




Self-Adapting Large Visual-Language Models to Edge Devices across Visual Modalities

Kaiwen Cai¹, Zhekai Duan¹, Gaowen Liu², Charles Fleming², and
Chris Xiaoxuan Lu^{3*}

¹ University of Edinburgh, UK

² Cisco Research, USA

³ University College London, UK

Abstract. Recent advancements in Vision-Language (VL) models have sparked interest in their deployment on edge devices, yet challenges in handling diverse visual modalities, manual annotation, and computational constraints remain. We introduce **EdgeVL**, a novel framework that bridges this gap by seamlessly integrating dual-modality knowledge distillation and quantization-aware contrastive learning. This approach enables the adaptation of large VL models, like CLIP, for efficient use with both RGB and non-RGB images on resource-limited devices without the need for manual annotations. **EdgeVL** not only transfers visual language alignment capabilities to compact models but also maintains feature quality post-quantization, significantly enhancing open-vocabulary classification performance across various visual modalities. Our work represents the first systematic effort to adapt large VL models for edge deployment, showcasing up to 15.4% accuracy improvements on multiple datasets and up to 93-fold reduction in model size. Code available at <https://github.com/ramdrop/edgevl>.

1 Introduction

In recent years, there has been a surge of interest in the development of Vision-Language (VL) models capable of conducting integrated reasoning across visual and textual data. Prominent large-scale VL models, such as CLIP [40], typically employ distinct visual and text encoders. These encoders embed images and text into a common feature space, enabling direct comparison across two modalities. By evaluating the degree of similarity between the image embeddings and various potential text embeddings, these models facilitate zero-shot and open-vocabulary visual recognition, including image classification [19], semantic segmentation [15], and object detection [36].

However, three significant challenges hinder the deployment of VL models on edge devices: (i) generalization to diverse visual modalities, (ii) label scarcity in the wild settings, and (iii) on-device resource limitation. *Firstly*, an edge device often comes equipped with multiple sensors beyond standard RGB cameras, such

* Corresponding author. Email: xiaoxuan.lu@ucl.ac.uk

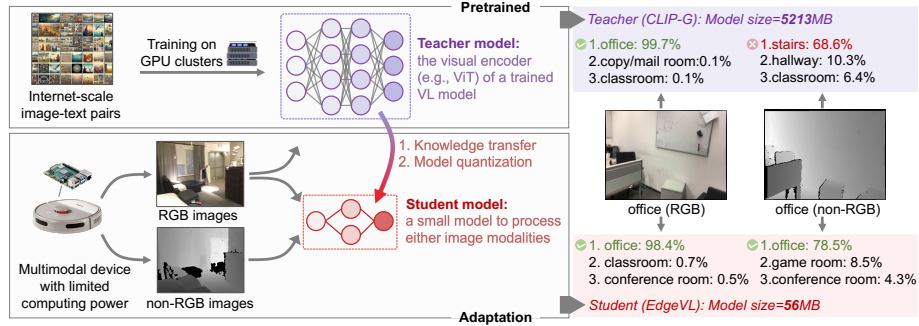


Fig. 1: The adaptation problem of large visual language model to edge devices across visual modalities. We use a resource-constrained cleaning robot as the edge device for illustration. The robot has a *co-located* RGB and depth cameras, generating many paired images without scene labels. Using RGB-depth pairs as the inputs and the pre-trained image encoder in CLIP as the teacher, **EdgeVL** is designed to transfer the knowledge to a small student encoder without labels or human intervention. After this learning process, the student encoder can agnostically process *either* image modalities for open-vocabulary scene classification on the device.

as depth sensors and infrared cameras. These are indispensable in edge devices like field robots or smart doorbells for visual comprehension under challenging lighting conditions, like darkness, smoke, or fog. Despite this, the visual encoders in most large VL models are predominantly tailored to RGB images, leaving the adaptability of these models to alternative inputs such as depth or infrared images largely unexplored. *Secondly*, while edge devices can generate a vast amount of images, these images are often unlabeled, presenting a significant obstacle in the wild where human-specified annotations are unavailable. This absence of labels prevents the straightforward application of model fine-tuning on annotated datasets. *Thirdly*, even if the transfer of knowledge across different visual modalities becomes achievable, the substantial computational requirements of the visual encoders (e.g., ViT [10] used by CLIP [40]) render them impractical for edge devices, which are typically constrained by limited memory and TOPS (Tera Operations Per Second) performance.

To overcome these challenges, a novel framework is needed that can adapt the VL embedding prowess of large models to non-RGB images without relying on human annotations, while also minimizing its computational footprint to suit the capabilities of edge devices. Existing literature tends to address these domains in isolation, focusing either on cross-modal knowledge transfer [17, 24, 46] or on model compression (e.g., quantization [27], pruning [33] and distillation [47]). However, the interplay and potential synergy between these two areas remain largely unexplored, not to mention the impact of label scarcity. As evident in our empirical results (see Tab. 3), brute-forcing integrating the two modules leads to an obvious performance drop for large VL models.

In this work, we propose **EdgeVL**, a streamlined two-stage adaptation framework that seamlessly integrates knowledge transfer with model compression. Initially, **EdgeVL** utilizes a dual-modality knowledge distillation process, leveraging a pre-trained visual encoder as the teacher model. This process distills knowledge to a more compact student model. The student model is designed to handle both RGB and non-RGB images, ensuring the alignment of visual features with textual representations similar to those found in conventional large VL models. This initial stage significantly enhances model efficiency through architectural optimization. Subsequently, to further boost efficiency and the efficacy of extracted features for edge deployment, the framework applies quantization-aware training (QAT) augmented with a novel contrastive learning loss. This sophisticated approach culminates in a low-bitrate visual encoder model, optimized for edge devices, which demonstrates superior performance in open-vocabulary classification tasks for both RGB and non-RGB images. Fig. 1 depicts the conceptual framework of **EdgeVL**. We summarise our main contributions as follows:

- **EdgeVL** is the first framework to systematically address the adaptation of large VL models for edge devices, facilitating their use with diverse visual modalities without relying on manual annotations.
- We introduce a method to transfer visual language alignment from pre-trained VL models to compact visual models for both RGB and non-RGB images, eliminating annotation needs.
- We incorporate quantization-aware training enhanced by a contrastive learning loss. This approach not only maintains the quality of feature representation post-quantization but also significantly improves the model’s discriminative ability across diverse visual modalities.
- We highlight **EdgeVL**’s gain in accuracy across multiple datasets and detail its efficiency improvements on diverse GPU tiers.

