

OmniVLA: Physically-Grounded Multimodal VLA with Unified Multi-Sensor Perception for Robotic Manipulation

Heyu Guo^{1*}, Shanmu Wang^{2*}, Ruichun Ma^{3†}, Shiqi Jiang³
Yasaman Ghasempour¹, Omid Abari², Baining Guo³, Lili Qiu³

¹Princeton University ²University of California, Los Angeles ³Microsoft Research Asia

Abstract—Vision-language-action (VLA) models have shown strong generalization in robotic manipulation through large-scale vision-language pretraining. However, most existing models rely solely on RGB cameras, limiting their perception and, consequently, manipulation capabilities. We present OmniVLA, an omni-modality VLA model that integrates novel sensing modalities to enable beyond-RGB robotic perception and manipulation. The core of our approach is the *sensor-masked image*, a unified representation that overlays physically meaningful, spatially grounded masks onto the RGB images. These masks are derived from sensors including an infrared camera, a mmWave radar, and a microphone array. This image-native unification keeps sensor input close to RGB statistics to facilitate training, provides a uniform interface across sensor hardware, and enables data-efficient learning with lightweight per-sensor projectors. Building on this, we design a multimodal vision-language-action model architecture and train OmniVLA by extending an RGB-pretrained VLA backbone. We evaluate OmniVLA on challenging real-world tasks that require sensor-modality perception to guide the manipulation. OmniVLA achieves an average task success rate of 84%, significantly outperforms both RGB-only and raw-sensor-input baseline models by 59% and 28% respectively, meanwhile showing higher learning efficiency and stronger generalization capability.


I. INTRODUCTION

Vision-language-action (VLA) models [1], [2] recently emerged as a powerful paradigm for generalist robotic policies. They leverage vision-language pretraining to map user prompts and camera observations to robot actions, showing great generalization and instruction following capability. However, most VLA models are limited to RGB camera image input, which constrains their perception capabilities and prevents them from handling tasks that require non-RGB cues. This undermines the potential of robots to utilize additional sensory hardware and perform challenging tasks that require perception capability similar to or even beyond humans. For example, Infrared (IR) cameras reveal temperature contrast for tasks such as search and rescue; millimeter wave (mmWave) radars penetrate occlusions, such as cardboard or clothing, with radio-frequency (RF) signals to localize hidden objects [3], [4]; acoustic microphone arrays enable human-like environmental awareness to ambient sound and react accordingly. Adding per-sensor encoders or directly feeding raw sensor data to VLA models, however, increases


*Co-primary authors; work completed during internships at Microsoft Research Asia.

†Corresponding author (ruichunma@microsoft.com).

Task: "Give me the *cold drink*."



Task: "Find my *item inside one of the closed boxes*."



Task: "Find my *ringing phone buried under things*."




Fig. 1: Instead of relying solely on RGB cameras, OmniVLA equips robots with multi-sensor perception. We use beam-forming heatmaps as acoustic and mmWave sensor images to highlight sound source and hidden item respectively.

system complexity, exacerbates data scarcity, and reduces compatibility with pretrained vision-language models.

In this paper, we aim to equip VLA models with unified multi-sensor perception, utilizing sensor modalities including infrared, acoustic, and mmWave. This enables robots to combine the strong generalization of foundation models and physical information from various sensors seamlessly, to enable physically-grounded spatial intelligence.

There are several challenges in integrating diverse sensors with a VLA model. First, VLA models need to effectively interpret heterogeneous sensor information and use them to guide the action; naively feeding raw sensor streams leads to poor performance and data efficiency as shown in Section IV, because existing VLA backbones are trained primarily on RGB images. Second, sensors differ in format, field of view, and resolution, calling for a scalable, uniform representation rather than training sensor fusion models that depend on specific hardware. Lastly, sensor modality data are much more scarce compared to web-scale image-text data pairs,

so we require a data-efficient approach for training. To solve these challenges, we take inspiration from how the human brain interprets sensor information: as we are used to RGB images, we naturally anchor other sensor cues to the RGB view—for example, interpreting infrared camera images by associating temperatures with objects visible in the image.

We present **OmniVLA**, the first multisensory VLA that integrates novel sensing modalities to enable beyond-RGB robotic perception and manipulation by unifying heterogeneous sensors into an image-native space. The core of our design is an intermediate representation, *sensor-masked images*, which is produced by semantically segmenting the RGB image and overlaying the relevant sensor information as colored masks. Such representation makes sure sensor information is spatially grounded and semantically aligned with the RGB image to ease the integration with VLA models. This brings several benefits that solve the challenges above: (i) making sensor information spatially grounded in RGB pixel coordinates to facilitate robotic manipulation on target objects, (ii) remaining close to RGB statistics so existing vision encoders and VLA backbones can be reused for further training, (iii) providing a uniform representation across sensors, resolutions, and hardware variants.

Building on the sensor-masked images, we propose a tailored VLA model architecture (Figure 2). We first convert all raw sensor measurements to image-like 2D spatial representations, specifically, we perform beamforming for mmWave radar and acoustic array data to acquire heatmap-like raw sensor images. Then, to generate masks for interested objects in the scene, we use a cloud-based Vision-Language Model (VLM) to interpret task request and generate a prompt for grounded SAM [5] to provide semantic-based segmentation masks. We further overlay sensor information on the RGB images at masked regions to acquire the sensor-masked images as the input to a frozen vision encoder. After the vision encoder, we add lightweight projection layers for each sensor to generate better aligned tokens for sensor images. Finally, a LLM backbone processes the input tokens and conditions the diffusion-based action expert for the robot action output.

