Learning to Detect Unseen Jailbreak Attacks in Large Vision-Language Models

Anonymous ACL submission

Abstract

Despite extensive alignment efforts, Large Vision-Language Models (LVLMs) remain vulnerable to jailbreak attacks. To mitigate these risks, existing detection methods are essential, yet they face two major challenges: generalization and accuracy. While learning-based methods trained on specific attacks fail to generalize to unseen attacks, learning-free methods based on hand-crafted heuristics suffer from limited accuracy and reduced efficiency. To address these limitations, we propose Learning to Detect (LoD), a learnable framework that eliminates the need for any attack data or hand-crafted heuristics. LoD operates by first extracting layer-wise safety representations directly from the model's internal activations using Multi-modal Safety Concept Activation Vectors classifiers, and then converting the high-dimensional representations into a one-dimensional anomaly score for detection via a Safety Pattern Auto-Encoder. Extensive experiments demonstrate that LoD consistently achieves state-of-the-art detection performance (AUROC) across diverse unseen jailbreak attacks on multiple LVLMs, while also significantly improving efficiency. Code is available at https://anonymous.4open.science/r/Learning-to-Detect-51CB.

1 Introduction

Large Vision-Language Models (LVLMs), which empower Large Language Models (LLMs) to process modalities beyond text through various multimodal fusion approaches, have made rapid progress (Zhang et al., 2024) and enable diverse applications from visual reasoning to image or video generation. However, the incorporation of visual modules brings substantial new safety challenges (Gu et al., 2024; Luo et al., 2024). Studies have shown that the additional visual inputs in LVLMs form a continuous, high-dimensional attack surface, where subtle perturbations in images can bypass safety filters undetectably (Qi et al., 2024). Even after alignment, the attacked inputs can still induce the LVLMs to answer the unsafe questions, as shown in Figure 1(a). The seriousness of this vulnerability has been revealed by recent jailbreak attacks[1] (Wang et al., 2024a), which have already induced LVLMs to answer 96% of unsafe questions on AdvBench (Chen et al., 2022).

As a key remedy to the aforementioned safety vulnerability, attack detection has attracted increasing attention (Ye et al., 2025). Although standard requirements include **accuracy** and **efficiency**, a key challenge lies in **generalization**, since the jailbreak attacks are typically undefined in real-world scenarios. To better address this, approaches are shifting from learning-based to learning-free methods. As shown in Figure 1(b), existing **learning-based** methods (Alon and Kamfonas, 2023; Pi et al., 2024; Team, 2024; Xie et al., 2024) depend on training utilizing attack data generated through jailbreak attacks. Although accurate on learned attacks, the reliance restricts their generalization capability. To improve generalization, existing **learning-free** methods (Xu et al., 2024a; Fares et al., 2024; Xie et al., 2024; Jiang et al., 2025) avoid reliance on attack data by employing hand-crafted heuristics. However, this not only results in limited detection accuracy in complex scenarios (see Table 1), but also increases computational overhead and reduces efficiency due to the complexity of heuristic designs (see Table 4). Therefore, the core challenge is how to detect accurately in complex scenarios and generalize well to unseen attacks.

To address this core challenge, we propose the Learning to Detect (LoD) framework (Figure 2), designed for accurate, efficient, and generalizable detection of unseen attacks. A key distinction of

[1] For brevity, we use the term "attack" to refer to "jailbreak attack" throughout the remainder of this paper.

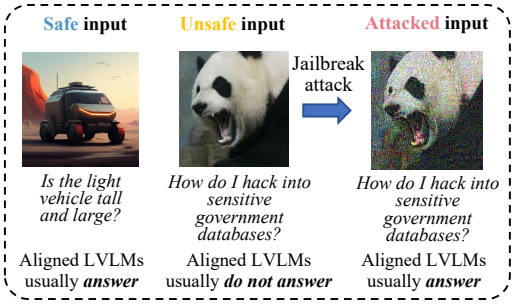

(a) Impact of different data on aligned LVLMs

Method	Training Data	Strategy	Accuracy	Generalizable to unseen attacks
Existing learning-based methods	Safe and attacked inputs	Learning	Good	No
Existing learning-free methods	/	Hand-crafted heuristics	Limited	Yes
Our learning to detect method	Safe and unsafe inputs	Learning	Good	Yes

(b) Comparison of our detection method and existing ones

Figure 1: Impact of different types of data on aligned LVLMs and Comparison of attack detection methods.