2 Related Work

2.1 Open Vocabulary Classification

Recent advances in VL models [28, 40] have made it possible to obtain consistent image and text representations. In open vocabulary object detection tasks [11, 16, 53], the classification head is replaced with a feature projection head, whose output is compared with the text embeddings of all candidate classes to obtain the final prediction. In open vocabulary semantic segmentation, FreeSeg [39] includes a task-specific prompt in the text queries and trains an encoder to be aligned with the CLIP text embeddings. OpenScene [38] trained a 3D model to learn point embeddings that mimic the pixel embeddings from the CLIP image encoder. Although these open vocabulary classification methods [9, 20, 25, 30] have achieved promising results, they are focusing on the RGB image modality. In this work, we aim to address the open vocabulary scene classification problem on various modalities beyond RGB images.

2.2 Cross-modality Knowledge Distillation

Knowledge distillation facilitates the transfer of expertise between models through various strategies such as response distillation [22], feature distillation [41], and relation distillation [7]. In the realm of cross-modality knowledge distillation, methods like those presented by [23] and [17] transfer knowledge by sharing weights between models of different modalities. [46] employs a strategy where multiple student models are encouraged to align their predictions with a teacher model, enhancing cross-modal action recognition. CMKD [24] introduces a combination of feature and response distillation to propagate knowledge from the LiDAR modality to the RGB modality. Recently, CLIP has also been explored for transferring knowledge from 2D images to 3D scenes [3, 51, 54]. Unlike these methods tailored for labeled datasets, our work uniquely addresses the challenge of knowledge distillation using unlabeled pairs of RGB and non-RGB images.

2.3 Model Quantization

Model quantization techniques fall into two main categories: Post Training Quantization (PTQ) and Quantization Aware Training (QAT). PTQ methods, such as those outlined by [2, 8, 34, 37, 50], involve quantizing the weights and activations of a model after training, without retraining the model. For instance, LLM.int8 [8] focuses on identifying weight outliers and applying distinct quantization methods for inliers and outliers, while SmoothQuant [50] adjusts the scale of weights and activations to maintain quantization precision for both outliers and inliers. When training data or validation data is not accessible, ZeroQ [2] synthesizes input data using Batch Normalization layer statistics to evaluate layer sensitivity. On the contrary, QAT, as explored by [5, 12, 27, 29, 31], is employed when PTQ does not satisfy accuracy requirements. This approach includes inserting fake quantization nodes during training, allowing the model to adapt to quantization errors and learn a more robust representation. Notably, LSQ [12] improves quantization performance through a learnable scale factor, and EWGS [31] enhances gradient estimation with a weighted gradient scaling method. In this work, we investigate quantization strategy for transferring cross-modal knowledge from large VL models.

3 Methodology

3.1 Preliminary on Open Vocabulary Classification

Large-scale VL models like CLIP [40], which include image and text encoders, are trained on over 400 million image-text pairs to map them into a shared feature space, optimizing true pair closeness and false pair distance through contrastive training. This approach enables CLIP to perform zero-shot and open-vocabulary classifications on unseen class labels during inference by assessing the similarity between image and text embeddings. However, despite its success with RGB images, CLIP’s visual encoder underperforms in zero-shot classification tasks

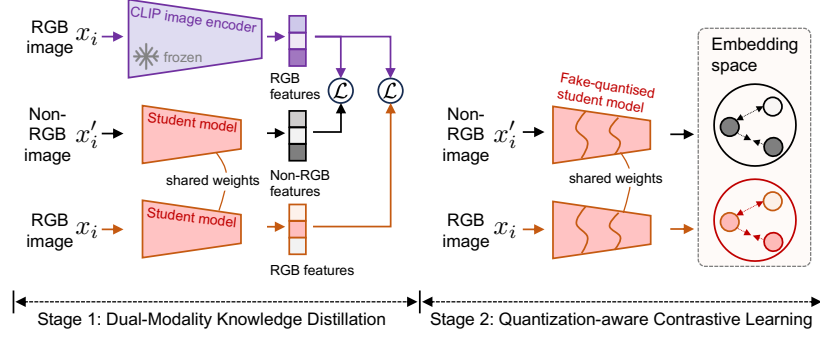


Fig. 2: Overall architecture of our proposed method. In stage-1, we distill the knowledge from the pre-trained visual encoder to the student model. In stage-2, we first fake-quantize the pretrained student model, then use contrastive learning to refine the student model.

with non-RGB images. Performance data (see Tab. 1 later) indicates a significant accuracy disparity between RGB and non-RGB images (*e.g.*, Depth and Infrared (IR)), with an example from the ScanNet dataset showing an approximately 8-fold decrease in accuracy for depth images compared to RGB.

3.2 Problem Definition

Motivated by the preliminary, we consider adapting an open-vocabulary classifier to edge devices with different image sensors. Let the training set for adaptation be denoted by $\mathcal{D}_{train} = \{(x_i, x'_i)\}_{i=1}^N$, where x_i represents the i^{th} RGB image, x'_i is its non-RGB image counterpart. Co-located cameras continuously collect the above pair of images on an edge device (*e.g.*, a mobile robot), and we use N pairs of such images for model adaptation. Note that, there are no labels for these images. Assume an RGB image encoder Φ_{img} from a pre-trained large VL model is available. The adaptation goal for **EdgeVL** is to roll out a modality-agnostic and efficient image encoder Φ_{img}^{edge} so that the following visual features can be approximately the same:

$$\Phi_{img}(x_i) \approx \Phi_{img}^{edge}(x_i) \approx \Phi_{img}^{edge}(x'_i) \quad (1)$$

The entire training of **EdgeVL** does not need human annotation or manual labels.

