all objects within the same scene share a consistent global context while being distinguishable from objects in different scenes. Let $\mathbf{e}_i^{\text{txt}}$ denote the text embedding of the i -th object, obtained by max-pooling the prompt distributions $\hat{z}_{i,j}^{(t)}$, and let $\mathbf{e}_i^{\text{img}}$ denote the image embedding of the same object, extracted from the monocular 3D encoder and spatially aligned with a 2D detector. The pair $(\mathbf{e}_i^{\text{txt}}, \mathbf{e}_i^{\text{img}})$ forms a positive sample. The contrastive loss is defined as:

$$\mathcal{L}_{\text{contrast}} = \frac{1}{N} \cdot (\ell_1 + \ell_2 + \dots + \ell_N), \quad (7)$$

where ℓ_i denotes the cross-entropy loss between $\mathbf{e}_i^{\text{txt}}$ and $\mathbf{e}_i^{\text{img}}$, and N is the number of objects in the batch. This objective strengthens semantic coherence among co-occurring objects and yields scene-aware priors, thereby enforcing intra-scene consistency and inter-scene separation. Consequently, the monocular encoder learns richer contextual dependencies. Additional implementation details are provided in the supplementary material.

3.4. Learning Objectives

3.4.1. Probabilistic Prompt Learning Loss

To preserve the expressivity and semantic disentanglement of probabilistic prompt embeddings, we introduce a composite loss comprising two components: a diversity loss and a KL divergence regularizer. We introduce an orthogonality-based diversity loss to explicitly encourage semantic differentiation among scenario prompts. Concretely, $\tilde{P}_i \in \mathbb{R}^{N_p \times D}$ denotes the normalized scenario embeddings for the i -th object. The diversity loss is formulated as:

$$\mathcal{L}_{\text{div}} = \frac{1}{K} \sum_{i=1}^K \left\| \tilde{P}_i \tilde{P}_i^\top - \mathbf{I} \right\|_2^2, \quad (8)$$

where $\mathbf{I} \in \mathbb{R}^{N_p \times N_p}$ is the identity matrix and K is the number of object RoI. This loss encourages the scenario embeddings to be as decorrelated as possible, promoting diverse semantics across sampled prompt variants. To further stabilize learning and prevent variance collapse, we impose a prior-matching constraint via KL divergence. Specifically,

| Method | Source | Supervision | AP _{BEV} /AP _{3D} @ IoU=0.5 _{R₄₀} | | |
|-----------------------|-----------|-----------------------------|---|-----------------------------|-----------------------------|
| | | | Easy | Moderate | Hard |
| CenterNet [63] | CVPR 2021 | Full | 34.36 / 20.00 | 27.91 / 17.50 | 24.65 / 15.57 |
| MonoGRNet [34] | AAAI 2019 | | 52.13 / 47.59 | 35.99 / 32.28 | 28.72 / 25.50 |
| M3D-RPN [2] | ICCV 2019 | | 53.35 / 48.53 | 39.60 / 35.94 | 31.76 / 28.59 |
| MonoPair [5] | CVPR 2020 | | 61.06 / 55.38 | 47.63 / 42.39 | 41.92 / 37.99 |
| MonoDLE [30] | CVPR 2021 | | 60.73 / 55.41 | 46.87 / 43.42 | 41.89 / 37.81 |
| GUPNet [27] | ICCV 2021 | | 61.78 / 57.62 | 47.06 / 42.33 | 40.88 / 37.59 |
| Kinematic [3] | ECCV 2020 | | 61.79 / 55.44 | 44.68 / 39.47 | 34.56 / 31.26 |
| MonoDistill [9] | ICLR 2022 | | 71.45 / 65.69 | 53.11 / 49.35 | 46.94 / 43.49 |
| MonoDETR [58] | ICCV 2023 | | 72.34 / 68.05 | 51.97 / 48.42 | 46.94 / 43.48 |
| VS3D [35] | ACM 2020 | | Weak (w/o 2D GT) | 31.59 / 22.62 | 20.59 / 14.43 |
| Autolabels [54] | CVPR 2020 | 50.51 / 38.31 | | 30.97 / 19.90 | 23.72 / 14.83 |
| WeakM3D [31] | ICLR 2022 | 58.20 / 50.16 | | 38.02 / 29.94 | 30.17 / 23.11 |
| CAW3D [24] | IROS 2025 | 52.99 / 46.30 | | 38.54 / 30.69 | 30.29 / 23.28 |
| VirPro+WeakM3D | - | 55.09 / 50.97 | | 38.76 / 31.95 | 31.12 / 24.27 |
| WeakMono3d [40] | CVPR 2023 | Weak (w/ 2D GT) | 54.32 / 49.37 | 42.83 / 39.01 | 40.07 / 36.34 |
| GGA+PGD [42, 55] | ECCV 2024 | | 57.20 / 51.48 | 40.11 / 35.73 | 34.96 / 30.49 |
| VirPro+GGA+PGD | - | | 60.11 / 54.72 | 42.95 / 39.49 | 37.50 / 33.32 |

Table 1. Performance comparison conducted on the KITTI *val* set for the "Car" category. All results are evaluated using the AP_{R40} metric with an IoU threshold of 0.5. For the result of "Pedestrian" and "Cyclist" category, please refer to the supplementary. The best and second-best results are highlighted in red and blue, respectively.

the learned prompt distributions are regularized toward a standard Gaussian prior as follows:

$$\mathcal{L}_{\text{prompt}} = \mathcal{L}_{\text{div}} + \frac{1}{N_p} \sum_{t=1}^{N_p} \text{KL} \left(\mathcal{P}(\hat{z}_i^{(t)} | p_i^{(t)}) \parallel \mathcal{N}(\mathbf{0}, \mathbf{I}) \right), \quad (9)$$

where $\hat{z}_i^{(t)}$ is the reparameterized embedding sampled from the Gaussian prompt distribution for scenario t . Together, these terms guide the prompt space to be both diverse and distributionally regularized, improving semantic grounding and enhancing downstream spatial reasoning.

3.4.2. Total Loss

We adopt a two-stage optimization strategy to effectively learn scene-aware visual–language knowledge and distill it into the monocular 3D detector.

Stage 1 focuses on probabilistic prompt learning and object-level alignment. We jointly optimize the RoI contrastive loss and prompt regularization loss to encourage discriminative correspondence between image embeddings and probabilistic text prompts:

$$\mathcal{L}_{\text{stage1}} = \mathcal{L}_{\text{contrast}} + \alpha \mathcal{L}_{\text{prompt}}, \quad (10)$$

where α is a weighting coefficient that balances the contributions of each loss component.

Stage 2 employs the Dual-to-One Distillation (D2OD) scheme from CAW3D [24], which preserves the baseline inference cost without add-on module designs. An MSE loss

(\mathcal{L}_{mse}) transfers the contextual semantic knowledge learned in Stage 1 to the trainable monocular 3D encoder. We retain the high-confidence 3D pseudo-labels used in existing weakly supervised monocular 3D detectors to strengthen spatial awareness and enforce geometric consistency:

$$\mathcal{L}_{\text{stage2}} = \mathcal{L}_{\text{mse}} + \lambda \mathcal{L}_{3D}, \quad (11)$$

where \mathcal{L}_{3D} denotes the pseudo-label based 3D supervision and λ balances the two terms. Additional details of \mathcal{L}_{3D} within the two WS-M3D frameworks, WeakM3D [31] and GGA [55], are provided in the supplementary material.

4. Experiments and Results


4.1. Experimental Setup

We conduct comprehensive evaluations on the KITTI [12] benchmark. We follow the CAW3D [24] protocol in stage 1 and adopt the KITTI RAW split containing 33,530 unlabeled images captured from diverse real-world driving scenarios. An off-the-shelf F-PointNet 2D detector [32] is employed to produce 2D RoI proposals for all scenes. In stage 2, we use the official KITTI 3D dataset, consisting of 3,711 training, 3,769 validation, and 7,518 test images, and report validation results using AP₄₀ with an IoU = 0.5.