We build a multi-sensor robot arm prototype to collect RGB camera and sensor data, paired with action demonstrations, and train OmniVLA with the collected datasets. We evaluate OmniVLA extensively with several manipulation tasks that require sensor-modality guidance, including thermal-based pick-and-place, mmWave-based see-through boxes and opening the non-empty one, and acoustic-based uncovering ringing phone beneath clothes. OmniVLA achieves an average task success rate of 84%, significantly outperforms 25% success rate of RGB-only baseline and 56% success rate of raw-sensor-input baseline. This shows the benefit from our unique sensing capabilities, and the performance gain from our sensor-masked image representation. We also highlight the data efficiency of our approach by achieving similar success rate with only 50% training data compared to raw-sensor-input baseline. Moreover, we show that our approach provides strong generalization capability

across three unseen tasks, outperforming base VLA model and raw-sensor-data based model by 59% and 28% respectively on average success rate.

To summarize, we make the following contributions:

- 1) To our knowledge, OmniVLA is the first VLA model that unifies multiple sensing modalities, including infrared, mmWave, acoustic, to enable robotic manipulation tasks beyond RGB-based perception capability.
- 2) We introduce *sensor-masked images*, a spatially grounded and semantically aligned representation that allows reusing pre-trained vision encoders, provides a uniform representation across sensor hardware, and improves learning efficiency.
- 3) We present a lightweight OmniVLA model architecture and evaluate system performance with extensive experimental evaluation. We open-source the project at <https://github.com/GuoHeyu/OmniVLA>.

II. RELATED WORK

A. Vision-Language-Action Models

Vision-language-action (VLA) models have been a popular research paradigm for robotic manipulation, using the language prompt and video feed as input and generating robot actions in an end-to-end manner. Conventional robotic manipulation policies using reinforcement learning and simulators provide great performance on specific tasks [6]–[11], while VLA models have shown great improvement in few-shot task generalization and instruction following by leveraging web-scale pretraining. The majority of VLA models only take video from RGB cameras as visual input [1], [2], [12]–[15]. However, these works are inherently limited by the RGB camera input and unable to finish tasks that require perception capabilities beyond RGB.

To address this issue, researchers propose novel Vision-Language-Action models with additional sensor input. Depth information is widely employed to enhance the capability of VLA models with better spatial-temporal understanding [16]–[20]. Other works incorporate tactile perception in VLA models for better task planning, grounding and reasoning capabilities [21], [22]. Moreover, Vlas [23] integrates speech information for convenient and personalized human-robot interaction. MultiPLY [24] introduces an embodied LLM for planning multisensory interactions with the environment, but it is based on simulator only and cannot generate actions for manipulation tasks.

Prior works have not explored the integration of novel sensing modalities like thermal, acoustic, and mmWave, which provide unique sensing information for manipulation tasks. Also, existing approaches often require training sensor encoders for each modality, which needs a large amount of data. In contrast, we propose a data-efficient approach towards beyond-RGB perception. Moreover, previous works propose complex and specific model architecture for extra sensor and image fusion, which is not generalizable to diverse sensors. In contrast, we propose sensor-masked images, a simple and unified sensor fusion strategy for VLA models.

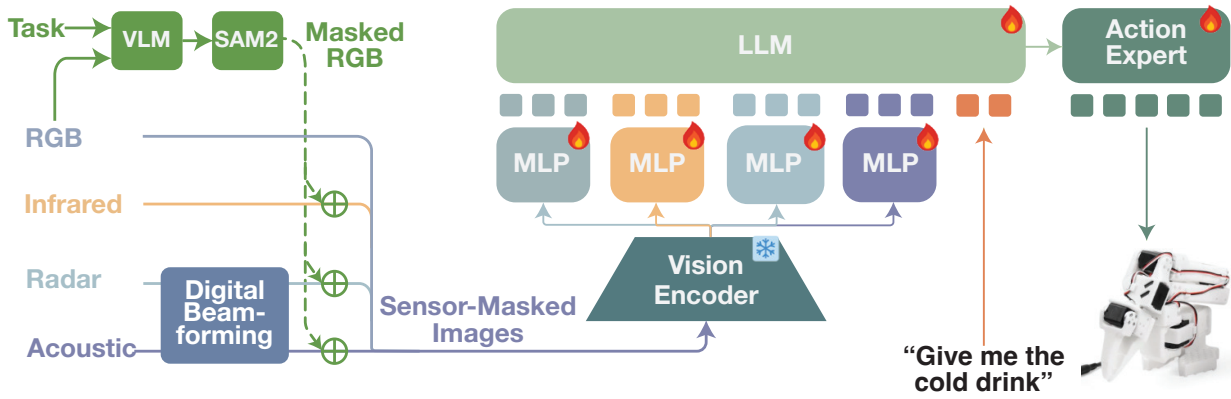


Fig. 2: System Overview. OmniVLA processes diverse sensor data into image-like 2D spatial representations, and then overlays sensor information onto RGB images to produce spatially grounded and semantically aligned *sensor-masked images*. We train OmniVLA by extending an RGB-pretrained VLA backbone to perform challenging tasks beyond RGB perception.

B. Multi-sensor Fusion based Perception

Multi-sensor fusion based perception has received significant interest recently, especially for 3D detection and other downstream tasks for autonomous driving. The benefits for multi-sensor fusion are comprehensive environmental understanding for complementary sensors [25]. One typical type of multi-sensor fusion is mmWave radar and camera sensor fusion, widely used in autonomous driving because of its reliability in rainy and foggy environments [3], [26]–[30]. Another thread of recent work fuses RF signals for robotic perception to achieve non-line-of-sight object finding [31]–[33], but it requires extra RFID tags attached on the objects and are designed only for RF signals.