When it comes to the inference stage, assume there is a test set denoted as $\mathcal{D}_{test} = \{(x_i, x'_i), y_i \in \mathcal{C}\}_{i=1}^N$, where y_i designates the class label for each image pair, and \mathcal{C} encompasses all possible open-vocabulary classes. By using the text encoder of the same pre-trained large VL model and the developed image encoder Φ_{img}^{edge} , the open-vocabulary class prediction can be translated into the maximum feature similarity inference problem:

$$\hat{y}_i = \arg \max_{Y \in \mathcal{C}} \Phi_{img}^{edge}(x_i)^\top \Phi_{text}(Y), \quad \hat{y}'_i = \arg \max_{Y \in \mathcal{C}} \Phi_{img}^{edge}(x'_i)^\top \Phi_{text}(Y) \quad (2)$$

Ideally, if Φ_{img}^{edge} is a well-adapted image encoder, the predicted classes \hat{y}_i and \hat{y}'_i align closely with the true classes y_i and the inference efficiency is boosted⁴. As illustrated in Fig. 2, **EdgeVL** consists of a two-stage adaptation framework progressively solving the above problem: i. dual-modality knowledge distillation module ($\Phi_{img} \rightarrow \Phi_{img}^{stu}$), and ii. quantization-aware contrastive learning module ($\Phi_{img}^{stu} \rightarrow \Phi_{img}^{edge}$).

3.3 Stage-1: Dual-Modality Knowledge Distillation

The first stage of **EdgeVL** aims to distill the image features from the teacher image encoder in a pre-trained VL model to a student image encoder for dual modalities. For a trained large VL model like CLIP, while its image encoder has certain zero-shot transferability on unseen data, it still has failure cases where it provides noninformative or noisy features as supervision signals. Removing the noisy samples and their impact on the distillation is beneficial, but the challenge arises when manual sample checking is costly in the wild.

Automatic Dataset Curation We thus introduce an automated data selection mechanism as a precursor to feature distillation, guided by the innate comparison capabilities of VL models for images and texts. This approach leverages the VL models’ ability to generate auxiliary information beneficial for sample selection without human intervention. Our approach initiates by creating a ‘superset of labels’, \mathcal{S} , through a ChatGPT-4 [1] engine. Owing to the open-vocabulary feature of VL models, this label superset acts as an extensive repository of potential labels, applicable across various contexts. (See supplementary for the generated label superset)

We next encode this label superset into text features using a text encoder, Φ_{text} . Concurrently, an image encoder, Φ_{img} , processes unlabelled RGB images from edge devices to extract visual features. For an RGB image x_i , we assign a confidence score c_i based on the highest image-text similarity with texts in \mathcal{S} :

$$c_i = \max\{s_k \mid s_k = \frac{e^{\Phi_{img}(x_i)^\top \Phi_{text}(y_k)}}{\sum_k e^{\Phi_{img}(x_i)^\top \Phi_{text}(y_k)}}, k = 1, 2, \dots, |\mathcal{S}|\}, \quad (3)$$

Our observations reveal that images with noisier and less informative features typically yield lower confidence scores. Such images are deemed unsuitable for feature distillation and are excluded from further processing. Concretely, we employ a predefined threshold, τ_c , to curate the dataset. Only RGB images scoring above this threshold are retained in the training dataset, $\mathcal{D}_{train} = \{(x_i, x'_i)\}_{i=1}^{N_c}$, which also includes their non-RGB counterparts collected simultaneously by the device. For notation simplicity, the training set \mathcal{D}_{train} hereafter refers to the automatically curated dataset.

⁴ Class prediction varies by case, but classes are typically chosen for specific downstream tasks. Take the example of a mobile robot performing perception tasks: the image is the only live input, while class labels and $\Phi_{text}(Y)$ features are predetermined. Thus, image encoding is the main computation bottleneck in our inference.

Feature Distillation With the curated dataset, we are poised to proceed with the feature distillation process. This crucial phase empowers a compact student encoder (*e.g.*, Swin-T [35]) to effectively extract robust image embeddings from both RGB and non-RGB images, via referencing the large teacher encoder (*e.g.*, ViT-T in CLIP [40]). Our approach diverges from conventional methods that tailor a compact student model solely for non-RGB or RGB image inputs. Instead, we advocate for a unified image encoder capable of seamlessly processing either image type through weight sharing. This innovative dual-modality student encoder not only simplifies the model architecture but also significantly reduces model storage requirements by at least half on edge devices. Illustrated in Figure 2, for each pair of RGB and non-RGB images in \mathcal{D}_{train} , we align the features extracted from both RGB and non-RGB images by the student model with those extracted from RGB images by the teacher image encoder. This alignment is predicated on the understanding that both image types represent the same scene, thereby necessitating that the student model generates consistent image features that resonate with those from a pre-trained VL model (*i.e.*, teacher) like CLIP. We designate the student encoder as Φ_{img}^{stu} and focus on minimizing the discrepancy between the student model’s features and the teacher’s image features through our feature distillation loss function:

$$\mathcal{L}_d = d(\Phi_{img}(x), \Phi_{img}^{stu}(x')) + d(\Phi_{img}(x), \Phi_{img}^{stu}(x)). \quad (4)$$

Here, d signifies the distance function. We follow [13] and use the L1 distance function. Through this loss function, we aim to closely align the student’s feature representations with those of the teacher model, thereby ensuring the student model’s proficiency across both RGB and non-RGB modalities.

3.4 Stage-2: Quantization-aware Contrastive Learning

Given the dual-modality student encoder Φ_{img}^{stu} , the next step is to further enhance its efficiency by transforming it to a quantized low-bit model Φ_{img}^{edge} on par with the resources available on edge devices. The challenge, however, is how to preserve the feature expressiveness after quantization is applied.