4.2. Implementation Details

All experiments are conducted on a single NVIDIA RTX 4090 GPU. In the first stage, the model is trained for 25

WeakM3D



VirPro+WeakM3D





Figure 3. **Qualitative results on the KITTI validation set comparing ours to the WeakM3D baseline.** WeakM3D [31] is a WS-M3D work with pure 3D pseudo-labels. Predicted boxes are rendered in **green**, and ground-truth boxes are shown in **red**.

GGA+PGD



VirPro+GGA+PGD




Figure 4. **Qualitative results on the KITTI validation set comparing ours to the GGA+PGD baseline.** GGA [55] + PGD [42] is a WS-M3D baseline employing both 3D pseudo-labels and static textual prompts. Predicted boxes are rendered in **green**, and ground-truth boxes are shown in **red**.

| Test Result | AP _{BEV} | | | AP _{3D} | | |
|-----------------------|-------------------|-------------|-------------|------------------|-------------|-------------|
| | Easy | Mod | Hard | Easy | Mod | Hard |
| WeakM3D [31] | 11.82 | 5.66 | 4.08 | 5.03 | 2.26 | 1.63 |
| VirPro+WeakM3D | 12.23 | 5.92 | 4.33 | 5.41 | 2.52 | 1.81 |
| WeakMono3D [40] | 12.31 | 8.80 | 7.81 | 6.98 | 4.85 | 4.45 |
| GGA+PGD [42, 55] | 14.87 | 9.26 | 7.09 | 7.09 | 4.27 | 3.26 |
| VirPro+GGA+PGD | 15.59 | 9.58 | 7.29 | 7.95 | 4.96 | 3.64 |

Table 2. **Comparison on the KITTI test set (Car category).** GGA+PGD is the baseline method using weak 2D-3D alignment and textual prompts generated from LLM for weak supervision. The best and second-best results are highlighted in red and blue.

epochs with a batch size of 16 using the AdamW optimizer and a fixed learning rate of 1×10^{-4} . The text encoder is a frozen CLIP ViT-B/32. For each RoI, 32 learnable prompts are initialized, and 8 of them are randomly sampled and normalized to form the RoI-specific text embedding for contrastive matching. In addition, 4 RoIs are randomly selected from each scene to construct contrastive pairs. We initialize the temperature of the contrastive loss to $\tau = 0.07$ and optimize its logarithmic form $\log(1/\tau)$. In stage 2, we follow the default settings of the baseline WS-M3D frameworks. Additional implementation details are provided in the supplementary material.

4.3. Quantitative Analysis

Tab. 1 and Tab. 2 report results on the KITTI validation split and test split, respectively. WeakM3D [31] introduces weak supervision by aligning detected 2D bounding boxes with LiDAR point clouds. GGA [55] leverages ground-truth 2D boxes, LiDAR, and LLM generated prompts to train a point encoder, while GGA+PGD uses its predicted 3D boxes as pseudo labels within the fully-supervised PGD framework [42]. The results of GGA+PGD reported in both tables are reproduced by us. VirPro+GGA+PGD is slightly lower than WeakMono3D [40] on the hard split, as WeakMono3D benefits from 2D direction labels that explicitly guide rotation estimation. Overall, VirPro integrates seamlessly into existing weak-supervision pipelines. Compared with pure 3D pseudo labels or static textual prompts, VirPro provides richer, scene-aware semantic cues, yielding consistent performance. Results for the *Cyclist* and *Pedestrian* categories are provided in the supplementary material.

4.4. Qualitative Analysis

As shown in Fig. 3 and 4, we visualize the 3D bounding box predictions by projecting them onto 2D images and LiDAR point clouds in the BEV view. The GGA+PGD pipeline employs GGA [55] to predict 3D bounding boxes as pseudo-labels to train PGD [42], a fully supervised monocular 3D detector. GGA is a point encoder supervised by 3D pseudo-labels, obtained from the alignment of 2D ground-truth boxes with LiDAR points and static LLM-generated textual prompts that provide object size priors. Integrated with this pipeline, our VirPro introduces visual-

referred probabilistic prompts as auxiliary. As shown in Fig. 3, our approach yields more accurate 3D boxes in size, position, and orientation, demonstrating the effectiveness of the VirPro pretraining paradigm in modeling diverse visual semantics across scenes. Additional qualitative results are provided in the supplementary.

| Prompt Design | AP _{BEV} | | | AP _{3D} | | |
|---------------|-------------------|--------------|--------------|------------------|--------------|--------------|
| | Easy | Mod | Hard | Easy | Mod | Hard |
| H.C.P | 52.13 | 35.12 | 27.97 | 45.77 | 29.67 | 22.89 |
| S.P.P | 52.99 | 36.98 | 29.58 | 46.30 | 29.92 | 23.28 |
| M.P.P | 53.46 | 37.03 | 28.55 | 46.33 | 30.13 | 23.50 |

Table 3. **Ablations on Prompts Design.** H.C.P denotes a hand-crafted prompt, S.P.P represents a single probabilistic prompt per RoI, and M.P.P refers to multiple probabilistic prompts per RoI. The best and second-best results are highlighted in red and blue.

| Prompt Fusion | AP _{BEV} | | | AP _{3D} | | |
|-------------------|-------------------|--------------|--------------|------------------|--------------|--------------|
| | Easy | Mod | Hard | Easy | Mod | Hard |
| MLP | 51.47 | 35.97 | 28.56 | 45.91 | 30.20 | 23.42 |
| Concat+MLP | 53.76 | 36.17 | 28.83 | 47.71 | 30.66 | 23.90 |
| Add | 53.80 | 36.38 | 28.93 | 47.56 | 30.82 | 23.89 |
| Maxpooling | 54.06 | 37.17 | 29.09 | 47.87 | 32.11 | 25.05 |

Table 4. **Ablations on Prompts Fusion Strategies.** Comparison of four prompt fusion strategies for integrating multiple probabilistic prompts sampled from the multi-Gaussian distributions prior to RoI contrastive matching. The best and second-best results are highlighted in red and blue, respectively.

4.5. Ablation Experiments

Effect of Prompt Design. We compare three prompt configurations: hand-crafted category-level prompts (e.g., “Car”) shared across all RoIs, a single probabilistic prompt per RoI, and multiple probabilistic prompts per RoI. As reported in Tab. 3, replacing hand-crafted prompts with a single probabilistic prompt yields noticeable AP gains, confirming that instance-conditioned probabilistic prompts encode finer semantic cues than static category descriptions. Introducing multiple probabilistic prompts per RoI further improves performance, highlighting the benefit of modeling diverse semantic modes within each object. This richer prompt set enables the model to better approximate complex vision-language distributions, ultimately providing more robust supervision and stronger generalization for comprehensive scene understanding.