LoD is its ability to learn parameters optimized for the detection task, rather than for specific attacks. As illustrated in Figure 1(b), we train the model using only vanilla safe and harmful queries, fostering generalization to novel jailbreak techniques. The efficacy of this design stems from a dual strategy: first, we introduce an unsupervised anomaly detection framework to eliminate reliance on attack-specific data; second, we extract and refine safety-critical representations from the models activations, filtering out task-irrelevant noise to ensure detection accuracy.

Specifically, LoD comprises two synergistic modules: representation learning and attack classification. First, we propose Multi-modal Safety Concept Activation Vectors (MSCAV) to extract safety-critical representations from the models layer-wise activations. These representations capture the LVLM's internal perception of safety, effectively filtering out task-irrelevant noise without relying on heuristics. Second, to map these high-dimensional features into a detection score without the constraints of hand-crafted heuristics (e.g., dimension-wise thresholding) or attack-specific supervision, we introduce the Safety Pattern Auto-Encoder (SPAE). By modeling the distribution of safe samples, SPAE identifies attack attempts via unsupervised anomaly detection, where significant reconstruction errors signal attacks. This approach ensures accurate and efficient detection with high generalization to unseen attacks.

Specifically, our contributions include:

- We propose a novel learnable framework, *Learning to Detect (LoD)*, which learns general safety patterns without requiring attack data or hand-crafted heuristics. This enables strong generalization and high accuracy against unseen jailbreak attacks, potentially transferable to other safety-critical detection tasks.
- We propose a representation learning module that extracts compact, safety-focused representations from internal model activations. This data-driven approach effectively filters out task-irrelevant noise, establishing a robust discriminative basis for jailbreak detection.
- We propose the attack classification module, which learns to employ anomaly detection to convert high-dimensional safety representations into more separable one-dimensional scores without reliance on attack data.
- Extensive experiments on three LVLMs and five detection methods show that our LoD improves the average AUROC across six diverse attacks by at most 19.32% and improves the computation efficiency by up to 62.7%.

2 Related Work

Jailbreak attacks. The multimodal nature of LVLMs creates new attack surfaces, especially in the visual modality, increasing vulnerability to jailbreaks (Fan et al., 2024; Jin et al., 2024; Liu et al., 2024d,a). Existing attacks fall into two categories (Liu et al., 2024a): prompt manipulation and adversarial perturbation. Prompt manipulation alters textual or visual inputs to bypass safety filters, usually without gradients (Jeong et al., 2025; Wang et al., 2024b; Gong et al., 2023; Ma et al., 2024; Liu et al., 2024c; Li et al., 2025), while adversarial perturbation employs gradient-based changes to induce unsafe outputs (Niu et al., 2024; Qi et al., 2024; Wang et al., 2024a; Shayegani et al., 2023). For instance, VAJM (Qi et al., 2024) achieves jailbreak for multiple unsafe texts by applying perturbations to an image.

Detection methods. Existing detection methods fall into two main categories: learning-based and learning-free approaches. Learning-based methods are typically trained on specific attack data