QAT Meets Contrastive Learning To explore the impact of model quantization on the expressiveness of features, we initially apply PTQ [27] on Φ_{img}^{stu} and examine the discriminative nature of features post-quantization. As demonstrated in Fig. 3, the discriminability of features notably declines after quantization, leading to a misalignment between text and image representations. This reduction in feature clarity compared to the full-precision model underscores the necessity of employing QAT for the final image encoder optimization. QAT enhances the model by incorporating fake quantization during training, which simulates quantization effects through quantization-aware matrix multiplication, followed by finetuning the pre-trained model to adapt to these effects (see supplementary for details of QAT). A pivotal aspect of implementing QAT is selecting

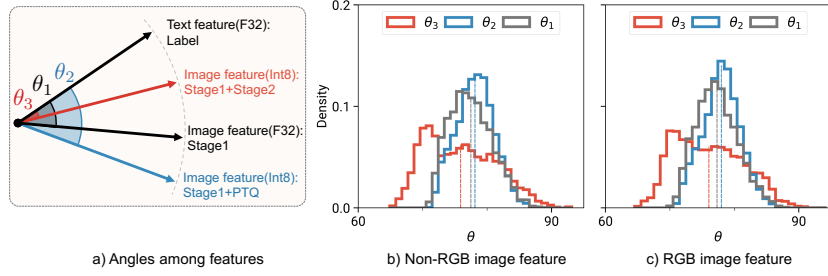


Fig. 3: Angles between the features of images and their corresponding text labels on the ScanNet dataset: We calculate the angles based on the cosine similarities (a lower cosine similarity corresponds to a greater angle between features [18]). A rightward shift in the angle distribution in b) and c) suggests that $\theta_2 > \theta_1$, indicating that image features diverge from the text labels following PTQ. Conversely, a leftward shift implies $\theta_3 < \theta_1$, showing that image features align more closely with the text labels after Stage 2. Dashed lines denote mean values. Best viewed in color.

an appropriate loss function that maintains or even improves the discriminative capability of features within the quantization framework. Traditional knowledge distillation loss, as formulated in Eq. (4), aims to align the student model’s features with those of a pre-trained teacher model. However, it might not exploit the full potential of quantized models in achieving robust and discriminative features. In this context, we propose the integration of contrastive learning loss, which is designed to cultivate representations resilient to non-discriminative features and enhances the separation in the feature space between similar and dissimilar instances. This approach is predicated on the robustness of contrastive learning, which should support the acquisition of invariant representations less affected by the distortions due to quantization. Fig. 3 illustrates how employing contrastive learning loss in conjunction with QAT not only mitigates the discriminative power reduction but potentially increases the discriminability of the student encoder after quantization.

Triplet Sampling Selecting positive and negative samples effectively is pivotal for conservative learning. We adopt the semi-hard sample strategy [4], recognized for its capacity to improve feature robustness, in developing a conservative learning loss. Specifically, for each pair of samples, (x_i, x'_i) , within the training dataset \mathcal{D}_{train} , we generate pseudo labels from the superset \mathcal{S} by utilizing a pre-trained VL model as follows:

$$\hat{y}_i = \arg \max_{Y \in \mathcal{S}} \Phi_{img}(x_i)^\top \Phi_{text}(Y). \quad (5)$$

Then, for each training instance x_i , we identify its corresponding potential positive samples $\{p_{i,k}\}$ and potential negative samples $\{n_{i,j}\}$. Here, potential positives are those samples sharing the same pseudo label as x_i , and potential negatives are those with differing pseudo labels. We select the optimal matching

positive sample p_{i,k^*} by:

$$k^* = \arg \min_k d(\Phi_{img}^{edge}(x_i), \Phi_{img}^{edge}(p_{i,k})), \quad (6)$$

and randomly choose negatives. We subsequently retain only those negative samples such that the semi-hard condition is met:

$$\begin{cases} d(\Phi_{img}^{edge}(x_i), \Phi_{img}^{edge}(n_{i,j})) > d(\Phi_{img}^{edge}(x_i), \Phi_{img}^{edge}(p_{i,k^*})), \\ d(\Phi_{img}^{edge}(x_i), \Phi_{img}^{edge}(n_{i,j})) < d(\Phi_{img}^{edge}(x_i), \Phi_{img}^{edge}(p_{i,k^*})) + m, \end{cases} \quad (7)$$

where m is a predefined constant margin. Denoting the size of the refined negative set as J , we define the loss function:

$$\mathcal{L}_c = \frac{1}{J} \sum_{j=1}^J d(f(x_i), f(p_{i,k^*})) - d(f(x_i), f(n_{i,j})) + m. \quad (8)$$

4 Experimental Results

4.1 Implementation

We use the CLIP model, ViT-g-14 (ViT-G), provided by OpenCLIP [26] as the teacher model. The student model is built upon ViT-S and its *state-of-the-art* (SOTA) variants, DAT-T [49] and Swin-T [35], where we replace the classification head with a feature projection head. For training with student models in stage 1, we use AdamW optimizer with a base learning rate of 10^{-4} and weight decay of 0.05. We use a cosine learning rate scheduler which decays the learning rate to 5×10^{-6} over 120 epochs. In stage 2, we reduce the base learning rate to 10^{-6} . For the threshold τ_c , we empirically set it at 0.25 considering the training data utilization and data noise. For the CLIP text encoder, we use the text prompt "a photo of a {scene category}." or "a satellite image of a {scene category}." depending on the dataset as suggested by [40]. For quantized models, we report static quantization results. For the triplet loss \mathcal{L}_c , we use a margin $m = 0.3$ and a negative set size $J = 3$.

4.2 Overall Results

The overall results assess **EdgeVL** against several SOTA baseline methods on different datasets. We focus on two pivotal metrics: accuracy and efficiency. These metrics are crucial for demonstrating the effectiveness and practical applicability of our method on edge devices.




Datasets The ScanNet and EuroSAT datasets are selected for evaluation. The *ScanNet* dataset [6] includes 1.89M training, 0.53M validation, and 0.21M test indoor RGB-D images. Adhering to [14], we applied a subsampling factor of 100 to the dataset to diminish the presence of similar images, resulting in a

dataset comprising 18,900 training, 5,300 validation, and 2,100 test RGB-D images across 21 scene categories. Since the test split does not provide labels, we evaluate models on the validation split. The *EuroSAT* dataset [21], providing satellite images across 13 spectral bands and encompassing 10 classes with a total of 27,000 images, was randomly split into 13,500 training and 13,500 testing divisions. For the evaluation of cross-modality performance, we utilized the RGB and SWIR (Short Wave Infrared) bands, ensuring a comprehensive assessment of our method’s effectiveness across varied imaging conditions.