Prior works focus on environmental perception tasks like 3D object detection, semantic scene understanding tasks [34], [35], while we aim for robot manipulation by generating actions. These prior works also heavily rely on model architectures tailored for domain-specific downstream tasks with specific sensors, unable to provide strong generalization or instruction following of VLA models. Moreover, OmniVLA provides a general framework for more diverse sensors.

III. SYSTEM DESIGN

A. System Overview

OmniVLA contains two parts: sensor-masked image generation and multi-sensor vision-language-action model architecture. The first part first extracts raw sensor data, then preprocesses it to 2D sensor images/heatmaps. After that, we use a semantic segmentation model to generate the masks from the RGB images following a VLM-generated prompt. Then we overlay the sensor images on the masked regions of RGB images to output sensor-masked images, which are the input for our multi-sensor vision-language-action model. The second part is a multi-sensor vision-language-action model backbone that is designed for sensor-masked images to capture each sensor’s input and avoid requiring a large amount of sensor data for training. We utilize the existing frozen vision encoders to encode sensor-masked images. For each sensor modality, we use individual multi-layer

perceptron layers to align sensor image tokens with language and RGB image tokens. The tokens are concatenated together with language tokens as input for the large language model in the architecture, and then we generate the final action predictions using the action expert.

B. Sensor Data Representation

The first part of our work is sensor-masked image generation, which includes sensor data preprocessing, segmentation mask processing, and sensor-masked image blending.

Preprocessing. We convert all raw measurements into a *camera-like 2D spatial representation*. Thermal camera already outputs raster images—infrared intensity, defined over image coordinates (u, v) . In contrast, the mmWave radar and the microphone array provide complex signal samples per array element, $x_{i,k} = A_{i,k}e^{j\psi_{i,k}}$ for the k -th element of array i , where $i \in \{\text{mmWave, acoustic}\}$. To obtain inputs with *consistent spatial mapping* (horizontal and vertical viewing axes on a 2D grid), we compute azimuth–elevation heatmaps via conventional delay-and-sum beamforming [36]:

$$I_i(\theta, \phi) = 20 \log_{10} \left\| \sum_{k=1}^K A_{i,k} e^{j\psi_{i,k}} e^{-j\Phi_{i,k}} \right\|, \quad (1)$$

$$\Phi_{i,k} = \frac{2\pi}{\lambda_i} (x_{i,k} \cos \phi \sin \theta + y_{i,k} \sin \phi).$$

where θ is the azimuth angle, ϕ is the elevation angle, λ_i is the wavelength received by sensor i , $(x_{i,k}, y_{i,k})$ is the position for k^{th} array element (antenna or microphone) in array i . One example is shown in Fig. 1. Similar to the principle of RGB camera, the azimuth-elevation heatmap reveals the environmental information in a direct way for human understanding and acts as the sensor images in the following steps.

Segmentation and overlay. The next step is generating segmentation masks based on the task description and RGB images. We first send the task request text and RGB image to a vision-language model, GPT-4o [37], to generate the segmentation prompts that describe objects related to the task in the current scene, such as ‘red block/drink’, ‘black phone’,

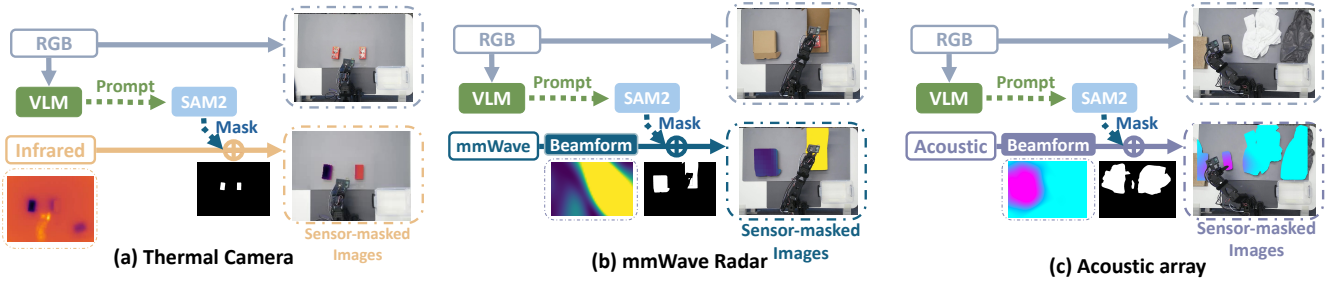


Fig. 3: Sensor Data Processing Illustration. We propose a general sensor data processing pipeline applicable to various sensors, including (a) thermal camera, (b) mmWave radar, and (c) acoustic microphone array, by overlaying sensor information on top of RGB images as VLA model input. We update the prompt input to the SAM2 model when the task begins and then asynchronously check for updates, so that VLM delay does not affect the real-time processing of sensor-masked images.

and ‘cardboard boxes’. These prompts are initialized when the task begins and are updated periodically to accommodate scene dynamics, such as the introduction of new objects. We run these updates in the background at a lower frequency than the robot’s primary action-generation loop, so that VLM output delay does not affect real-time robot manipulation actions. Then we input the segmentation prompts and RGB image into a segmentation model to generate image masks for task-related objects. Specifically, we use Grounded SAM 2 [5], which combines segment anything model [38] and open-set object detection model, grounding DINO [39]. The final output is a 0-1 matrix, where ‘1’ labels the masked regions for objects. The second step can be expressed as:

$$\begin{aligned} l &= \text{VLM}(T_{\text{task}}, I_{\text{RGB}}), \\ \text{mask} &= \text{SAM2}(l, I_{\text{RGB}}), \end{aligned} \quad (2)$$