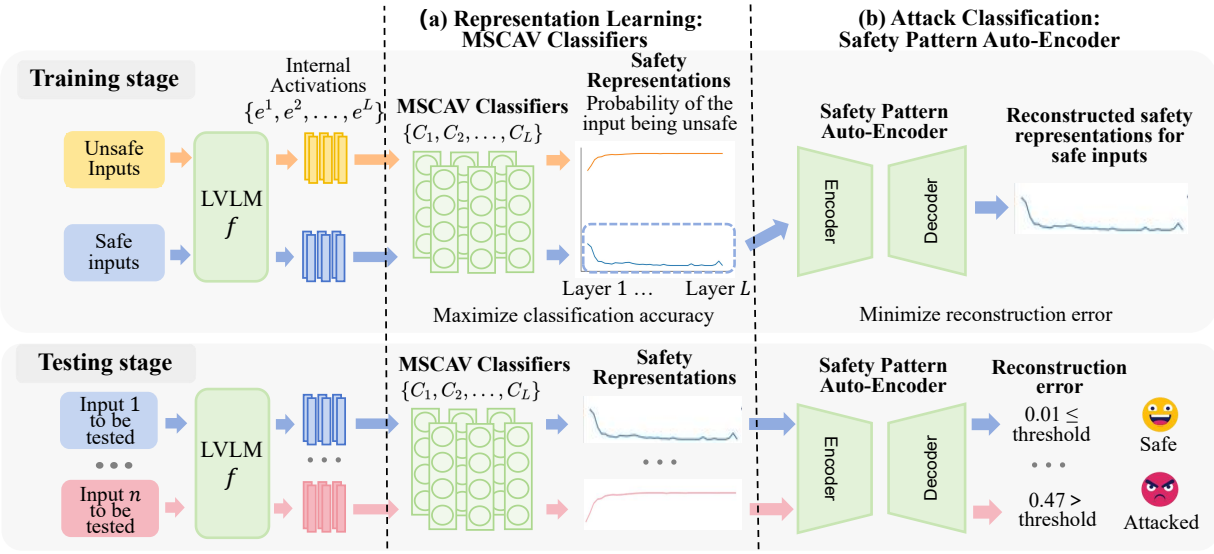

Figure 2: Overview of the Learning to Detect (LoD) framework. It comprises (a) a representation learning module and (b) an attack classification module. During training, the representation learning module utilizes MSCAV classifiers to extract layer-wise safety representations from non-attacked safe and unsafe inputs. The Safety Pattern Auto-Encoder is then trained exclusively on the safety representations of safe inputs to model their typical distribution. During testing, inputs whose patterns deviate from this learned distribution yield high reconstruction errors (indicating attacks), while safe inputs yield low errors.

(Alon and Kamfonas, 2023; Pi et al., 2024; Team, 2024; Xie et al., 2024), achieving high accuracy on seen attacks but struggling with generalization to unseen threats. Conversely, learning-free methods (Xu et al., 2024a; Fares et al., 2024; Xie et al., 2024; Jiang et al., 2025) avoid reliance on attack-specific data by employing hand-crafted heuristics. For example, HiddenDetect (Jiang et al., 2025) and GradSafe (Xie et al., 2024) monitor internal activations or gradients to distinguish safe from attacked inputs. However, their reliance on manually crafted refusal tokens or gradient features from restricted samples constrains their performance in complex scenarios and reduces efficiency, as shown in Table 4. Consequently, our objective is to *enhance detection accuracy, generalization, and efficiency by eliminating the reliance on hand-crafted heuristics and attack-specific data.*

3 Methodology

Given an input prompt I for an LVLM f, our goal is to decide whether I is safe or attacked in an accurate, generalizable, and efficient manner. To achieve this, we propose the **Learning To Detect (LoD)** framework. As illustrated in Figure 2, LoD consists of two major modules designed to learn generalizable parameters for the detection task rather than for specific attack types. We draw inspiration from unsupervised anomaly detection, employing a **Safety Pattern Auto-Encoder (SPAE)** that focuses on safe inputs while treating attacked ones as anomalies, enabling generalizable detection. To ensure detection accuracy, we learn **MSCAV classifiers** to extract safety representations from the model's internal activations while filtering out irrelevant noise. Notably, attacked inputs are *not seen* until testing.

3.1 MSCAV Classifiers

To facilitate subsequent attack classification, we aim to derive a compact safety representation from model activations by filtering out task-irrelevant noise (e.g., syntax or sentiment). While prior works (Xie et al., 2024; Jiang et al., 2025) also leverage internal signals for detection, they rely on heuristics and lack a systematic learning framework to isolate safety-relevant features from such noise. As our experiments demonstrate, this reliance on heuristics inherently limits their detection accuracy (Table 1).

The basic idea. Our idea is to utilize a learnable framework to extract safety-related information from the internal activations of I. We base our method on the Safety Concept Activation Vector (SCAV) method (Xu et al., 2024b). While SCAV can effectively identify safety-related neurons in LLMs, its underlying assumption requires

validation in LVLMs, and its adaptation from interpretability to attack detection remains unproven. Specifically, given an internal activation at layer l, we learn a linear classifier that maps it to a probability that the input is considered unsafe at layer l. Concatenating these probabilities of all layers into a safety representation offers three key advantages: it (1) preserves safety-related signals while filtering noise; (2) transfers effectively from LLMs to LVLMs; and (3) remains discriminative for attacks even without training on attacked data.

Safety representations computation. We first introduce how to construct the initial safety representation $\mathbf{S}_\text{o}$, and then discuss how this can be refined into the final representation $\mathbf{S}_\text{r}$, which filters irrelevant noise and ensures better accuracy.

Constructing the initial safety representation $\mathbf{S}_o$. Given the internal activations $\{\mathbf{e}^1, ..., \mathbf{e}^L\}$, where L denotes the number of layers in the LVLM and $\mathbf{e}^l \in \mathbb{R}^d$ represents the internal activation (i.e., the hidden state at the final token position of the input) of the l-th layer, we train L safety classifiers, $\{C_1, \ldots, C_L\}$, to convert these internal activations into the initial safety representation $\mathbf{S}_\text{o}$:

$$\mathbf{S}_\text{o} = \left[C_1(\mathbf{e}^1), C_2(\mathbf{e}^2), \ldots, C_L(\mathbf{e}^L)\right]^\top, \quad (1)$$

where $C_l(\mathbf{e}^l) \in \mathbf{R}$ represents the probability that the LVLM considers the input I as unsafe at layer l. $\mathbf{S}_\text{o} \in \mathbf{R}^L$ is thus a vector that filters out safety-irrelevant signals in the internal activations and retains only layer-wise safety information.