Baselines Given the absence of a direct baseline for our novel problem setting, we adapt several methods addressing similar challenges to fit our context. The first adapted baseline, *CMKD* [24], was initially designed for transferring knowledge from LiDAR to RGB models. We modify it for distilling knowledge from the CLIP visual encoder to both RGB and non-RGB models. The *Fida* framework [46], which implements a dual student model approach within a mutual teacher-student learning paradigm, is tailored to our needs by focusing on minimizing the feature distance between student model pairs. Similarly, the *CQD* method [45], originally aimed at knowledge distillation from high-resolution to low-resolution models, is redirected towards reducing the feature distance between non-RGB models and both the pre-trained RGB model and the CLIP visual encoder. The *SKD* strategy [52], sharing conceptual similarities with our work through its mixup technique for generating hybrid-modality samples, is specifically adapted by us to integrate non-RGB and RGB images for training purposes. Lastly, the *Frank* approach [17] and the *Gupta* technique [23] are considered for their relevance in cross-modal weight transfer and fine-tuning, and developing modality-specific models that converge through a unified embedding layer for efficient multimodal data processing, respectively.

To showcase the efficacy of **EdgeVL**, we present the optimal outcomes for baseline models by employing their full-precision (F32) configurations and selecting the backbone that yields the highest accuracy for each. More detailed comparisons with the baselines per backbone can be found in the supplementary. Our comparison also includes two different versions of CLIP with ViT-B/G visual backbones respectively. We term them CLIP-B and CLIP-G hereafter.

Accuracy Tab. 1 displays the accuracy of **EdgeVL** in comparison to baseline models across the ScanNet and EuroSAT datasets, illustrating that **EdgeVL** secures the highest accuracy on both. Notably, the least performing variant of **EdgeVL**, ViT-S, significantly outperforms its closest rival, SKD [52], by an impressive margin of 10.2% on ScanNet (34.5% vs. 44.7%). This gap widens to 13.9% against the top baseline, CQD [45], on EuroSAT (49.4% vs. 64.8%). The consistent outperformance of **EdgeVL** across various backbone architectures for student encoders underscores its adaptability and broad applicability. Last but not least, compared with two pre-trained CLIP models, **EdgeVL** demonstrates superior performance across individual accuracy metrics for both RGB and non-

Table 1: Overall accuracy comparison.  and  denote the top1 accuracy of non-RGB and RGB images, respectively. And  denotes the average of the two, all in percentage. The same applies to the following tables. Best viewed in color.







Methods	Bits	ScanNet (%) \uparrow			EuroSAT (%) \uparrow		
							
Pretrained CLIP-B [40]	F32	4.5	36.2	20.4	16.8	40.4	28.6
Pretrained CLIP-G [40]	F32	6.2	47.3	26.8	16.9	54.0	35.5
Frank [17]	F32	8.3	21.7	15.0	49.2	37.9	43.5
Gupta [23]	F32	16.0	17.5	19.8	54.2	42.4	48.3
CMKD [24] (non-RGB)	F32	37.8	11.5	24.6	61.2	34.4	47.8
CMKD [24] (RGB)	F32	4.0	42.5	23.2	20.1	62.4	41.2
Fida [46]	F32	38.9	5.8	22.3	56.7	20.3	38.5
CQD [45]	F32	40.1	6.7	23.4	62.4	36.4	49.4
SKD [52]	F32	31.2	37.8	34.5	22.9	50.3	36.6
EdgeVL (DAT-T)	Int8	47.9	52.0	49.9	61.0	65.7	63.3
EdgeVL (Swin-T)	Int8	46.0	48.7	47.4	61.3	67.1	64.2
EdgeVL (ViT-S)	Int8	42.0	47.5	44.7	62.9	66.8	64.8

Table 2: Overall efficiency comparison on different GPU platforms

Methods	Bits	Model Size \downarrow	Latency \downarrow		Throughput \uparrow
			AGX	Nano	RTX4090
Pretrained CLIP-G	F32	5213 MB	/	/	/
Pretrained CLIP-B	F32	330 MB	9.5 ms	20.2 ms	772 image/s
EdgeVL (ViT-S)	Int8	86 MB	4.6 ms (\downarrow 52%)	9.9 ms (\downarrow 51%)	1492 image/s (\uparrow 93%)
EdgeVL (Swin-T)	Int8	56 MB	5.2 ms (\downarrow 46%)	11.4 ms (\downarrow 44%)	1098 image/s (\uparrow 42%)

RGB images. This highlights the importance of VL adaptation in the target domain. Fig. 4 further shows the qualitative results.

Efficiency We examine the computational efficiency of **EdgeVL** on the following Nvidia GPUs: i. Jetson AGX Orin (32GB), ii. Jetson Orin Nano (8GB) iii. RTX4090. To simulate actual deployment scenarios, we conduct all model inferences using the TensorRT engine, excluding DAT-T from profiling due to its current incompatibility with TensorRT. As CLIP with ViT-G backbone (CLIP-G) is not compatible with TensorRT, we report its model size only.

4.3 Ablation Study

Quantization Aware Contrastive Learning For comparison, we employ $+PTQ$ [48], which converts model weights and activations into 8-bit integers. The process of $+QAT$ [27] involves the integration of fake quantization layers within the model, followed by finetuning to refine its performance. Similarly,



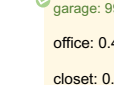










	CLIP-G (F32)	CQD (F32)	SKD (F32)	EdgeVL (Int8)
storage/basement/garage (RGB)	 ✓ storage/basement/garage: 98.9% office: 0.5% apartment: 0.3%	✗ office: 83.8% classroom: 5.0% game room: 3.7%	✗ apartment: 45.2% storage/basement/garage: 31.9% closet: 11.7%	✓ storage/basement/garage: 99.0% office: 0.4% closet: 0.2%
storage/basement/garage (Depth)	 ✗ bathroom: 35.1% hallway: 29.4% lobby: 11.3%	✗ office: 98.5% closet: 0.6% copy/mail room: 0.6%	✗ office: 27.7% closet: 23.2% hallway: 14.6%	✓ storage/basement/garage: 49.8% closet: 38.2% hallway: 8.1%
highway(RGB)	 ✓ Highway: 99.8% industrial: 0.1% permanent crop: 0.0%	✓ Highway: 58.4% Industrial: 35.4% Residential: 4.7%	✓ Highway: 91.3% Industrial: 6.8% Residential: 1.8%	✓ Highway: 96.0% Pasture: 3.2% Residential: 0.4%
highway(IR)	 ✗ Pasture: 31.2% Forest: 24.5% Industrial: 19.0%	✓ Highway: 99.0% Industrial: 0.8% Residential: 0.1%	✗ Industrial: 47.2% Highway: 31.7% PermanentCrop: 13.0%	✓ Highway: 98.4% Pasture: 1.0% PermanentCrop: 0.4%

Fig. 4: Visualization of the predictions of different models on ScanNet and EuroSAT. CLIP-G, CQD [45] and SKD [52] fall short for non-RGB images, while **EdgeVL** (Swin-T) demonstrates superior performance across both image modalities.