Effect of Prompts Fusion Strategy. We evaluate four strategies for fusing multiple prompts into a unified text object embedding: a Conv1D-based MLP with ReLU, concatenation followed by an MLP, element-wise addition, and max pooling. As shown in Tab. 4, max pooling consistently yields the best performance. This indicates that a simple, parameter-free aggregation is better suited for probabilistic prompts, as it preserves the most salient activation along each feature dimension without introducing additional projection layers or optimization noise. In

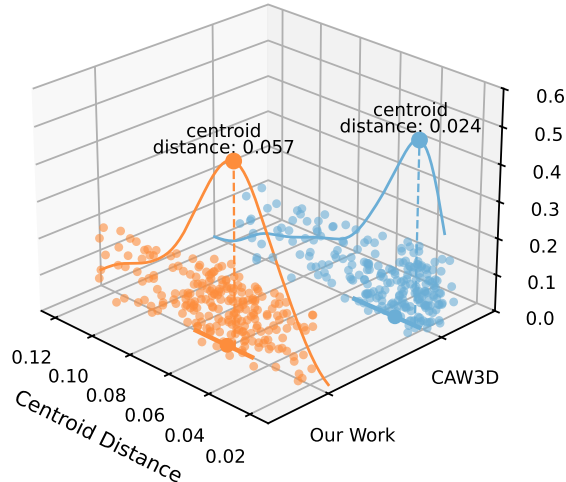


Figure 5. **Comparison of Inter-Scene Centroid Distances in Latent Space Between CAW3D and Our Proposed VirPro.** We extract RoI visual embeddings for the “Car” category from 15 scenes randomly chosen from KITTI *val* set. Then we compute the centroid of the embedding distribution and calculate pairwise distances between scene centroids.

| Image-Text Fusion | AP _{BEV} | | | AP _{3D} | | |
|-------------------|-------------------|--------------|--------------|------------------|--------------|--------------|
| | Easy | Mod | Hard | Easy | Mod | Hard |
| Add | 50.28 | 35.14 | 27.83 | 44.68 | 26.29 | 22.37 |
| Concat | 51.76 | 36.23 | 28.39 | 44.77 | 28.04 | 21.88 |
| I.P.A. | 52.15 | 36.62 | 28.64 | 45.74 | 29.41 | 22.92 |
| C.A. | 54.06 | 37.17 | 29.09 | 47.87 | 32.11 | 25.05 |

Table 5. **Ablations on Image-Text Fusion Strategies.** C.A. denotes cross-attention as adopted in MGPM. I.P.A. denotes the Image-Pooling Attention module introduced in YOLO-World [6]. The best and second-best results are highlighted in red and blue.

contrast, MLP-based fusion and concatenation may over-smooth or duplicate information on top of already aligned vision–language embeddings, while element-wise addition can suppress sparse but informative signals.

Effect of Image-Text Fusion Strategy. Our Visual Prompt Enricher (VPE) in the Multi-Gaussian Prompt Modeling (MGPM) module employs a cross attention mechanism to directly inject scene-aware visual cues into the prompt space. We compare this strategy with element-wise addition, concatenation, and I.P.A., an image-pooling attention module from YOLO-World [6] that integrates multi-scale image features into textual embeddings. As shown in Tab. 5, the proposed cross-attention consistently outperforms all alternatives. This confirms that conventional fusion mechanisms are limited for deterministic and scene-invariant descriptions, whereas our method effectively captures instance-level visual uncertainties across diverse scenes for adaptive prompts.

Latent Embeddings Space Structuring. We extract RoI visual embeddings for the Car category from the monocular 3D detector pretrained in Stage 1. For each scene in

| Latent Space | Calinski–Harabasz | Silhouette Score |
|---------------|-------------------|------------------|
| CAW3D [24] | 2.54 | -0.630 |
| VirPro | 2.62 | -0.498 |

Table 6. **Latent space structure comparison.** We evaluate the compactness and separability of RoI embeddings clusters from CAW3D and VirPro in latent space. Higher values (in red) indicate better-structured representations.

the KITTI validation set, we compute the centroid of its embedding distribution and measure all pairwise centroid distances. We compare CAW3D [24], which uses hand-crafted textual prompts, with our VirPro that leverages visually referred probabilistic prompts. As shown in Tab. 6, VirPro achieves consistently higher Calinski–Harabasz (CH) and Silhouette scores, indicating that its RoI embeddings are more compact within scenes and more separable across scenes. This validates that probabilistic prompts enriched with visual context yield cleaner and better-structured latent spaces. Moreover, Fig. 5 reveals that VirPro’s inter-scene distance distribution exhibits a heavier tail toward larger distances, suggesting stronger scene-level discrimination and improved contextual modeling. Overall, both quantitative and qualitative results demonstrate that our probabilistic prompt design induces a clearer and more compact latent structure across scenes, which enhances robust and generalizable WS-M3D. Metric definitions and additional visualizations are provided in the supplementary material.

5. Discussion and Conclusion

We introduce **Visual-referred Probabilistic Prompt Learning (VirPro)**, an adaptive multimodal pretraining paradigm compatible with diverse WS-M3D frameworks to alleviate label scarcity. Despite delivering consistent gains on KITTI, VirPro still faces several limitations. The quality of probabilistic prompts is fundamentally constrained by the reliability of region-level visual features extracted from the monocular detector. Current RoI embeddings entirely depend on 2D detector outputs, implicitly assuming accurate bounding box alignment. When 2D detections are inaccurate, the resulting visual cues become biased, yielding noisy Gaussian prompt distributions. Moreover, most objects in real scenes are not perfectly rectangular. Using 2D bounding boxes to crop RoI features inevitably introduces background noisy. In addition, RoI feature extraction is restricted by fixed image resolution and predefined cropping strategies, limiting robustness across diverse input domains. Future work may introduce more flexible RoI modeling. For example, leveraging attention maps or dense feature maps to aggregate visual evidence beyond rigid bounding boxes and dynamically emphasize object-relevant regions.

6. Acknowledgements

We thank Ziyuan Tao (ziyuan.tao@students.mq.edu.au) for his coding contributions to the experiments on the nuScenes dataset.