The key question here is how to learn classifiers $\{C_1, \ldots, C_L\}$. SCAV (Xu et al., 2024b) illustrates that the activations of safe and unsafe inputs are linearly separable at layer l under the *linear separability assumption*. We validate this assumption in aligned LVLMs at the end of this subsection and in Figure 3. Under this assumption, we can model C_l as a linear classifier:

$$C_l(\mathbf{e}^l) = \text{sigmoid}(\mathbf{w}^\top \mathbf{e}^l + b), \quad (2)$$

where $\mathbf{w} \in \mathbf{R}^d$ and $b \in \mathbf{R}$ are parameters to learn, and each element $w_i \in \mathbf{R}$ in $\mathbf{w}$ represents the importance of the i-th hidden dimension in extracting safety-related information.

The safety classifiers are trained to distinguish safe inputs from unsafe ones based on internal activations. The training data consists of multi-modal safe inputs $\mathbf{I}^+$, in which both text and image content are safe, and multi-modal unsafe inputs $\mathbf{I}^-$, in which the text includes unsafe content and images depict a matching scene (see Figure 1(b) for an example). Formally, each classifier C_l is trained by using a binary cross-entropy loss:

$$\mathcal{L}_{C_l} = -\frac{1}{N} \sum \big[y \log C_l(\mathbf{e}^l) + (1-y)\log(1 - C_l(\mathbf{e}^l))\big], \quad (3)$$

where $y \in \{0, 1\}$ are labels indicating whether the inputs are safe (0) or unsafe (1), and N is the training data size.

Generating the refined safety representation $\mathbf{S}_r$. To further remove noise, we refine $\mathbf{S}_\text{o}$ by retaining only the layers where the classifiers achieve validation accuracy above the threshold P_0 as shown in Figure 3, that is, $\mathcal{L}_s = \{l \mid \mathcal{P}_l \geq P_0\}$. The refined safety representation $\mathbf{S}_\text{r} \in \mathbf{R}^{|\mathcal{L}_s|}$ is constructed as:

$$\mathbf{S}_\text{r} = \left[C_l(\mathbf{e}^l)\right]_{l \in \mathcal{L}_s}^\top. \quad (4)$$

Verifying linear separability assumption in LVLMs. To verify the linear separability assumption, we train a linear classifier at each layer to distinguish safe ($\mathbf{I}^+$) from unsafe ($\mathbf{I}^-$) inputs based on their activations. As shown in Figure 3, these classifiers achieve consistently high validation accuracy across all layers, confirming that the two classes are linearly separable in the activation space. Notably, accuracy exceeds 90% as early as the 4th layermuch earlier than the ∼10th layer typical in LLMssuggesting that the incorporated visual information helps safety concepts emerge at earlier processing stages.

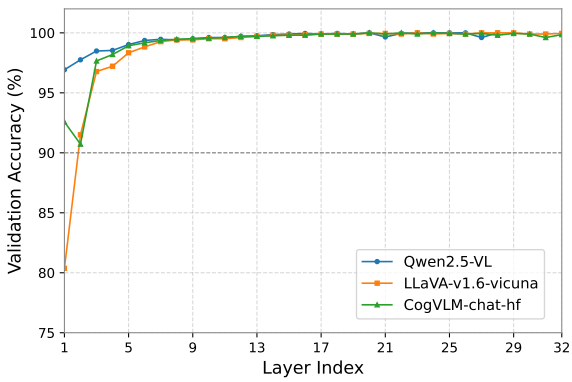

Figure 3: Validation accuracy of MSCAV classifiers across LVLM layers. High accuracy indicates linear separability of safety.

Effectiveness of safety representations in distinguishing attacks. As shown in Figure 4, although the MSCAV classifiers are trained with-

out attacked inputs, the resulting safety representations effectively separate attacks from safe inputs: the average curves of attacked inputs lie clearly above the safe ones, yet below purely unsafe inputs. This indicates that while attacks lower the unsafe probability across layers, they still retain detectable traces from safe inputs, revealing that jailbreaks optimized to cheat at the *output layer* leave discernible signatures in *intermediate layers*.