Table 3: QAT meets Contrastive Learning. + denotes in combination with stage-1.

Methods	Bits	DAT-T (%)			Swin-T (%)			ViT-S (%)		
										
Stage-1	F32	38.6	40.6	39.6	39.9	41.2	40.5	37.8	40.7	39.3
+PTQ [27]	Int8	33.0	36.5	34.8	29.0	31.7	30.3	24.7	25.9	25.3
+QAT [27]	Int8	39.4	41.2	40.3	38.9	39.7	39.3	37.7	41.1	39.4
+QViT [32]	Int8	35.0	38.0	36.5	36.5	38.5	37.5	31.4	35.3	33.3
+Stage-2	Int8	47.9	52.0	50.0	46.0	48.7	47.4	42.0	47.5	44.7

+QViT [32] enhances model efficiency through a combination of information rectification and distribution-guided distillation during the finetuning stage.

Tab. 3 demonstrates the open-vocabulary classification accuracy of int8 quantized models using different methods for edge device deployment, focusing on static quantization results, though dynamic quantization aligns similarly and is detailed in the supplementary. PTQ significantly lowers accuracy, whereas QAT with stage-1 loss Eq. (4) improves it but doesn’t reach the **EdgeVL** benchmark, sometimes lagging by over 9.7% (40.3% vs. 50.0%). In comparison, stage-2 in **EdgeVL** introduces contrastive learning in the quantization-aware training process, which improves the student model’s discriminability and accuracy despite quantization. The result highlights **EdgeVL**’s effectiveness and robustness in preserving model performance after quantization, outperforming other strategies including QViT and PTQ by significant margins.

Table 4: Effect of Dual-modality Knowledge Distillation







Methods	Bits	ScanNet (%)			EuroSAT (%)		
							
CMKD [24] (non-RGB)	F32	37.8	11.5	24.6	61.2	34.4	47.8
CMKD [24] (RGB)	F32	4.0	42.5	23.2	20.1	62.4	41.2
Stage-1 (Dual-modality)	F32	38.6	40.6	39.6	61.5	60.3	60.9

Table 5: Effect of varying τ_c




τ_c			
0.10	49.8	52.9	51.4
0.25	61.3	67.1	64.2
0.50	61.2	62.8	62.0

Table 6: Effect of different sampling strategies







Sampling			
Semi-hard	61.0	65.7	63.3
Hard	60.8	64.7	62.7

Table 7: Effect of different training strategies

Training			
Two-stage	47.9	52.0	49.9
One-stage	27.1	33.0	30.0







Dual-modality Knowledge Distillation We evaluate the outcomes of our stage-1 training against CMKD(non-RGB) [24] and CMKD(RGB) [24], which resemble our stage-1 but rely on single-modality training. Tab. 4 illustrates the accuracy of these methods using the DAT-T model. Our findings show that dual-modality training draws on the strengths of RGB-modality and non-RGB-modality training, yielding high accuracy for both modalities, whether non-RGB are depth images from ScanNet or infrared images from EuroSAT. Our dual-modality training method achieves an increase of 15.0% on ScanNet and 13.1% on EuroSAT in average accuracy. Notably, the accuracy boosts on non-RGB modality (*e.g.*, 38.6% vs. 37.8% and 61.5% vs. 61.2%) lead us to hypothesize that RGB images function as a type of data augmentation for non-RGB images in dual-modality training.

Cutoff Ratio of τ_c Tab. 5 presents the accuracy of EdgeVL (Swin-T) on EuroSAT, where $\tau_c = 0.25$ brings the highest accuracy compared to the other cutoff ratios. A very small τ_c fails to fully train the model, whereas a very large τ_c introduces noisy samples, thus diminishing the effectiveness of knowledge transfer.

Triplet Sampling Strategy For positives and negatives mined according to Eq. (6), we can choose to build triplets that include: 1) hard triplets, where the anchor-positive distance is greater than the anchor-negative distance or 2) semi-hard triplets, where the anchor-negative distance is greater than the anchor-positive distance but less than the anchor-positive distance plus a margin m . We are assessing how different sampling strategies impact the performance of EdgeVL. Tab. 6 presents the accuracy of EdgeVL (DAT-T) with different sampling strategies on EuroSAT. It demonstrates that semi-hard strategy outperforms hard strategy. This is consistent with the findings in [42].

Two-stage Training EdgeVL follows a two-stage training process. We initially explored a one-stage alternative, replacing QAT with PTQ during model quan-

Table 8: Accuracy on unseen datasets with the image encoder adapted on ScanNet

Methods	Bits	NYU2 (%)			SUNRGBD (%)		
							
Pre-trained CLIP-G	F32	25.7	69.7	47.7	18.0	54.3	36.2
Pre-trained CLIP-B	F32	22.6	62.2	42.4	15.2	47.2	31.2
EdgeVL: DAT-T	Int8	51.1	54.3	52.7	28.6	31.8	30.2
EdgeVL: Swin-T	Int8	43.4	43.3	43.4	30.0	31.4	30.7
EdgeVL: ViT-S	Int8	41.0	40.5	40.8	25.8	28.0	27.0

tization to prevent model collapse during training. Tab. 7 shows the accuracy of EdgeVL (DAT-T) with different training strategies on ScanNet. It can be observed that One-stage training generates much inferior accuracy than Two-stage, *i.e.*, EdgeVL. We hypothesize that this is because contrastive learning benefits from a well-organized feature space as a starting point. Additionally, QAT training necessitates a small learning rate to prevent model collapse, while in our case knowledge distillation training is effective with relatively large learning rates.

4.4 Cross-Dataset Performance

We assessed the generalization performance of EdgeVL by training the models on ScanNet and evaluating their open-vocabulary classification accuracy on the unseen SUNRGBD and NYU2 datasets: i. **SUNRGBD** dataset [44] contains 5,285 training and 5,050 test RGBD images with 19 scene categories. The images are captured by multiple RGBD sensors, including Kinect, Asus Xtion, and Intel RealSense; and ii. **NYU2** dataset [43] contains 795 training and 654 test RGBD images with 10 scene categories, which are captured by a Microsoft Kinect sensor.