References

- [1] Luca Barsellotti, Lorenzo Bianchi, Nicola Messina, Fabio Carrara, Marcella Cornia, Lorenzo Baraldi, Fabrizio Falchi, and Rita Cucchiara. Talking to dino: Bridging self-supervised vision backbones with language for open-vocabulary segmentation. In *Proceedings of the CVPR*, pages 22025–22035, 2025. 1
- [2] Garrick Brazil and Xiaoming Liu. M3d-rpn: Monocular 3d region proposal network for object detection. In *Proceedings of the ICCV*, pages 9287–9296, 2019. 5, 3
- [3] Garrick Brazil, Gerard Pons-Moll, Xiaoming Liu, and Bernt Schiele. Kinematic 3d object detection in monocular video. In *Proceedings of the ECCV*, pages 135–152. Springer, 2020. 5
- [4] Guangyi Chen, Weiran Yao, Xiangchen Song, Xinyue Li, Yongming Rao, and Kun Zhang. Plot: Prompt learning with optimal transport for vision-language models. In *Proceedings of the ICLR*, 2023. 2
- [5] Yongjian Chen, Lei Tai, Kai Sun, and Mingyang Li. Monopair: Monocular 3d object detection using pairwise spatial relationships. In *Proceedings of the CVPR*, pages 12093–12102, 2020. 5
- [6] Tianheng Cheng, Lin Song, Yixiao Ge, Wenyu Liu, Xinggang Wang, and Ying Shan. Yolo-world: Real-time open-vocabulary object detection. In *Proceedings of the CVPR*, pages 16901–16911, 2024. 1, 2, 8
- [7] Youngjae Cho, HeeSun Bae, Seungjae Shin, Yeo Dong Youn, Weonyoung Joo, and Il-Chul Moon. Make prompts adaptable: Bayesian modeling for vision-language prompt learning with data-dependent prior. In *Proceedings of the AAAI 2024*, 2024. 2
- [8] Hyungyu Choi, Young Kyun Jang, and Chanho Eom. Goal: Global-local object alignment learning. In *Proceedings of the CVPR*, pages 4070–4079, 2025. 1
- [9] Zhiyu Chong, Xinzhu Ma, Hong Zhang, Yuxin Yue, Haojie Li, Zhihui Wang, and Wanli Ouyang. Monodistill: Learning spatial features for monocular 3d object detection. *Proceedings of the ICLR*, 2022. 5
- [10] MMDetection3D Contributors. Mmdetection3d: Openmmlab next-generation platform for general 3d object detection. <https://github.com/open-mmlab/mmdetection3d>, 2020. 3
- [11] Yuxin Fan, Junbiao Cui, and Jiye Liang. Learning textual prompts for open-world semi-supervised learning. In *Proceedings of the CVPR*, pages 14756–14765, 2025. 1
- [12] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Proceedings of the CVPR*, 2012. 5
- [13] Fusheng Hao, Fengxiang He, Fuxiang Wu, Tichao Wang, Chengqun Song, and Jun Cheng. Task-aware clustering for prompting vision-language models. In *Proceedings of the CVPR*, pages 14745–14755, 2025. 1
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the CVPR*, pages 770–778, 2016. 3
- [15] Kuan-Chih Huang, Yi-Hsuan Tsai, and Ming-Hsuan Yang. Weakly supervised 3d object detection via multi-level visual guidance. In *Proceedings of the ECCV*, pages 175–191, 2024. 1
- [16] Yujun Jiang, Yecheng Deng, Shuo Shi, Xiaoyang Wang, and Yizhou Shen. Weakly supervised monocular 3d detection with a single-view image. In *Proceedings of the CVPR*, pages 3323–3332, 2024. 1, 2, 5
- [17] Pengkun Jiao, Na Zhao, Jingjing Chen, and Yu-Gang Jiang. Unlocking textual and visual wisdom: Open-vocabulary 3d object detection enhanced by comprehensive guidance from text and image. In *Proceedings of the ECCV*, pages 376–392, 2024. 1
- [18] Cijo Jose, Théo Moutakanni, Dahyun Kang, Federico Baldassarre, Timothée Darcet, Hu Xu, Daniel Li, Marc Szafraniec, Michaël Ramamonjisoa, Maxime Oquab, et al. Dinov2 meets text: A unified framework for image-and pixel-level vision-language alignment. In *Proceedings of the CVPR*, pages 24905–24916, 2025. 1
- [19] Muhammad Uzair Khattak, Syed Talal Wasim, Muzammal Naseer, Salman Khan, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Self-regulating prompts: Foundational model adaptation without forgetting. In *Proceedings of the ICCV*, pages 15190–15200, 2023. 2
- [20] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 2
- [21] Donghyeon Kwon, Youngseok Yoon, Hyeongseok Son, and Suha Kwak. Memdistill: Distilling lidar knowledge into memory for camera-only 3d object detection. In *Proceedings of the ICCV*, pages 6828–6838, 2025. 1
- [22] Hyeongjun Kwon, Taeyong Song, Somi Jeong, Jin Kim, Jinhyun Jang, and Kwanghoon Sohn. Probabilistic prompt learning for dense prediction. In *Proceedings of the CVPR*, pages 6768–6777, 2023. 2, 3
- [23] Yunheng Li, Yuxuan Li, Quan-Sheng Zeng, Wenhao Wang, Qibin Hou, and Ming-Ming Cheng. Unbiased region-language alignment for open-vocabulary dense prediction. In *Proceedings of the ICCV*, pages 23795–23805, 2025. 1
- [24] Chupeng Liu, Runkai Zhao, and Weidong Cai. Ca-w3d: Leveraging context-aware knowledge for weakly supervised monocular 3d detection. In *Proceedings of the IROS*, 2025. 1, 2, 3, 5, 8
- [25] Hou-I Liu, Christine Wu, Jen-Hao Cheng, Wenhao Chai, Shian-Yun Wang, Gaowen Liu, Hugo Latapie, Jih-Ciang Wu, Jenq-Neng Hwang, Hong-Han Shuai, and Wen-Huang Cheng. Monotakd: Teaching assistant knowledge distillation for monocular 3d object detection. In *Proceedings of the CVPR*, pages 22266–22275, 2025. 1
- [26] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with

- grounded pre-training for open-set object detection. In *Proceedings of the ECCV*, pages 38–55, 2024. 1, 2
- [27] Yan Lu, Xinzhu Ma, Lei Yang, Tianzhu Zhang, Yating Liu, Qi Chu, Junjie Yan, and Wanli Ouyang. Geometry uncertainty projection network for monocular 3d object detection. In *Proceedings of the ICCV*, pages 3111–3121, 2021. 5
- [28] Yuning Lu, Jianzhuang Liu, Yonggang Zhang, Yajing Liu, and Xinmei Tian. Prompt distribution learning. In *Proceedings of the CVPR*, pages 5206–5215, 2022. 2, 3
- [29] Wufei Ma, Luoxin Ye, Celso M de Melo, Alan Yuille, and Jieneng Chen. SpatialLM: A compound 3d-informed design towards spatially-intelligent large multimodal models. In *Proceedings of the CVPR*, pages 17249–17260, 2025. 1
- [30] Xinzhu Ma, Yinmin Zhang, Dan Xu, Dongzhan Zhou, Shuai Yi, Haojie Li, and Wanli Ouyang. Delving into localization errors for monocular 3d object detection. In *Proceedings of the CVPR*, pages 4721–4730, 2021. 5
- [31] Liang Peng, Senbo Yan, Boxi Wu, Zheng Yang, Xiaofei He, and Deng Cai. Weakm3d: Towards weakly supervised monocular 3d object detection. In *Proceedings of the ICLR*, 2022. 1, 2, 5, 6, 7
- [32] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *Proceedings of the CVPR*, pages 918–927, 2018. 5, 3
- [33] Rui Qian, Divyansh Garg, Yan Wang, Yurong You, Serge Belongie, Bharath Hariharan, Mark Campbell, Kilian Q Weinberger, and Wei-Lun Chao. End-to-end pseudo-lidar for image-based 3d object detection. In *Proceedings of the CVPR*, pages 5881–5890, 2020. 1
- [34] Zengyi Qin, Jinglu Wang, and Yan Lu. Monogrnnet: A geometric reasoning network for monocular 3d object localization. *Proceedings of the AAAI*, 2019. 5
- [35] Zengyi Qin, Jinglu Wang, and Yan Lu. Weakly supervised 3d object detection from point clouds. In *Proceedings of the ACM MM*, pages 4144–4152, 2020. 2, 5
- [36] Alec Radford, Jong Wook Kim, Hallacy, et al. Learning transferable visual models from natural language supervision. In *Proceedings of the ICML*, pages 8748–8763, 2021. 1, 2
- [37] Jiyong Rao, Brian Nlong Zhao, and Yu Wang. Probabilistic prompt distribution learning for animal pose estimation. In *Proceedings of the CVPR*, pages 29438–29447, 2025. 2, 3
- [38] Huiyang Shao, Xin Xia, Yuxi Ren, Xing Wang, and Xuefeng Xiao. Labridge: Text–image latent alignment framework via mean-conditioned ou process. In *Proceedings of the NIPS*, 2025. 1
- [39] Jan Skvrna and Lukas Neumann. Monosowa: Scalable monocular 3d object detector without human annotations. In *Proceedings of the ICCV*, 2025. 1
- [40] Runzhou Tao, Wencheng Han, Zhongying Qiu, Cheng-zhong Xu, and Jianbing Shen. Weakly supervised monocular 3d object detection using multi-view projection and direction consistency. In *Proceedings of the CVPR*, pages 7674–7683, 2023. 2, 5, 7
- [41] Mengxiao Tian, Xinxiao Wu, and Shuo Yang. Llm-enhanced action-aware multi-modal prompt tuning for image-text matching. In *Proceedings of the ICCV*, 2025. 1
- [42] Tai Wang, ZHU Xinge, Jiangmiao Pang, and Dahua Lin. Probabilistic and geometric depth: Detecting objects in perspective. In *Proceedings of the CoRL*, pages 1475–1485, 2022. 5, 6, 7, 2, 3
- [43] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *Proceedings of the CVPR*, pages 8445–8453, 2019. 1
- [44] Yan Wang, Bin Yang, Rui Hu, Ming Liang, and Raquel Urtasun. PLUMENet: Efficient 3D Object Detection from Stereo Images. In *Proceedings of the IROS*, pages 3383–3390, 2021.
- [45] Xinhao Weng and Kris Kitani. Monocular 3d object detection with pseudo-lidar point cloud. In *Proceedings of the CVPRW*, pages 0–0, 2019.
- [46] Hai Wu, Hongwei Lin, Xusheng Guo, Xin Li, Mingming Wang, Cheng Wang, and Chenglu Wen. Motal: Unsupervised 3d object detection by modality and task-specific knowledge transfer. In *Proceedings of the ICCV*, pages 6284–6293, 2025. 1
- [47] Lin Wu, Wei Wei, Peizhuo Yu, and Jianglin Lan. Open-vocabulary 3d affordance understanding via functional text enhancement and multilevel representation alignment. In *Proceedings of the ACM*, pages 7988–7997, 2025. 1
- [48] Yongjian Wu, Yang Zhou, Jiya Saiyin, Bingzheng Wei, and Yan Xu. Visual textualization for image prompted object detection. In *Proceedings of the ICCV*, 2025. 1
- [49] Shaoan Xie, Lingjing Lingjing, Yujia Zheng, Yu Yao, Zeyu Tang, Eric P. Xing, Guangyi Chen, and Kun Zhang. Smartclip: Modular vision-language alignment with identification guarantees. In *Proceedings of the CVPR*, pages 29780–29790, 2025. 1
- [50] Yung-Hsu Yang, Luigi Piccinelli, Mattia Segu, Siyuan Li, Rui Huang, Yuqian Fu, Marc Pollefeys, Hermann Blum, and Zuria Bauer. 3d-mood: Lifting 2d to 3d for monocular open-set object detection. In *Proceedings of the ICCV*, pages 7429–7439, 2025. 1
- [51] Yurong You, Yan Wang, Wei-Lun Chao, Divyansh Garg, Geoff Pleiss, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar++: Accurate depth for 3d object detection in autonomous driving. *arXiv preprint arXiv:1906.06310*, 2019. 1
- [52] Rui Yu, Runkai Zhao, Jiagen Li, Qingsong Zhao, Songhao Zhu, HuaiCheng Yan, and Meng Wang. Unleashing the potential of mamba: Boosting a lidar 3d sparse detector by using cross-model knowledge distillation. *arXiv preprint arXiv:2409.11018*, 2024. 1
- [53] Rui Yu, Runkai Zhao, Cong Nie, Heng Wang, HuaiCheng Yan, and Meng Wang. Future does matter: Boosting 3d object detection with temporal motion estimation in point cloud sequences. *arXiv preprint arXiv:2409.04390*, 2024. 1
- [54] Sergey Zakharov, Wadim Kehl, Arjun Bhargava, and Adrien Gaidon. Autolabeling 3d objects with differentiable