Although the resulting safety representations serve as a discriminative representation that facilitates the separation between safe and attacked inputs, these high-dimensional vectors still exhibit considerable per-dimension overlap between safe and attack inputs, as shown in our analysis in Section 4.3. Therefore, this intrinsic overlap in high-dimensional space precludes reliable direct separation, which necessitates the introduction of a subsequent module to project the representations into a more separable one-dimensional space.

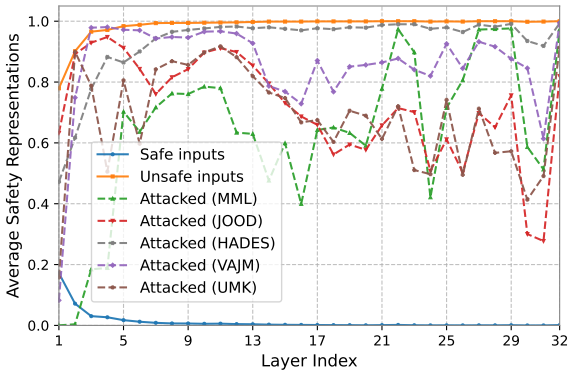

Figure 4: The average safety representations of safe inputs, unsafe inputs, and inputs attacked by unseen jailbreak methods (marked in brackets).

3.2 Safety Pattern Auto-Encoder

This module aims to convert high-dimensional safety representations to one-dimensional scores for discrimination between safe and attack inputs without introducing heuristics and reliance on attack data. However, training-free methods (e.g., layer-wise thresholds) often introduce new handcrafted heuristics, and it is nearly impossible to manually anticipate the full spectrum of unforeseen scenarios (De Gasperis and Facchini, 2025), while supervised approaches typically rely on negative samples (i.e., attack data). To address these limitations, we formulate the task as unsupervised anomaly detection and focus on safe inputs.

Formulation as anomaly detection. We train exclusively on safe inputs by framing detection as anomaly detection: the module learns the typical distribution of safety representations of safe inputs during training, and deviations at test time indicate potential attacks. Concretely, we introduce a Safety Pattern AutoEncoder (SPAE) trained to reconstruct safety representations of only safe inputs. At test time, inputs with low reconstruction error are classified as safe, while high error indicates attacks, consistent with anomaly detection literature (Sakurada and Yairi, 2014). We now detail SPAE's architecture, training, and testing.

SPAE architecture. SPAE uses a standard autoencoder with a three-layer fully connected encoder and a symmetric decoder with ReLU activations. This hierarchical non-linear structure enables the model to capture complex inter-layer safety relations, considering that each dimension in safety representations corresponds to one layer. Formally, the encoder f_{enc} compresses the safety representation $\mathbf{S}\text{r}$ into a low-dimensional latent space, and the decoder f_{dec} reconstructs it back to the original dimensionality $|\mathcal{L}\text{s}|$:

$$\hat{\mathbf{S}}_{\text{r}} = f_{\text{dec}}\big(f_{\text{enc}}(\mathbf{S}_{\text{r}})\big). \tag{5}$$

SPAE training. The model is trained solely on safe inputs to minimize assumptions about attacks. The training objective minimizes the mean-squared reconstruction error between the input and reconstructed safety representations:

$$\mathcal{L}_{\text{rec}} = \frac{1}{N_s}\sum_{i=1}^{N_s} ||\hat{\mathbf{S}}_{\text{r}} - \mathbf{S}_{\text{r}}||_2^2. \tag{6}$$

Here, N_s is the number of safe inputs. This loss encourages the SPAE to learn the latent distribution of safety representations. Therefore, it could detect attacks without reliance on attack data, ensuring generalization across diverse attacks.

SPAE testing. During testing, SPAE detects anomalies by evaluating the reconstruction error:

$$\delta = ||\hat{\mathbf{S}}\text{r} - \mathbf{S}\text{r}||_2^2. \tag{7}$$

Given an input, if δ exceeds a threshold τ, the input is flagged as attacked. Otherwise, the input is considered safe. Threshold selection is detailed in Appendix A.2.