where l represents the segmentation prompts generated by VLM, T_{task} is the input task description, and I_{RGB} is the RGB image. To generate the final sensor-masked images, we first spatially align each sensor modality with the RGB frame through a one-time calibration involving rotation and cropping. While high-precision pixel matching is not strictly required for the model to learn effectively, this ensures the sensor heatmaps are roughly centered on their physical counterparts. We then isolate the regions of interest by applying the masks generated in the previous step. This fusion process is formalized as:

$$\begin{aligned} I_i^c &= \text{Calibration}(I_i), \\ I_i^m &= \text{mask} \odot (\alpha I_i^c + (1 - \alpha) I_{\text{RGB}}) \\ &\quad + (1 - \text{mask}) \odot I_{\text{RGB}} \end{aligned} \quad (3)$$

where $i \in \{\text{mmWave}, \text{acoustic}, \text{thermal}\}$ and $\alpha \in [0, 1]$ is the hyper-parameter for image blending. The higher α is, the more sensor information remains for comprehensive understanding, but less correlation between the masked part and the unmasked RGB image remains. We set α as 1 by default based on our empirical testing. These consolidated sensor images then serve as the unified input for our multi-sensor VLA backbone.

C. Model architecture and training

The second part is the architecture design and training strategy of our multi-sensory vision-language-action model.

Existing VLA models typically encode RGB inputs with a vision encoder, followed by a multi-layer perceptron (MLP) to generate image tokens. Simultaneously, the task description is processed by a language tokenizer. These tokens are concatenated and fed into a large language model (LLM) backbone, which generates action tokens for robot manipulation. We reuse and adapt this architecture design.

We present a generalizable and efficient multi-sensory VLA model architecture design as shown in Fig. 2. We enable effective sensor data understanding by feeding sensor-masked images to the existing vision encoder. Then for each sensor modality, we input the encoded results into individual multi-layer perceptron (MLP) modules for projection to align with the language tokens. The projected tokens for each sensor are concatenated together with language tokens from the task description after passing through a language tokenizer. The LLM backbone takes these concatenated tokens as input. To generate the final robotic actions, we follow the architecture used in SmoVLA [2] and π_0 [1], where the LLM backbone is combined with a diffusion or flow matching based Action Expert for robotic action generation. We utilize conditional flow matching or diffusion within the action expert to model the continuous distribution of actions, which is inherited from the chosen base model. This allows the model to generate a full action chunk step by step from random noise. We use SmoVLA as our base model by default, modify its architecture, and train our model from its pretrained weights. Our approach is potentially compatible with various VLA backbones because our sensor-masked image representation allows processing sensor data as standard visual tokens.

The overall data flow can be described with:

$$\begin{aligned} t_i &= \text{MLP}_i(E_I(I_i^m)), \\ t_{\text{task}} &= E_L(T_{\text{task}}), \\ \text{action} &= \text{VLA}([t_1, t_2, \dots, t_m, t_{\text{task}}]) \end{aligned} \quad (4)$$

where m is the number of sensors we use, E_I is the image encoder, and E_L is the language embedding layer, t_i is

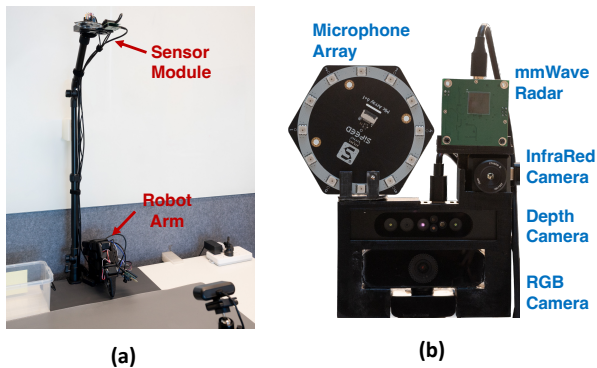


Fig. 4: Hardware Implementation. (a) robot arm and sensor setup (b) sensor module, integrating multiple sensors and cameras.

the embedding of i th sensor image. We underline that our architecture does not require all sensors shown here. Instead, we allow flexible sensor setup according to the deployment scenario, budget, etc, for example using a single sensor like an infrared camera. This design is a general and flexible framework to process beyond-RGB sensor-masked images, compatible with existing Vision-Language-Action models. It means we can utilize the vision-language capability in existing model’s pretraining and understand unique features from each sensor easily with a few demonstration data.

During the model training, we start from the pretrained base model weight. We freeze the vision encoder, typical in various Vision-Language-Action models [1], [2], [40] and Vision-Language models finetuning [41], [42]. Meanwhile, we set all other weights trainable for multi-sensor perception in robotic manipulation. We initialize the individual MLP projectors for each sensor using weights from the pre-trained RGB projection layer of the base VLA model. This strategy leverages established visual feature mappings as a strong prior, allowing the model to adapt quickly to the unique sensor-masked images. Finally, we co-fine-tune these sensor-specific MLP modules alongside the unfrozen backbone weights using the collected demonstration datasets.

IV. EVALUATION

We evaluate OmniVLA with a real-world prototype across several sensor-related manipulation tasks. First, we show our unique beyond-RGB perception capabilities, using daily tasks that require non-visible cues, including thermal-based pick-and-place, see-through and open the box containing the object, and uncovering hidden ringing items. OmniVLA significantly outperforms RGB-only VLA models and VLA models trained with unprocessed sensor images. We also highlight the data efficiency of our approach. Second, we show that our approach provides superior generalization capability for sensor-related tasks, outperforming baselines.

A. Experimental Setup

Implementation. Our prototype includes a SO101 manipulator arm [43] with a standard top-down RGB camera, a front camera, and an arm camera, augmented by our multi-modal




Fig. 5: Examples of Robotic Manipulation Task Completion over Time. (a) Thermal: finding the cold drink. (b) mmWave: opening the box with an object inside. (c) Acoustic: uncovering the ringing phone. The first three images are sensor-masked images; the rest of the images are raw RGB images for visibility.

sensor suite comprising: a depth camera, an infrared thermal camera, an mmWave radar sensor, and a six-microphone circular array (Figure 4). While the RGB camera provides standard visual perception, our additional sensor modalities capture unique physical information that extends perception capability significantly.