- rendering of sdf shape priors. In Proceedings of the CVPR, pages 12224–12233, 2020. [2](#), [5](#), [3](#)
- [55] Guowen Zhang, Junsong Fan, Liyi Chen, Zhaoxiang Zhang, Zhen Lei, and Lei Zhang. General geometry-aware weakly supervised 3d object detection. In Proceedings of the ECCV, 2024. [1](#), [2](#), [5](#), [6](#), [7](#)
- [56] Jiacheng Zhang, Jiaming Li, Xiangru Lin, Wei Zhang, Xiao Tan, Junyu Han, Errui Ding, Jingdong Wang, and Guanbin Li. Decoupled pseudo-labeling for semi-supervised monocular 3d object detection. arXiv:2403.17387, 2024. Accepted to CVPR 2024. [1](#)
- [57] Kun Zhang, Jingyu Li, Zhe Li, and S Kevin Zhou. Dh-set: Improving vision-language alignment with diverse and hybrid set-embeddings learning. In Proceedings of the CVPR, pages 24993–25003, 2025. [1](#)
- [58] Renrui Zhang, Han Qiu, Tai Wang, Xuanzhuo Xu, Ziyu Guo, Yu Qiao, Peng Gao, and Hongsheng Li. Monodetr: Depth-aware transformer for monocular 3d object detection. Proceedings of the ICCV, 2023. [5](#)
- [59] Weijia Zhang, Dongnan Liu, Chao Ma, and Weidong Cai. ODM3D: Alleviating Foreground Sparsity for Semi-Supervised Monocular 3D Object Detection. In Proceedings of the WACV, pages 7542–7552, 2024. [1](#)
- [60] Runkai Zhao, Yuwen Heng, Heng Wang, Yuanda Gao, Shilei Liu, Changhao Yao, Jiawen Chen, and Weidong Cai. Advancements in 3d lane detection using lidar point clouds: From data collection to model development. In ICRA, pages 5382–5388. IEEE, 2024. [1](#)
- [61] Runkai Zhao, Heng Wang, and Weidong Cai. Lanecmkt: Boosting monocular 3d lane detection with cross-modal knowledge transfer. In ACM MM, pages 4283–4291, 2024. [1](#)
- [62] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. Proceedings of the IJCV, 130(9):2337–2348, 2022. [2](#)
- [63] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. arXiv preprint arXiv:1904.07850, 2019. [5](#)

Supplementary Material

7. Overview

This supplementary material provides additional technical details and extended results supporting the proposed Visual-referred Probabilistic Prompt Learning (**VirPro**) framework. We first introduce the computation details of RoI contrastive learning objective (Sec. 8), which improves semantic coherence within scenes and enhances inter-scene discriminability in the latent space. We then summarize the pseudo-label generation pipelines of the weakly supervised baselines, WeakM3D and GGA (Sec. 9), clarifying their geometric, semantic, and alignment constraints. Implementation details for all components in these two baselines are provided in Sec. 10, followed by definitions of the clustering metrics used for latent space analysis (Sec. 11). We present additional quantitative (Sec. 12) and qualitative (Sec. 13) results. Finally, we demonstrate additional ablations in Sec. 14.

8. RoI Contrastive Learning

To reinforce the semantic coherence among co-occurring objects within the same scene in the latent space while discriminating scene-specific traits, we follow CAW3D [24] to adopt an object-level matching paradigm based on the traditional image-text contrastive learning. The associated loss is defined below.

Let $\mathbf{e}_i^{\text{txt}}$ denote the text embeddings of the i -th object normalized from prompt distributions $\hat{z}_{i,j}^{(t)}$ by maxpooling, and $\mathbf{e}_j^{\text{img}}$ denote the image embeddings of the j -th object, extracted from the Monocular 3D Encoder and spatially aligned using a 2D detector. The cosine similarity between these embeddings, along with the corresponding contrastive loss for the i -th sample, is formulated as follows:

$$\text{sim}_{ij} = \frac{\langle \mathbf{e}_i^{\text{txt}}, \mathbf{e}_j^{\text{img}} \rangle}{\|\mathbf{e}_i^{\text{txt}}\|_2 \cdot \|\mathbf{e}_j^{\text{img}}\|_2}, \ell_i = -\log \frac{\exp(\text{sim}_{ij}/\tau)}{\sum_{k=1}^N \exp(\text{sim}_{ik}/\tau)}, \quad (12)$$

$$\mathcal{L}_{\text{contrast}} = \frac{1}{N} \sum_{i=1}^N \ell_i, \quad (13)$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product, ℓ_i denotes the Cross-Entropy Loss between $\mathbf{e}_i^{\text{txt}}$ and $\mathbf{e}_i^{\text{img}}$, τ is a temperature scaling factor. N is the total number of objects in the batch.