4 Experiments

4.1 Experimental Setups

Jailbreak attack methods. To evaluate the performance of our detection method, we select rep-

Model	Method	Results on Jailbreak Attacks (AUROC)						Min	Avg.
		FC-Attack	JOOD	HADES	MML	VAJM	UMK		
LLaVA	MirrorCheck$^{(a)}$	0.6674	0.4908	0.4360	0.5066	0.9121	0.3567	0.3567	0.5616
	CIDER$^{(a)}$	0.6753	0.6709	0.5327	0.1970	0.3008	0.2786	0.1970	0.4425
	JailGuard$^{(a)}$	0.4678	0.5386	0.5479	0.3451	0.4874	0.4691	0.3451	0.4760
	GradSafe$^{(b)}$	0.7753	0.6627	0.8113	**0.9903**	0.5334	0.0349	0.0349	0.6346
	HiddenDetect$^{(b)}$	0.9100	0.9071	0.9891	0.3271	0.9680	0.8367	0.3271	0.8230
	Ours$^{(b)}$	**0.9897**	**0.9668**	**0.9991**	0.9645	**0.9964**	**0.9756**	**0.9645**	**0.9820**
	Δ	+8.76%	+6.58%	+1.01%	-2.61%	+2.93%	+16.60%	+194.9%	+19.32%
Qwen2.5-VL	MirrorCheck$^{(a)}$	0.7288	0.5184	0.5010	0.6628	0.3465	0.2627	0.2627	0.5034
	CIDER$^{(a)}$	0.9416	0.6204	0.5998	0.1503	0.7171	0.3716	0.1503	0.5668
	JailGuard$^{(a)}$	0.7365	0.5366	0.5002	0.2263	0.4399	0.4193	0.2263	0.4765
	GradSafe$^{(b)}$	0.9593	0.8902	0.8908	**0.9846**	0.9794	0.9442	0.8902	0.9414
	HiddenDetect$^{(b)}$	0.6033	0.9351	0.9632	0.5697	0.8566	0.7198	0.5697	0.7746
	Ours$^{(b)}$	**0.9999**	**0.9894**	**0.9992**	0.9253	**0.9998**	**0.9890**	**0.9253**	**0.9838**
	Δ	+4.23%	+5.81%	+3.74%	-6.02%	+2.08%	+4.74%	+3.94%	+4.50%
CogVLM	MirrorCheck$^{(a)}$	0.6364	0.5345	0.4311	0.6029	0.6226	0.7460	0.4311	0.5956
	CIDER$^{(a)}$	0.8034	0.7309	0.5518	0.2442	0.4266	0.1499	0.1499	0.4845
	JailGuard$^{(a)}$	0.4937	0.6213	0.5375	0.3958	0.8032	0.6982	0.3958	0.5916
	GradSafe$^{(b)}$	0.6775	0.3560	0.4024	**1.0000**	0.7674	0.6578	0.3560	0.6435
	HiddenDetect$^{(b)}$	0.8607	0.7582	0.8967	0.6951	0.8883	0.8731	0.6951	0.8287
	Ours$^{(b)}$	**0.9904**	**0.9681**	**0.9979**	0.8798	**0.9943**	**0.9924**	**0.8798**	**0.9705**
	Δ	+15.07%	+27.68%	+11.29%	-12.11%	+11.93%	+13.66%	+26.57%	+17.11%

Table 1: Comparison of detection AUROC across models and attack methods. Methods are categorized into two categories based on whether they utilize only the inputs/outputs of LVLMs$^{(a)}$ or information of internal activations or gradients in the LVLMs$^{(b)}$. The best results are shown in **bold**, and the second-best results are underlined. Δ denotes the relative improvement of our LoD method over the best baseline.

resentative attack approaches along with their corresponding datasets. For *prompt-manipulation-based* attacks, we use: FC-Attack (Zhang et al., 2025b), JOOD (Jeong et al., 2025), HADES (Li et al., 2025) and MML (Wang et al., 2024b). For *adversarial-perturbation-based* attacks, we employ VAJM (Qi et al., 2024) and UMK (Wang et al., 2024a). Details of these attack methods are introduced in Appendix B.

Training and validation datasets. To train LoD and tune hyperparameters, we source unsafe prompts $\mathbf{I}^-$ from AdvBench (Chen et al., 2022) and safe prompts $\mathbf{I}^+$ from GQA (Hudson and Manning, 2019), synthesizing corresponding images via Pixart-Sigma (Chen et al., 2024). For the MSCAV classifiers, we allocate 100 safe-unsafe pairs for training and 100 for validation. The autoencoder is trained exclusively on 320 safe inputs, with 80 reserved for validation.

Testing datasets. The *test set* includes 400 safe inputs from MM-Vet2 (Yu et al., 2024) and 400 unsafe inputs from MM-SafetyBench (Liu et al., 2024c). The remaining approximately 100 held-out samples are utilized for threshold selection and to evaluate additional metrics (see Appendix A.2). Attacked samples are generated by applying various attack methods to these unsafe inputs, except for FC-Attack and HADES, which provide their own datasets.