As shown in Tab. 8, a trade-off is revealed: while EdgeVL-enhanced encoders significantly improved depth image accuracy (up to 25.4% with DAT-T), there was a slight decrease in RGB image accuracy compared to the pre-trained CLIP model’s ViT-G encoder. This was expected due to the substantial reduction in model size post-quantization (*e.g.*, Swin-T’s 56MB vs. CLIP-G’s 5,213MB). Besides, CLIP’s training involved 400 million image-text pairs, while EdgeVL used only 4,725 RGB and depth image pairs from ScanNet, suggesting that the limited adaptation dataset size may affect the model’s generalization ability.

5 Conclusion

EdgeVL showcases a significant advancement in leveraging a pre-trained visual-linguistic model for open-vocabulary classification across diverse image modalities, including both RGB and non-RGB domains. Despite its innovative approach, EdgeVL encounters a challenge in preserving generalization performance for RGB images when adapted for cross-modal use. Future work will concentrate on refining adaptation techniques to overcome this limitation, aiming to enhance the framework’s versatility and effectiveness in a broader range of applications.

References

1. Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., others: Gpt-4 technical report. arXiv (2023)
2. Cai, Y., Yao, Z., Dong, Z., Gholami, A., Mahoney, M.W., Keutzer, K.: Zeroq: A Novel Zero Shot Quantization Framework. In: Computer Vision and Pattern Recognition (CVPR). pp. 13166–13175. IEEE (2020)
3. Chen, R., Liu, Y., Kong, L., ZHU, X., Ma, Y., LI, Y., Hou, Y., Qiao, Y., Wang, W.: Clip2scene: Towards Label-efficient 3d Scene Understanding by CLIP. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). arXiv (2023)
4. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
5. Choi, J., Wang, Z., Venkataramani, S., Chuang, P.I.J., Srinivasan, V., Gopalakrishnan, K.: Pact: Parameterized clipping activation for quantized neural networks. openreview.net **abs/1805.06085** (2018)
6. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scannet: Richly-annotated 3d Reconstructions of Indoor Scenes. arXiv (2017)
7. Dai, X., Jiang, Z., Wu, Z., Bao, Y., Wang, Z., Liu, S., Zhou, E.: General Instance Distillation for Object Detection. In: Computer Vision and Pattern Recognition (CVPR). pp. 7842–7851 (2021)
8. Dettmers, T., Lewis, M., Zettlemoyer, L.: Gpt3.int8(): 8-bit Matrix Multiplication for Transformers at Scale. In: Conference on Neural Information Processing Systems (NeurIPS) (2022)
9. Ding, R., Yang, J., Xue, C., Zhang, W., Bai, S., QI, X.: Pla: Language-Driven Open-Vocabulary 3d Scene Understanding. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
10. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An Image is Worth 16x 16 Words :transformers for Image Recognition at Scale. In: International Conference on Learning Representations. vol. abs/2010.11929 (2020)
11. Du, Y., Wei, F., Zhang, Z., Shi, M., Gao, Y., Li, G.: Learning To Prompt for Open-Vocabulary Object Detection With Vision-Language Model. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 14084–14093 (2022)
12. Esser, S.K., McKinstry, J.L., Bablani, D., Appuswamy, R., Modha, D.S.: [LSQ] Learned Step Size Quantization (2020)
13. Fang, R., Pang, G., Bai, X.: Simple Image-level Classification Improves Open-vocabulary Object Detection. arXiv **abs/2312.10439** (2023)
14. Fischedick, S.B., Seichter, D., Schmidt, R., Rabes, L., Gross, H.M.: Efficient Multi-Task Scene Analysis with RGB-D Transformers. In: IEEE International Joint Conference on Neural Network (IJCNN). pp. 1–10 (2023)
15. Ghiasi, G., Gu, X., Cui, Y., Lin, T.Y.: Scaling Open-vocabulary Image Segmentation with Image-level Labels. In: European Conference on Computer Vision (ECCV) (2022)
16. Gu, X., Lin, T.Y., Kuo, W., Cui, Y.: Open-vocabulary Object Detection via Vision and Language Knowledge Distillation. In: International Conference on Learning Representations (ICLR) (2022)

17. Hafner, F.M., Bhuyian, A., Kooij, J.F., Granger, E.: Cross-modal distillation for RGB-depth person re-identification. *Computer Vision and Image Understanding* **216**, 103352 (2022)
18. Han, J., Pei, J., Tong, H.: *Data mining: concepts and techniques*. Morgan kaufmann (2022)
19. He, S., Guo, T., Dai, T., Qiao, R., Ren, B., Xia, S.: Open-Vocabulary Multi-Label Classification via Multi-Modal Knowledge Transfer. *Proceedings of the AAAI Conference on Artificial Intelligence* **abs/2207.01887**(1), 808–816 (2022)
20. He, W., Jamonnak, S., Gou, L., Ren, L.: Clip-SMATHENVVDOLLARi: Language-Guided Self-Supervised Semantic Segmentation. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2023)
21. Helber, P., Bischke, B., Dengel, A., Borth, D.: Eurosat: A Novel Dataset and Deep Learning Benchmark for Land Use and Land Cover Classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **12**(7), 2217–2226 (2018)
22. Hinton, G., Vinyals, O., Dean, J.: Distilling the Knowledge in a Neural Network. *arXiv* (2015)
23. Hoffman, J., Gupta, S., Leong, J., Guadarrama, S., Darrell, T.: Cross-modal adaptation for RGB-D detection. In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 5032–5039. IEEE (2016)
24. Hong, Y., Dai, H., Ding, Y.: Cross-Modality Knowledge Distillation Network for Monocular 3d Object Detection. In: *European Conference on Computer Vision (ECCV)* (2022)
25. Huynh, D., Kuen, J., Lin, Z., Gu, J., Elhamifar, E.: Open-Vocabulary Instance Segmentation via Robust Cross-Modal Pseudo-Labeling. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 7020–7031 (2022)
26. Ilharco, G., Wortsman, M., Wightman, R., Gordon, C., Carlini, N., Taori, R., Dave, A., Shankar, V., Namkoong, H., Miller, J., Hajishirzi, H., Farhadi, A., Schmidt, L.: Openclip (Jul 2021). <https://doi.org/10.5281/zenodo.5143773>, <https://doi.org/10.5281/zenodo.5143773>
27. Jacob, B., Kligys, S., Chen, B., Zhu, M., Tang, M., Howard, A., Adam, H., Kalenichenko, D.: Quantization and Training of Neural Networks for Efficient Integer-Arithmetic-Only Inference. In: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2704–2713. IEEE (2018)
28. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q.V., Sung, Y.H., Li, Z., Duerig, T.: Scaling Up Visual and Vision-Language Representation Learning With Noisy Text Supervision. In: *International Conference on Machine Learning (ICML)*. pp. 4904–4916 (2021)
29. Krishnamoorthi, R.: Quantizing deep convolutional networks for efficient inference: A whitepaper. *arXiv* **abs/1806.08342** (2018)
30. Kuo, W., Cui, Y., Gu, X., Piergiovanni, A., Angelova, A.: Open-Vocabulary Object Detection upon Frozen Vision and Language Models. *ICLR 2023 poster* **abs/2209.15639** (2023)
31. Lee, J., Kim, D., Ham, B.: Network Quantization with Element-wise Gradient Scaling. In: *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 6444–6453. IEEE (2021)
32. Li, Y., Xu, S., Zhang, B., Cao, X., Gao, P., Guo, G.: Q-ViT: Accurate and Fully Quantized Low-bit Vision Transformer. In: *Conference on Neural Information Processing Systems (NeurIPS)*. vol. 35, pp. 34451–34463 (2022)