9. Pseudo-Labels in Baselines

9.1. WeakM3D

WeakM3D [31] generates pseudo 3D labels by projecting LiDAR point clouds onto the corresponding 2D object masks of each image, thereby extracting Region-of-Interest (RoI) points. These RoI points are subsequently aligned with the predicted 3D bounding boxes for loss calculation. To handle the inherent challenges in this process, WeakM3D incorporates essential loss functions as follows: **Geometric Alignment Loss** aims to minimize the discrepancy caused by using center loss alone to determine the center of predicted 3D bounding box. The formulation is given as:

$$\begin{aligned} \mathcal{L}_{\text{geo}} &= \|\mathbf{p}_i - \hat{\mathbf{p}}_i\|_1 \\ &= \left\| \mathbf{p}_i - \text{Intersect} \left(\overrightarrow{\mathbf{c} \rightarrow \mathbf{p}_i}, \hat{b}_{3d} \right) \right\|_1, \end{aligned} \quad (14)$$

where \mathbf{p}_i denotes the i -th RoI point, and $\hat{\mathbf{p}}_i$ is computed as the intersection between the ray from predicted 3D center \mathbf{c} to \mathbf{p}_i and the surface of predicted 3D bounding box \hat{b}_{3d} . **Ray Tracing Loss** is designed to mitigate surface uncertainty associated with RoI points by enforcing their accurate correspondence to the correct object surface, thereby enhancing geometric consistency and localization precision. The loss is formulated as:

$$\mathcal{L}_{\text{ray}} = \begin{cases} \|\mathbf{p}_i - \mathbf{p}_i^{(r)}\|_1, & \text{if } \text{Ray}(\mathbf{p}_{\text{cam}} \rightarrow \mathbf{p}_i) \cap \hat{b}_{3d} \neq \emptyset, \\ 0, & \text{otherwise,} \end{cases} \quad (15)$$

and $\mathbf{p}_i^{(r)}$ denotes the intersection point on the predicted 3D bounding box \hat{b}_{3d} that is closest to the camera along the ray from the camera center \mathbf{p}_{cam} through the RoI point \mathbf{p}_i .

Point-wise Balancing Loss compensates for non-uniform point cloud distributions by ensuring that sparse yet significant points are not overlooked, thereby improving the completeness of the overall object detection process. For each point \mathbf{p}_i , we compute its local neighborhood density as:

$$w_i = |\{\mathbf{p}_j \mid \|\mathbf{p}_i - \mathbf{p}_j\|_2 < R, j \neq i\}|, \quad (16)$$

where w_i is the neighborhood count of point \mathbf{p}_i , and R is a predefined distance threshold for determining neighborhood connectivity within the RoI point set. The final 3D supervision loss is then weighted inversely by this density, and formulated as:

$$\mathcal{L}_{3D} = \frac{1}{M} \sum_{i=1}^M \frac{1}{w_i} (\mathcal{L}_{\text{geo},i} + \mathcal{L}_{\text{ray},i} + \lambda \mathcal{L}_{\text{center},i}), \quad (17)$$

where M denotes the total number of RoI points and λ is a scalar hyperparameter used to balance the contribution of the center loss term.

By incorporating this series of 3D loss functions, the monocular detector is guided to acquire enhanced spatial awareness through supervision from 3D pseudo labels, thereby improving the accuracy of 3D object detection.

9.2. GGA

GGA [55] presents a unified weakly supervised 3D detection framework that integrates geometric constraints, 2D–3D consistency, and static textual prompts. For each 2D frustum, point clouds are cropped as *In-Box Points* and fed into a point-cloud backbone with a proposal head to estimate 3D bounding boxes, class scores, and auxiliary pseudo-scores. To strengthen geometric reliability, GGA incorporates the following components:

Boundary Projection Loss (BPL) enforces 2D-3D consistency by aligning each predicted 3D bounding box with its corresponding 2D annotation. Given the calibrated camera model, the eight corners of a predicted 3D box are first projected onto the image plane, and the minimum enclosing rectangle of these projected points forms a predicted 2D box. Formally, let $\mathbf{C}^p = \text{Proj}(\text{Corners}(B_{3d}^p))$ denote the set of projected corners, and define the predicted 2D bounds as $\mathbf{b}_{2d}^p = [\min(\mathbf{C}_x^p), \min(\mathbf{C}_y^p), \max(\mathbf{C}_x^p), \max(\mathbf{C}_y^p)]$. The BPL then minimizes the discrepancy between \mathbf{b}_{2d}^p and the ground-truth 2D box $\mathbf{b}_{2d} = [x_{\min}, y_{\min}, x_{\max}, y_{\max}]$ through an L_1 penalty:

$$\mathcal{L}_{\text{BPL}} = \|\mathbf{b}_{2d}^p - \mathbf{b}_{2d}\|_1. \quad (18)$$

This loss encourages the projected 3D box to tightly align with its 2D counterpart, thereby constraining the 3D box location and scale from the perspective of the image space.

Semantic Ratio Loss (SRL) leverages simple yet effective shape priors derived from GPT-4 to regularize the predicted 3D box dimensions. Instead of relying on handcrafted geometric rules or synthetic statistics, SRL uses the observation that the bird’s-eye-view width–height ratio provides sufficient semantic cues for constraining object shapes. Let the predicted 3D box be parameterized by $(x, y, z, l, w, h, \alpha)$. We compute the predicted ratio using the shorter side over the longer side:

$$r^p = \frac{\min(l, w)}{\max(l, w)}. \quad (19)$$

Given a category-level prior ratio r obtained from GPT-4, SRL penalizes deviations between the predicted and prior ratios using an L_1 loss:

$$\mathcal{L}_{\text{SRL}} = L(r^p, r), \quad (20)$$

where $L(\cdot)$ denotes the L_1 distance. By providing a lightweight semantic constraint on object shape, SRL helps the model converge faster and improves the stability of 3D box estimation.

Points-to-Box Alignment Loss (PAL) exploits the spatial relationship between the predicted box and the In-Box-Points to impose geometric supervision on predicted 3D boxes in the absence of full annotations. Since a valid 3D box should enclose the corresponding foreground points in the BEV space, PAL first computes the distances from each point to the four edges of the predicted BEV box. Let (l, w) be the predicted length and width, and let $(d_i^1, d_i^2, d_i^3, d_i^4)$ denote the distances from point i to the left, right, top, and bottom edges, respectively. A soft constraint encourages points to lie inside the predicted box through a ReLU activation $\phi(\cdot)$, yielding:

$$\mathcal{L}_{\text{PAL}_1} = \sum_{i=1}^N \left(\sum_{j \in \{1,2\}} \phi(d_i^j - \frac{l}{2}) + \sum_{k \in \{3,4\}} \phi(d_i^k - \frac{w}{2}) \right). \quad (21)$$

However, RGB-D and LiDAR observations often capture only one visible side of an object, causing points to cluster around box boundaries. To leverage this property for implicit supervision, PAL further minimizes the shortest edge-wise distance for each point, producing a tighter alignment:

$$\mathcal{L}_{\text{PAL}_2} = \sum_{i=1}^N \min(d_i^1, d_i^2, d_i^3, d_i^4). \quad (22)$$

Together, $\mathcal{L}_{\text{PAL}_1}$ and $\mathcal{L}_{\text{PAL}_2}$ constrain the predicted BEV box to geometrically align with the foreground point distribution, providing effective supervision for learning accurate 3D box dimensions and positions.

The overall training objective combines the proposed geometric, semantic, and alignment constraints with standard detection losses. Specifically, the final loss is defined as:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\text{BPL}} + \lambda_2 \mathcal{L}_{\text{SRL}} + \lambda_3 (\mathcal{L}_{\text{PAL}_1} + \mathcal{L}_{\text{PAL}_2}) + \lambda_4 \mathcal{L}_{\text{score}} + \lambda_5 \mathcal{L}_{\text{cls}}, \quad (23)$$

where λ_{1-5} are balancing weights. $\mathcal{L}_{\text{score}}$ denotes the objectness heatmap regression loss used in CenterPoint and the centerness loss in FCAF3D, while \mathcal{L}_{cls} is the cross-entropy loss for classification. The predicted 3D boxes are subsequently treated as pseudo labels to train the final 3D detector PGD [42] in a fully supervised manner.