Model training and inference. We use SmolVLA [2] as the base model by default to implement our design and use the pre-trained weights. We expect our approach to be compatible with most existing RGB-only VLA models. We use 8 NVIDIA A100 GPUs on a server for distributed training and use a local RTX 4090 GPU for model inference during system evaluation. Depending on the size of the demonstration dataset, the time required for training is around 14 hours for 50 K optimization steps with batch size 32. For real-time inference, we load both the VLA model and segmentation model on the local RTX 4090 machine, which is able to output 15 predictions per second for end-to-end action prediction. Note that prompting the VLM model for segmentation keywords is invoked at the beginning of the task and only updated infrequently to minimize overhead. We expect to have lower delays with code optimization.

Task setting. We evaluate OmniVLA on three types of manipulation tasks that need to leverage non-visual sensory modalities: (1) *Thermal modality*: Distinguishing between a cold and warm drink, picking up the cold drink and placing it into a plastic container; (2) *mmWave modality*: Seeing through enclosed cardboard/foam (non-metal) boxes with mmWave radar, opening one of the boxes with an object inside and exposing the target object. (3) *Acoustic modality*: Locating a ringing mobile phone concealed beneath opaque coverings using spatial audio cues from the microphone array, and removing the covering to uncover the phone. We show examples of successful action trajectories from OmniVLA in Figure 5.

Evaluation metrics and baselines. We evaluate model performance using task success rates computed over 25 independent trials per task with random object placement, complemented by task scores: 0.5 score for choosing the

	Success Rate				Task Score			
	Thermal	mmWave	Acoustic	Average	Thermal	mmWave	Acoustic	Average
VLA-RGB	28%	8%	40%	25%	0.62	0.34	0.70	0.55
VLA-RAW	52%	68%	48%	56%	0.74	0.84	0.62	0.73
OmniVLA	80%	84%	88%	84%	0.91	0.88	0.92	0.90

TABLE I: OmniVLA success rates and task scores across three types of sensor-modality tasks.

Base Model	Success Rate				Task Score			
	Thermal	mmWave	Acoustic	Average	Thermal	mmWave	Acoustic	Average
SmolVLA	80%	84%	88%	84%	0.91	0.88	0.92	0.90
Pi0	68%	60%	64%	64%	0.84	0.72	0.82	0.80

TABLE II: OmniVLA performance when using two different base models as backbone, evaluated across three types of sensor tasks.

	Thermal		mmWave		Acoustic	
	Stage 1	Stage 2	Stage 1	Stage 2	Stage 1	Stage 2
OmniVLA-Base	100%	24%	56%	40%	76%	16%
Pretrained VLA-RAW	76%	84%	52%	76%	60%	92%
Pretrained OmniVLA	100%	92%	92%	80%	92%	92%
Gains	+0%/+24%	+68%/+8%	+36%/+40%	+40%/+4%	+16%/+32%	+76%/+0%

TABLE III: Pretraining effectiveness by comparing per-stage success rates (%). Gains are improvements over baselines.

right item to interact with, 0.5 score for performing the correct manipulation, e.g., picking up and placing in the container, opening up the box, and removing the coverings. For baselines, we compare our approach against the following ablation baselines: (1) **VLA-RGB** (modality ablation): VLA models with standard RGB input only for training and inference, without our architecture changes. Since we use SmolVLA as the backbone, this baseline model is essentially a SmolVLA model fine-tuned on our task dataset. (2) **VLA-RAW** (representation ablation): VLA models with raw sensor data/images for training and inference input. It uses the same model architecture as OmniVLA, but skips the segmentation and overlay step for the sensor data processing. Note that we still apply beamforming for the VLA-RAW model mmWave and acoustic sensor input to have a strong baseline.

B. Multi-sensory Task Performance

We first evaluate OmniVLA performance on manipulation tasks compared with the baselines. Then we evaluate the performance over different foundational robotic models. Finally, we explore the impact of finetuning data amount on final performance.

Overall Performance. We evaluate OmniVLA after training on 100 expert demonstration episodes of thermal and acoustic modality tasks and 200 episodes of mmWave modality task individually. The mmWave task requires more training data because opening a box is more difficult to learn compared to typical pick-and-place type of actions, which seldom appears in VLA model pretraining datasets. For each demonstration, we randomize objects’ positions on a table within the robot’s workspace. As the SmolVLA pretraining dataset does not include any non-RGB sensor,

we consider the number of episodes reasonable and showing high learning efficiency of our approach. Table I shows success rates and task scores across tasks. OmniVLA finishes the tasks requiring extra sensors successfully at a high rate of 84% on average. OmniVLA consistently outperforms all baseline configurations across the three tasks, demonstrating the effectiveness of our unified multi-sensory perception approach. On average, OmniVLA outperforms VLA-RGB model and VLA-RAW model by 59% and 28% in success rate respectively. OmniVLA also improves the task score by 0.45 and 0.17, respectively. This shows that sensor-modalities effectively provide new capabilities for the VLA model and our sensor-masked image representation significantly boosts performance.

Comparing base models. To evaluate the applicability of our approach on different base VLA models, we apply it to Pi0 [1] and compare the performance across 3 tasks. Table II shows that both models work, while SmolVLA provides better performance. This is likely due to SmolVLA is pre-trained with lerobot robot arm dataset [2]. Overall, this shows OmniVLA’s potential for combining with various RGB-only VLA models to provide extra new sensor-modality capabilities without starting from scratch.