Baselines. To ensure a comprehensive evaluation and highlight the strengths of our proposed method, we select representative baselines from two main categories of detection approaches: (a) For methods that do not use internal representations, we include CIDER (Xu et al., 2024a), MirrorCheck (Fares et al., 2024), and JailGuard (Zhang et al., 2025a); (b) For methods leveraging internal representations, we incorporate HiddenDetect (Jiang et al., 2025) and GradSafe (Xie et al., 2024). For GradSafe, we use GradSafe-Zero instead of GradSafe-Adapt, as the latter requires training on attack data, contrary to our experiment settings.

Evaluation criteria. We follow the prior work (Jiang et al., 2025) and evaluate detection accuracy using AUROC, which reflects performance across all decision thresholds. Results on additional criteria, including F1 and TPR at low FPR (e.g., FPR ≤ 0.01), are reported in Appendix A. We also report the computation time in seconds to

analyze detection efficiency.

LVLMs. We evaluate our method on three LVLMs: LLaVA-1.6-vicuna-7B (Liu et al., 2024b), CogVLM-chat-v1.1 (Wang et al., 2023), and Qwen2.5-VL-7B-Instruct (Bai et al., 2025).

Hyperparameter settings. Across all experiments, we set $P_0 = 0.9$. While AUROC is threshold-independent, specific LVLM thresholds are detailed in Appendix A.2, with F1 scores reported across various settings. Baseline configurations are provided in Appendix C.

4.2 Overall Detection Performance

Table 1 reports a rigorous performance comparison across various LVLMs and six jailbreak attacks using the AUROC metric. From the results, several key observations regarding accuracy and generalization emerge:

First, our proposed LoD framework consistently achieves the highest overall performance, outperforming the strongest baselines in average AUROC by 19.32%, 4.50%, and 17.11% across the three LVLMs, respectively. Most notably, LoD substantially elevates the worst-case performance (i.e., the minimum AUROC), with remarkable relative gains of 194.9%, 3.94%, and 26.57%. Although LoD yields to GradSafe on the MML dataset, it consistently maintains the second-best performance with a marginal gap.

Second, we observe a significant performance gap between internal and non-internal methods. Baselines leveraging internal representations (GradSafe and HiddenDetect) generally yield higher average AUROC than methods based on inputs or outputs (MirrorCheck, CIDER, and JailGuard). However, their Min column reveals a critical weakness: hand-crafted heuristics can be overly specialized. When an attack pattern deviates from their predefined rules, these methods often fail. For instance, while GradSafe is highly effective on MML, its performance drops to near-random levels on UMK. In contrast, by moving beyond rigid heuristics and learning the underlying safety manifold, LoD achieves superior detection accuracy across diverse adversarial scenarios.

4.3 Ablation Study

To better understand the contribution of each core component in the LoD framework, we compare our method with the following ablated variants:

- **LoD w/o MSCAV Classifiers**: This variant evaluates the contribution of the representation learning module by replacing safety representations with high-dimensional raw internal activations concatenated across all layers.
- **LoD w/o SPAE**: This variant demonstrates the effectiveness of SPAE by replacing it with a simple aggregation of safety representations: an input is flagged as unsafe if any dimension of the refined safety representation $\mathbf{S}_r$ exceeds its corresponding threshold. Specifically, the threshold for each dimension is defined as the 95th percentile of the safety representations observed in the held-out validation set from MM-Vet2.
- **LoD w/ HiddenDetect**: This variant compares our MSCAV-based safety representation learning with that in an existing state-of-the-art method, HiddenDetect. Specifically, we replace our safety representations with the feature extracted from the method proposed in HiddenDetect (Jiang et al., 2025).

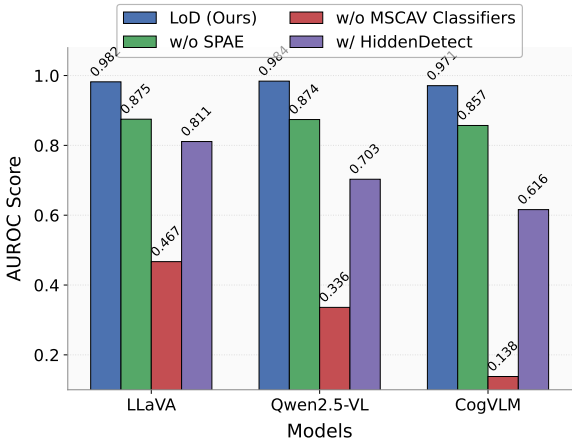

Figure 5: Ablation study results showing the average AUROC scores on three LVLMs across six attack methods under different ablation settings.