10. Implementation Details

WeakM3D optimized using the Adam optimizer [20] with an initial learning rate of 10^{-4} . The network is trained for

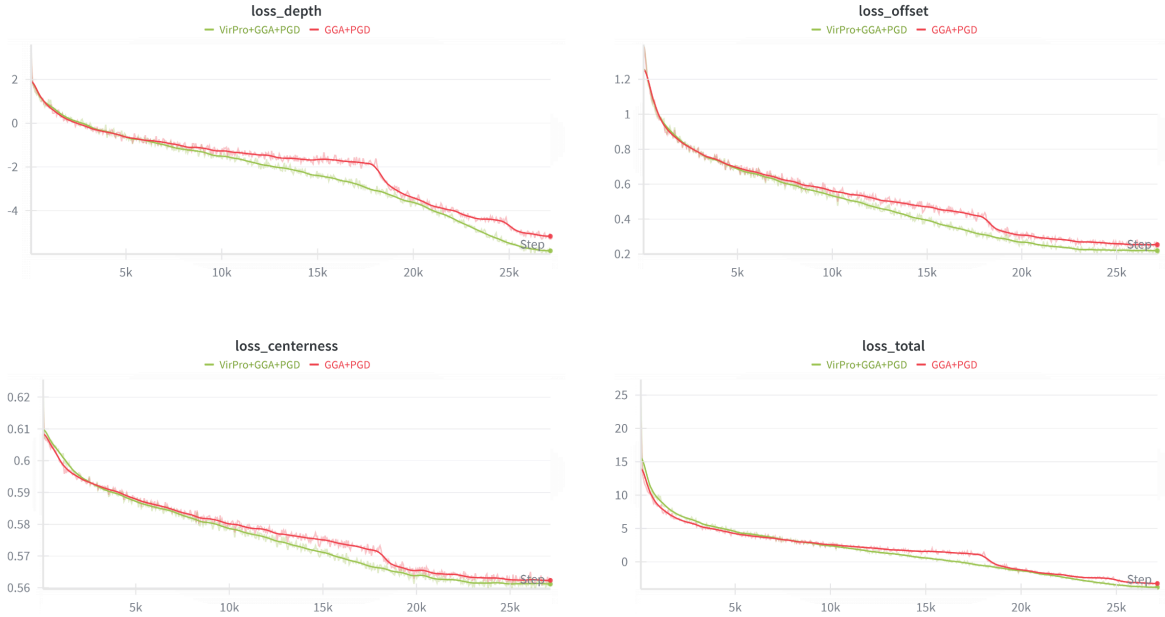


Figure 6. **Training loss comparison between VirPro+GGA+PGD and GGA+PGD.** The depth loss supervises the predicted 3D depth of the object to ensure accurate distance estimation from the camera. The offset loss constrains the projected 2D center offset. The centerness loss encourages confident predictions near object centers while suppressing noisy peripheral responses. The total loss is the weighted sum of all objectives.

50 epochs. To initialize the object point cloud, WeakM3D adopts an off-the-shelf 2D detector FPN [32]. For car-sized objects, the frozen dimensions are empirically set to height 1.6 m, width 1.8 m, and length 4.0 m. The point density threshold in Eq. 17 is fixed to 0.4. Following the 2D-3D alignment strategy of Brazil and Liu [2], the y -coordinate adjustment is applied to improve geometric consistency. The image backbone is ResNet34 [14].

GGA adopts CenterPoint [54] as the backbone networks. The framework is implemented in MMDetection3D [10] and optimized using the AdamW optimizer. The RANSAC thresholds for plane fitting are set to 0.2. Following the configuration of CenterPoint on KITTI, GGA omits \mathcal{L}_{cls} and assigns $\lambda_{1-4} = 0.3, 0.1, 0.1, 5$. The framework is trained for 120 epochs. The image backbone of PGD [42] is the ResNet101 [14].

11. Latent Space Evaluation Metrics

To qualitatively assess the impact of **VirPro** on the structure of latent embeddings, we perform a clustering-based analysis of RoI visual embeddings produced by our model and CAW3D [24], a deep semantic supervision work with hand-crafted prompts, after stage 1 training. Specifically, we extract all RoI features from the validation set and treat each as an individual point, grouped

by its originating scene. We adopt two standard clustering metrics: the **Calinski–Harabasz (CH)** index and the **average Silhouette Score (\bar{s})**. The CH index measures the ratio of between-cluster dispersion to within-cluster dispersion, with higher values indicating better cluster separation and compactness. It is defined as:

$$CH = \frac{\text{Tr}(B_k)}{\text{Tr}(W_k)} \cdot \frac{n - k}{k - 1} \quad (24)$$


where $\text{Tr}(B_k)$ and $\text{Tr}(W_k)$ denote the between-cluster and within-cluster dispersion, respectively; n is the number of samples and k is the number of clusters.

The Silhouette Score \bar{s} evaluates the consistency within clusters by comparing the intra-cluster distance $a(i)$ and the nearest-cluster distance $b(i)$ for each sample i :

$$s(i) = \frac{b(i) - a(i)}{\max(a(i), b(i))} \quad \bar{s} = \frac{1}{n} \sum_{i=1}^n s(i), \quad (25)$$

where $s(i)$ is the silhouette coefficient of sample i , $a(i)$ is the average distance to all other points in the same cluster, and $b(i)$ is the average distance to points in the nearest neighboring cluster.

GGA+PGD

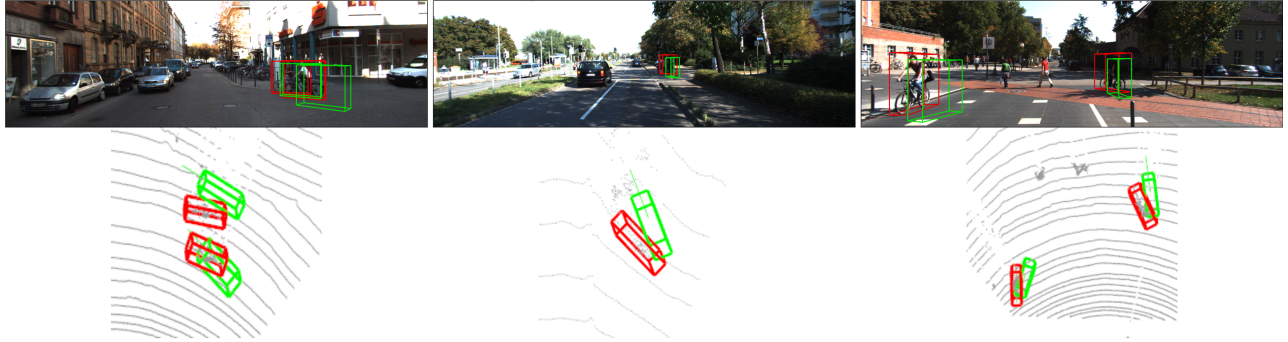


VirPro+GGA+PGD



Figure 7. Qualitative results on the KITTI validation set comparing ours on the "Pedestrian" category. Predicted boxes are rendered in green, and ground-truth boxes are shown in red.

GGA+PGD



VirPro+GGA+PGD




Figure 8. Qualitative results on the KITTI validation set comparing ours on the "Cyclist" category. Predicted boxes are rendered in green, and ground-truth boxes are shown in red.

Table 7. **Comparison on the KITTI validation set (Pedestrian category).** We report validation performance using the AP₄₀ at an IoU threshold of 0.5. The best results are highlighted in red.

| Pedestrian Validation | AP _{BEV} | | | AP _{3D} | | |
|-----------------------|-------------------|-------------|-------------|------------------|-------------|-------------|
| | Easy | Mod | Hard | Easy | Mod | Hard |
| GGA+PGD [42, 55] | 3.94 | 2.94 | 2.45 | 2.48 | 1.86 | 1.33 |
| VirPro+GGA+PGD | 5.19 | 4.11 | 3.15 | 3.38 | 2.52 | 1.91 |

Table 8. **Comparison on the KITTI validation set (Cyclist category).** We report validation performance using the AP₄₀ at an IoU threshold of 0.5. The best results are highlighted in red.

| Cyclist Validation | AP _{BEV} | | | AP _{3D} | | |
|-----------------------|-------------------|-------------|-------------|------------------|-------------|-------------|
| | Easy | Mod | Hard | Easy | Mod | Hard |
| GGA+PGD [42, 55] | 2.32 | 1.06 | 0.73 | 2.23 | 0.77 | 0.73 |
| VirPro+GGA+PGD | 4.13 | 2.03 | 1.67 | 3.11 | 1.86 | 1.51 |