Learning efficiency. We quantitatively compare the learning efficiency of using sensor-masked images and raw sensor data (no overlaying on RGB images), by training with an increasing number of thermal task episodes. As shown in Figure 6, OmniVLA constantly outperforms VLA-RAW model, achieving similar success rate with only around 50% of the training episodes. This shows the data efficiency of our proposed sensor-masked image representation, which

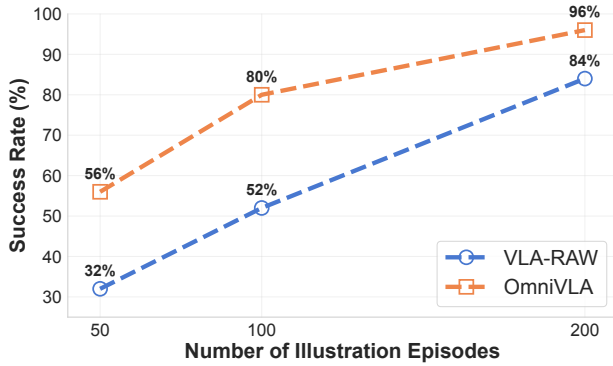


Fig. 6: Success rates over number of demonstration episodes.

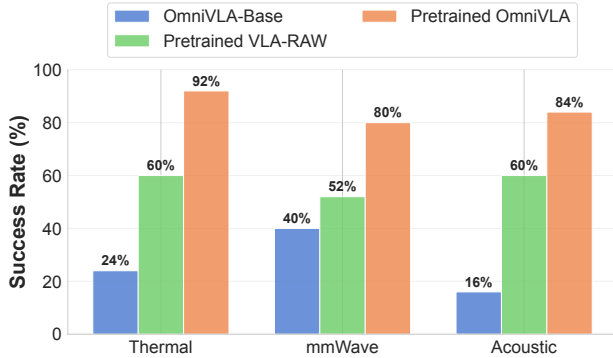


Fig. 7: Success rates when adapting to unseen tasks. We compare the pretrained OmniVLA model with two baselines, OmniVLA-Base (no pretraining), Pretrained VLA-RAW (pretrained with raw sensor images).

significantly reduces the required finetuning data amount compared with using raw sensor images.

C. Generalization Performance

Lastly, we evaluate how well our architecture can generalize to unseen tasks by pre-training on a mixed multi-sensory training dataset.

Multi-sensory pretraining. We construct a pretraining corpus of 800 demonstration episodes: 200 episodes for each of three sensor-modality tasks (total 600) and 200 episodes for generic pick-and-place with everyday objects.

We pre-train our OmniVLA model with the mixed dataset, then test the performance after performing few-shot learning with only 25 demonstration episodes of performing unseen tasks (matching common RGB-only VLA testing practice). We induce distribution shifts in object identity and materials based on tasks shown in Figure 5): (1) swap the drink type, (2) replace cardboard with foam boxes, and (3) substitute T-shirts with towels. The object locations are randomized within the feasible workspace of the robot arm and towels are randomly folded.

Baselines. To show the effectiveness, we implement two ablation baselines: (1) **OmniVLA-Base** — no multi-sensory pretraining, using OmniVLA model architecture with SmolVLA backbone weights, finetuned on the 25 demonstrations. (2) **Pretrained VLA-RAW** — OmniVLA

model pretrained on the same 800 episodes but feed raw sensor heatmaps/images directly (no segmentation/overlay), then perform the same 25-shot adaptation. Our method, OmniVLA, uses the identical pretraining and few-shot protocols but with sensor-masked images as input. To compare different methods in more detail, we decompose each task into two stages: **Stage 1** (select the correct target to interact with) and **Stage 2** (complete the subsequent manipulation). We report stage-wise success rates and scores (Table III), and overall task success rate (Figure 7).

As shown in Table III, OmniVLA shows substantially better generalization to unseen sensor-modality tasks than both baselines. In Stage 1, OmniVLA improves success rates by 17% over OmniVLA-Base and by 32% over Pretrained VLA-RAW on average. In Stage 2, OmniVLA yields gains of up to 76%. These results suggest two complementary effects: (i) the sensor-masked overlay makes sensor cues spatially aligned with RGB, which simplifies learning the selection policy (Stage 1); and (ii) large-scale multi-sensory pretraining supplies transferable manipulation priors that boost few-shot control (Stage 2). Consequently, the few-shot task success rate increases by 59% over OmniVLA-Base and 28% over Pretrained VLA-RAW on average, and up to 68% across the three unseen tasks, as shown in Figure 7. To conclude, both sensor-masked images and pretraining significantly improve the success rate on unseen tasks with a few episodes of finetuning, showing powerful generalization capability.

V. CONCLUSION

We introduce OmniVLA, a multi-modal vision-language-action model that equips robots with perception capabilities beyond the visible spectrum by integrating sensors, i.e., infrared, mmWave, and acoustic sensors. Our core design is the sensor-masked image, a unified representation that spatially grounds and semantically aligns diverse sensor data onto RGB images. This image-native approach allows reusing a pretrained vision encoder, enabling data-efficient learning with lightweight per-sensor projection layers. Through extensive real-world experiments, we demonstrated that OmniVLA significantly outperforms baseline models on challenging manipulation tasks, showing an average success rate of **84%**, exceeding the 25% of an RGB-only VLA and the 56% of a model trained on unprocessed sensor data. OmniVLA is a step towards creating more versatile and perceptive robots that can fully understand and interact with their physical surroundings.