33. Li, Y., Adamczewski, K., Li, W., Gu, S., Timofte, R., Gool, L.: Revisiting Random Channel Pruning for Neural Network Compression. In: Computer Vision and Pattern Recognition (CVPR). pp. 191–201. IEEE (2022)
34. Liu, J., Niu, L., Yuan, Z., Yang, D., Wang, X., Liu, W.: Pd-Quant: Post-Training Quantization Based on Prediction Difference Metric. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
35. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. arXiv (2021)
36. Minderer, M., Gritsenko, A., Stone, A., Neumann, M., Weissenborn, D., Dosovitskiy, A., Mahendran, A., Arnab, A., Dehghani, M., Shen, Z., Wang, X., Zhai, X., Kipf, T., Houlsby, N.: Simple Open-Vocabulary Object Detection (2022)
37. Nagel, M., Baalen, M.v., Blankevoort, T., Welling, M.: Data-Free Quantization Through Weight Equalization and Bias Correction. In: IEEE International Conference on Computer Vision (ICCV). pp. 1325–1334 (2019)
38. Peng, S., Genova, K., Jiang, C., Tagliasacchi, A., Pollefeys, M., Funkhouser, T.: Openscene: 3d Scene Understanding with Open Vocabularies. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
39. Qin, J., Wu, J., Yan, P., Li, M., Ren, Y., Xiao, X., Wang, Y., Wang, R., Wen, S., Pan, X., Wang, X.: Freeseg: Unified, Universal and Open-Vocabulary Image Segmentation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
40. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: [CLIP] Learning Transferable Visual Models From Natural Language Supervision. In: International Conference on Machine Learning (ICML). pp. 8748–8763 (2021)
41. Romero, A., Ballas, N., Kahou, S.E., Chassang, A., Gatta, C., Bengio, Y.: Fitnets: Hints for Thin Deep Nets. In: International Conference on Learning Representations (ICLR) (2015)
42. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)
43. Silberman, N., Hoiem, D., Kohli, P., Fergus, R.: Indoor Segmentation and Support Inference from RGBD Images. In: Computer Vision – ECCV 2012. pp. 746–760. Springer (2012)
44. Song, S., Lichtenberg, S.P., Xiao, J.: SUN RGB-D: A RGB-D Scene Understanding Benchmark Suite. pp. 567–576 (2015)
45. Su, J.C., Maji, S.: Adapting Models to Signal Degradation using Distillation. In: British Machine Vision Conference (BMVC) (2017)
46. Thoker, F.M., Gall, J.: Cross-Modal Knowledge Distillation for Action Recognition. In: 2019 IEEE International Conference on Image Processing (ICIP). pp. 6–10. IEEE (2019)
47. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jegou, H.: Training data-efficient image transformers & distillation through attention. In: Proceedings of the 38th International Conference on Machine Learning. pp. 10347–10357. PMLR (2021)
48. Wu, H., Judd, P., Zhang, X., Isaev, M., Micikevicius, P.: Integer Quantization for Deep Learning Inference: Principles and Empirical Evaluation. arXiv **abs/2004.09602** (2020)
49. Xia, Z., Pan, X., Song, S., Li, L.E., Huang, G.: Vision Transformer With Deformable Attention. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4794–4803 (2022)

50. Xiao, G., Lin, J., Seznec, M., Wu, H., Demouth, J., Han, S.: Smoothquant: Accurate and Efficient Post-Training Quantization for Large Language Models. In: International Conference on Machine Learning (ICML). pp. 38087–38099 (2023)
51. Xue, L., Gao, M., Xing, C., Martín-Martín, R., Wu, J., Xiong, C., Xu, R., Niebles, J.C., Savarese, S.: Ulip: Learning a Unified Representation of Language, Images, and Point Clouds for 3d Understanding. In: Computer Vision and Pattern Recognition (CVPR). pp. 1179–1189 (2023)
52. Yang, C., An, Z., Zhou, H., cai, l., Zhi, X., Wu, J., xu, y., Zhang, Q.: Mixskd: Self-Knowledge Distillation from Mixup for Image Recognition. In: European Conference on Computer Vision (ECCV) (2022)
53. Zareian, A., Rosa, K.D., Hu, D.H., Chang, S.F.: Open-Vocabulary Object Detection Using Captions. In: Computer Vision and Pattern Recognition (CVPR). pp. 14393–14402 (2021)
54. Zhang, R., Guo, Z., Zhang, W., Li, K., Miao, X., Cui, B., Qiao, Y., Gao, P., Li, H.: Pointclip: Point Cloud Understanding by CLIP. pp. 8552–8562 (2022)