Table 9. **Comparison on the KITTI test set (Pedestrian category).** GGA+PGD is the baseline method using weak 2D-3D alignment and textual prompts generated from LLM for weak supervision. The best results are highlighted in red.

| Pedestrian Test | AP _{BEV} | | | AP _{3D} | | |
|-----------------------|-------------------|-------------|-------------|------------------|-------------|-------------|
| | Easy | Mod | Hard | Easy | Mod | Hard |
| GGA+PGD [42, 55] | 0.87 | 0.61 | 0.46 | 0.57 | 0.36 | 0.27 |
| VirPro+GGA+PGD | 1.59 | 1.09 | 0.91 | 1.07 | 0.70 | 0.51 |

Table 10. **Comparison on the KITTI test set (Cyclist category).** GGA+PGD is the baseline method using weak 2D-3D alignment and textual prompts generated from LLM for weak supervision. The best results are highlighted in red.

| Cyclist Test | AP _{BEV} | | | AP _{3D} | | |
|-----------------------|-------------------|-------------|-------------|------------------|-------------|-------------|
| | Easy | Mod | Hard | Easy | Mod | Hard |
| GGA+PGD [42, 55] | 0.84 | 0.28 | 0.30 | 0.69 | 0.25 | 0.28 |
| VirPro+GGA+PGD | 1.29 | 0.52 | 0.43 | 1.08 | 0.45 | 0.49 |

Table 11. **Performances on nuScenes val set of "Car" class.**

| Method | AP _↑ | ATE _↓ | ASE _↓ | AAE _↓ |
|-----------------------|-----------------|------------------|------------------|------------------|
| WeakM3D [31] | 0.214 | 0.814 | 0.234 | 0.682 |
| SKD-WM3D [16] | 0.242 | 0.795 | 0.231 | 0.659 |
| GGA+PGD [55] | 0.227 | 0.661 | 0.173 | 0.238 |
| Virpro+GGA+PGD | 0.239 | 0.626 | 0.158 | 0.206 |

12. Quantitative Results

Tabs. 9, 7, 10, and 8 present that consistent with the trends observed on the "Car" category, our VirPro pretraining paradigm delivers clear and steady improvements over the baseline GGA+PGD [42, 55] on both "Pedestrian" and "Cyclist" categories under all difficulty levels. These results demonstrate that the proposed visually referred probabilistic prompts provide contextual and more

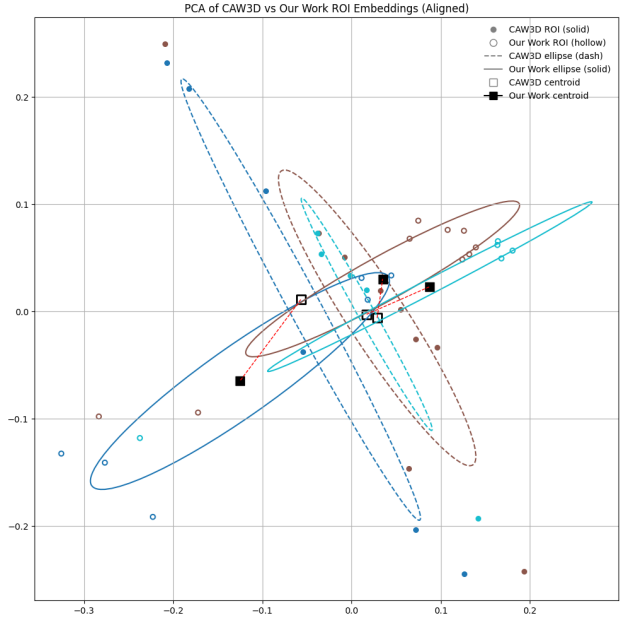


Figure 9. **PCA Visualization of RoI Embeddings from CAW3D and Our Proposed VirPro.** We compare the RoI embeddings distribution projected via PCA from CAW3D and our proposed VirPro, where our work exhibits better-separated clusters across scenes, indicating stronger scene discrimination and improved latent space structuring.

Table 12. **Ablations on Gaussian Sampling of Prompts.** G.S. denotes Gaussian Sampling. Both prompts are generated by the same prompt bank and visual conditioning.

| Gaussian sampling | AP _{BEV} | | | AP _{3D} | | |
|------------------------|-------------------|--------------|--------------|------------------|--------------|--------------|
| | Easy | Mod | Hard | Easy | Mod | Hard |
| Prompts w/o G.S. | 57.36 | 40.39 | 36.76 | 52.17 | 37.86 | 31.49 |
| Prompts w/ G.S. | 60.11 | 42.95 | 37.50 | 54.72 | 39.49 | 33.32 |

Table 13. **Ablations on the Quality of 2D Annotation.**

| 2D GT Annotation | AP _{BEV} | | | AP _{3D} | | |
|--------------------------|-------------------|--------------|--------------|------------------|--------------|--------------|
| | Easy | Mod | Hard | Easy | Mod | Hard |
| w/o 2D GT finetune | 59.31 | 43.32 | 37.17 | 52.62 | 38.65 | 33.51 |
| w/ 2D GT finetune | 60.11 | 42.95 | 37.50 | 54.72 | 39.49 | 33.32 |

informative supervisory signals, enabling stronger 3D localization and shape estimation for both Pedestrian and Cyclist instances as well. In addition, we evaluate our method on nuScenes dataset. As shown in Tab. 11, VirPro yields consistent gains across nuScenes' diverse scenes which demonstrates improved robustness beyond KITTI. We follow SKD-WM3D [16] to evaluate "Car" on validation set, since test is unreported. We train and validate our model on "CAM_FRONT" split.

13. Qualitative Results

3D Visualizations for Cyclist and Pedestrian. We provide qualitative 3D visualizations on the KITTI *validation* set for the Pedestrian and Cyclist categories. As illustrated in Fig. 7 and 8, VirPro+GGA+PGD generates noticeably more accurate and spatially coherent 3D bounding boxes than the GGA+PGD baseline. Across diverse urban scenes, our predictions exhibit improved scale estimation, orientation stability, and depth reasoning, yielding tighter alignment between predicted boxes (green) and ground-truth annotations (red). The gains are particularly clear for small, heavily occluded, and cluttered instances, highlighting the effectiveness of visually enriched probabilistic prompts.

Latent Space Distribution. As shown in Fig. 9, PCA on RoI embeddings from three randomly selected scenes in KITTI validation split shows more clearly separated clusters under VirPro. Specifically, the CAW3D embeddings (solid markers) form highly overlapping clusters, with large covariance ellipses indicating weak scene discrimination and significant intra-scene variation. In contrast, the RoI embeddings generated by VirPro (hollow markers) exhibit sharper and well-separated clusters across scenes. Moreover, the centroids of VirPro embeddings align more distinctly between scenes, suggesting improved inter-scene separability and a more structured latent space. The result verifies that visually referred probabilistic prompts yield a more structured and semantically discriminative latent space.

Train Loss Curve. Fig. 6 shows that VirPro demonstrates a noticeably smoother and more stable optimization trajectory, with reduced oscillation and consistently lower loss values throughout pretraining. This reflects VirPro’s strong guidance on geometric reasoning and convergence behavior, enabling a more stable and steady learning compared to the baseline. This improvement is mainly attributed to that our proposed VirPro encourages smoother modality alignment. Therefore, the model receives soft, probabilistic guidance rather than rigid supervision, which leads to steady convergence and mitigates training noise.

14. Ablation Experiments

Quality of 2D Annotation VirPro is designed to be robust because visually injected probabilistic prompts capture cross-scene diversity and uncertainty for each RoI, reducing sensitivity to imperfect 2D boxes. Empirically, Tab. 13 shows additional fine-tuning with 2D GT after Stage 1 yields only moderate effects.

Effectiveness of Gaussian Sampling. We added a controlled ablation on the KITTI benchmark by removing Gaussian sampling. Tab. 12 validates the Gaussian sampling with consistent gains across BEV and 3D.