REFERENCES

- [1] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, S. Jakubczak, T. Jones, L. Ke, S. Levine, A. Li-Bell, M. Mothukuri, S. Nair, K. Pertsch, L. X. Shi, J. Tanner, Q. Vuong, A. Walling, H. Wang, and U. Zhilinsky, “ π_0 : A vision-language-action flow model for general robot control,” 2024. [Online]. Available: <https://arxiv.org/abs/2410.24164>
- [2] M. Shukor, D. Aubakirova, F. Capuano, P. Kooijmans, S. Palma, A. Zouitine, M. Aractingi, C. Pascal, M. Russi, A. Marafioti *et al.*, “Smolvla: A vision-language-action model for affordable and efficient robotics,” *arXiv preprint arXiv:2506.01844*, 2025.

- [3] J. Guan, S. Madani, S. Jog, S. Gupta, and H. Hassanieh, "Through fog high-resolution imaging using millimeter wave radar," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 464–11 473.
- [4] L. Dodds, T. Boroushaki, K. Zhou, and F. Adib, "Non-line-of-sight 3d object reconstruction via mmwave surface normal estimation," in *Proceedings of the 23rd Annual International Conference on Mobile Systems, Applications and Services*, 2025, pp. 445–458.
- [5] T. Ren, S. Liu, A. Zeng, J. Lin, K. Li, H. Cao, J. Chen, X. Huang, Y. Chen, F. Yan *et al.*, "Grounded sam: Assembling open-world models for diverse visual tasks," *arXiv preprint arXiv:2401.14159*, 2024.
- [6] Y. Chebotar, Q. Vuong, K. Hausman, F. Xia, Y. Lu, A. Irpan, A. Kumar, T. Yu, A. Herzog, K. Pertsch *et al.*, "Q-transformer: Scalable offline reinforcement learning via autoregressive q-functions," in *Conference on Robot Learning*. PMLR, 2023, pp. 3909–3928.
- [7] J. Luo, P. Dong, J. Wu, A. Kumar, X. Geng, and S. Levine, "Action-quantized offline reinforcement learning for robotic skill learning," in *Conference on Robot Learning*. PMLR, 2023, pp. 1348–1361.
- [8] C. Ying, H. Zhongkai, X. Zhou, X. Xu, H. Su, X. Zhang, and J. Zhu, "Peac: Unsupervised pre-training for cross-embodiment reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 37, pp. 54 632–54 669, 2024.
- [9] A. Kumar, A. Singh, S. Tian, C. Finn, and S. Levine, "A workflow for offline model-free robotic reinforcement learning," *arXiv preprint arXiv:2109.10813*, 2021.
- [10] Y. Chebotar, K. Hausman, Y. Lu, T. Xiao, D. Kalashnikov, J. Varley, A. Irpan, B. Eysenbach, R. Julian, C. Finn *et al.*, "Actionable models: Unsupervised offline reinforcement learning of robotic skills," *arXiv preprint arXiv:2104.07749*, 2021.
- [11] A. Kumar, A. Singh, F. Ebert, M. Nakamoto, Y. Yang, C. Finn, and S. Levine, "Pre-training for robots: Offline rl enables learning new tasks from a handful of trials," *arXiv preprint arXiv:2210.05178*, 2022.
- [12] O. Mees, D. Ghosh, K. Pertsch, K. Black, H. R. Walke, S. Dasari, J. Hejna, T. Kreiman, C. Xu, J. Luo *et al.*, "Octo: An open-source generalist robot policy," in *First Workshop on Vision-Language Models for Navigation and Manipulation at ICRA 2024*, 2024.
- [13] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi *et al.*, "Openvla: An open-source vision-language-action model," *arXiv preprint arXiv:2406.09246*, 2024.
- [14] J. Wen, Y. Zhu, J. Li, Z. Tang, C. Shen, and F. Feng, "Dexvla: Vision-language model with plug-in diffusion expert for general robot control," *arXiv preprint arXiv:2502.05855*, 2025.
- [15] Y. Ze, Z. Chen, W. Wang, T. Chen, X. He, Y. Yuan, X. B. Peng, and J. Wu, "Generalizable humanoid manipulation with 3d diffusion policies," *arXiv preprint arXiv:2410.10803*, 2024.
- [16] C. Li, J. Wen, Y. Peng, Y. Peng, F. Feng, and Y. Zhu, "Pointvla: Injecting the 3d world into vision-language-action models," *arXiv preprint arXiv:2503.07511*, 2025.
- [17] V. Bhat, Y.-H. Lan, P. Krishnamurthy, R. Karri, and F. Khorrami, "3d cavla: Leveraging depth and 3d context to generalize vision language action models for unseen tasks," *arXiv preprint arXiv:2505.05800*, 2025.
- [18] D. Qu, H. Song, Q. Chen, Y. Yao, X. Ye, Y. Ding, Z. Wang, J. Gu, B. Zhao, D. Wang *et al.*, "Spatialvla: Exploring spatial representations for visual-language-action model," *arXiv preprint arXiv:2501.15830*, 2025.
- [19] M. A. Patratskiy, A. K. Kovalev, and A. I. Panov, "Spatial traces: Enhancing vla models with spatial-temporal understanding," *arXiv preprint arXiv:2508.09032*, 2025.
- [20] H. Zhen, X. Qiu, P. Chen, J. Yang, X. Yan, Y. Du, Y. Hong, and C. Gan, "3d-vla: A 3d vision-language-action generative world model," *arXiv preprint arXiv:2403.09631*, 2024.
- [21] J. Bi, K. Y. Ma, C. Hao, M. Z. Shou, and H. Soh, "Vla-touch: Enhancing vision-language-action models with dual-level tactile feedback," *arXiv preprint arXiv:2507.17294*, 2025.
- [22] J. Huang, S. Wang, F. Lin, Y. Hu, C. Wen, and Y. Gao, "Tactile-vla: Unlocking vision-language-action model's physical knowledge for tactile generalization," *arXiv preprint arXiv:2507.09160*, 2025.
- [23] W. Zhao, P. Ding, M. Zhang, Z. Gong, S. Bai, H. Zhao, and D. Wang, "Vlas: Vision-language-action model with speech instructions for customized robot manipulation," *arXiv preprint arXiv:2502.13508*, 2025.
- [24] Y. Hong, Z. Zheng, P. Chen, Y. Wang, J. Li, and C. Gan, "Multiply: A multisensory object-centric embodied large language model in 3d world," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 26 406–26 416.
- [25] Z. Liu, H. Tang, A. Amini, X. Yang, H. Mao, D. Rus, and S. Han, "Bevfusion: Multi-task multi-sensor fusion with unified bird's-eye view representation," *arXiv preprint arXiv:2205.13542*, 2022.
- [26] P. Wolters, J. Gilg, T. Teepe, F. Herzog, A. Laouichi, M. Hofmann, and G. Rigoll, "Unleashing hydra: Hybrid fusion, depth consistency and radar for unified 3d perception," *arXiv preprint arXiv:2403.07746*, 2024.
- [27] Z. Lin, Z. Liu, Z. Xia, X. Wang, Y. Wang, S. Qi, Y. Dong, N. Dong, L. Zhang, and C. Zhu, "Rcbevdet: Radar-camera fusion in bird's eye view for 3d object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024, pp. 14 928–14 937.
- [28] E. Palladin, R. Dietze, P. Narayanan, M. Bijelic, and F. Heide, "Samfusion: Sensor-adaptive multimodal fusion for 3d object detection in adverse weather," in *European Conference on Computer Vision*. Springer, 2024, pp. 484–503.
- [29] J. Kim, M. Seong, and J. W. Choi, "Crt-fusion: Camera, radar, temporal fusion using motion information for 3d object detection," *Advances in Neural Information Processing Systems*, vol. 37, pp. 108 625–108 648, 2024.
- [30] W. Xiong, Z. Zou, Q. Zhao, F. He, and B. Zhu, "Lxlvl2: Enhanced lidar excluded lean 3d object detection with fusion of 4d radar and camera," *IEEE Robotics and Automation Letters*, 2025.
- [31] T. Boroushaki, I. Perper, M. Nachin, A. Rodriguez, and F. Adib, "Rfusion: Robotic grasping via rf-visual sensing and learning," in *Proceedings of the 19th ACM conference on embedded networked sensor systems*, 2021, pp. 192–205.
- [32] T. Boroushaki, J. Leng, I. Clester, A. Rodriguez, and F. Adib, "Robotic grasping of fully-occluded objects using rf perception," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 923–929.
- [33] T. Boroushaki, L. Dodds, N. Naeem, and F. Adib, "Fusebot: Rf-visual mechanical search," *Robotics: Science and Systems 2022*, 2022.
- [34] T. Wang, X. Mao, C. Zhu, R. Xu, R. Lyu, P. Li, X. Chen, W. Zhang, K. Chen, T. Xue, X. Liu, C. Lu, D. Lin, and J. Pang, "Embodiedscan: A holistic multi-modal 3d perception suite towards embodied ai," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024, pp. 19 757–19 767.
- [35] X. Han, S. Chen, Z. Fu, Z. Feng, L. Fan, D. An, C. Wang, L. Guo, W. Meng, X. Zhang *et al.*, "Multimodal fusion and vision-language models: A survey for robot vision," *arXiv preprint arXiv:2504.02477*, 2025.
- [36] Q. Xu, C. Jiang, Y. Han, B. Wang, and K. R. Liu, "Waveforming: An overview with beamforming," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 1, pp. 132–149, 2017.
- [37] A. Hurst, A. Lerer, A. P. Goucher, A. Perelman, A. Ramesh, A. Clark, A. Ostrow, A. Welihinda, A. Hayes, A. Rafford *et al.*, "Gpt-4o system card," *arXiv preprint arXiv:2410.21276*, 2024.
- [38] N. Ravi, V. Gabeur, Y.-T. Hu, R. Hu, C. Ryal, T. Ma, H. Khedr, R. Rädle, C. Rolland, L. Gustafson *et al.*, "Sam 2: Segment anything in images and videos," *arXiv preprint arXiv:2408.00714*, 2024.
- [39] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, Q. Jiang, C. Li, J. Yang, H. Su *et al.*, "Grounding dino: Marrying dino with grounded pre-training for open-set object detection," in *European conference on computer vision*. Springer, 2024, pp. 38–55.
- [40] S. Karamcheti, S. Nair, A. Balakrishna, P. Liang, T. Kollar, and D. Sadigh, "Prismatic vlms: Investigating the design space of visually-conditioned language models," in *Forty-first International Conference on Machine Learning*, 2024.
- [41] H. Liu, C. Li, Q. Wu, and Y. J. Lee, "Visual instruction tuning," *Advances in neural information processing systems*, vol. 36, pp. 34 892–34 916, 2023.
- [42] J. Chen, D. Zhu, X. Shen, X. Li, Z. Liu, P. Zhang, R. Krishnamoorthi, V. Chandra, Y. Xiong, and M. Elhoseiny, "Minigt-v2: large language model as a unified interface for vision-language multi-task learning," *arXiv preprint arXiv:2310.09478*, 2023.
- [43] R. Cadene, S. Alibert, A. Soare, Q. Galloudec, A. Zouitine, S. Palma, P. Kooijmans, M. Aractingi, M. Shukor, D. Aubakirova, M. Russi, F. Capuano, C. Pascal, J. Choghari, J. Moss, and T. Wolf, "Lerobot: State-of-the-art machine learning for real-world robotics in pytorch," <https://github.com/huggingface/lerobot>, 2024.