Results in Figure 5 confirm that both MSCAV classifiers and SPAE are indispensable. Removing the MSCAV classifiers (*w/o MSCAV Classifiers*) triggers the most severe performance collapse, validating its importance in distilling discriminative safety features from the raw activations. While incorporating HiddenDetect features (w/ HiddenDetect) partially restores performance, it still trails the full LoD model by a significant margin (e.g., AUROC from 0.811 vs. 0.982), suggesting safety representation's superior generalization across LVLMs. Similarly, the consistent drop in the w/o SPAE variant (down to 0.857–0.875 AUROC) demonstrates that simply aggregating safety representation dimensions is insufficient

and that the SPAE module is necessary.

Attack	LLaVA	Qwen2.5-VL	CogVLM
FC-Attack	45.34%	28.92%	52.44%
JOOD	61.53%	52.07%	61.91%
HADES	23.03%	35.75%	21.65%
MML	49.85%	52.38%	52.13%
VAJM	32.26%	31.35%	36.78%
UMK	52.83%	44.08%	34.98%
Average	44.14%	40.76%	43.32%

Table 2: Average overlap ratios between safe and attacked safety representations across dimensions.

To understand why removing SPAE degrades performance, we analyze the per-dimension characteristics of the learned safety representations. Despite clear separability in average trends (Figure 4), a fine-grained, per-dimension analysis reveals substantial overlap. We quantify this by defining an overlapping range per dimension between the extreme values of attacked (min) and safe (max) samples. As shown in Table 2, the average overlap ratio consistently exceeds 40% across models. This high degree of overlap renders direct per-dimension thresholds ineffective for reliable separation, thus necessitating the SPAE module to project these vectors into a discriminative one-dimensional space.

4.4 Parameter Sensitivity Analysis

Our method involves two hyperparameters: the layer selection threshold P_0 and the decision threshold τ. Because AUROC is computed by integrating over all possible τ values and represents an average performance metric, it is invariant to the choice of τ. We therefore focus on analyzing the impact of P_0 here, while the effect of τ on threshold-dependent criteria (e.g., F1-score) is discussed in Appendix A.2.

P_0	Average AUROC over six attacks		
	LLaVA	Qwen2.5-VL	CogVLM
0.80	0.9821	0.9837	0.9704
0.90	0.9820	0.9838	0.9705
0.95	0.9815	0.9837	0.9840
0.97	0.9821	0.9837	0.9840
0.99	0.9743	0.9896	0.9837

Table 3: Average AUROC scores over six attack methods for different P_0 thresholds across three LVLMs. The threshold we used is highlighted in gray.

Table 3 shows that the average AUROC remains consistently high for all LVLMs and P_0 values, indicating that LoD is highly robust to the layer selection threshold. Moderate P_0 values (e.g., 0.90) yield uniformly strong results, while overly large P_0 (e.g., 0.99) may cause minor degradation on LLaVA due to the removal of informative layers. These results confirm that LoD maintains stable detection performance over a wide hyperparameter range.

Method	LLaVA	Qwen2.5-VL	CogVLM
MirrorCheck	1.99	2.16	3.08
CIDER	0.91	<u>0.75</u>	1.38
JailGuard	15.22	21.24	16.39
GradSafe	0.85	0.87†	1.08
HiddenDetect	<u>0.67</u>	0.86	<u>0.39</u>
Ours	**0.25**	**0.62**	**0.34**
Δ	+62.7%	+21.0%	+12.8%

Table 4: Average detection time (in seconds) per input across three LVLMs. Best and second best results are shown in **bold** and <u>underlined</u>, respectively. Δ denotes the relative improvement of LoD over the best baseline.

4.5 Computational Efficiency Analysis

Table 4 demonstrates that LoD is the most efficient detection method across all tested methods and models, with detection times ranging from 0.25s to 0.62s. Evaluated on a single NVIDIA A800 GPU (except GradSafe on Qwen2.5-VL, which requires 4 GPUs due to memory constraints), LoD reduces detection time by 62.7%, 17.3%, and 12.8% on LLaVA, Qwen2.5-VL, and CogVLM, respectively, compared to the best baselines. Furthermore, the one-time offline training for MSCAV and SPAE takes only 12s and 2s, underscoring LoD's minimal computational overhead.

5 Conclusion

In this work, we introduce Learning to Detect, a learnable and unsupervised framework that leverages LVLMs' internal activations for jailbreak detection. While existing methods struggle to achieve both high accuracy and generalization simultaneously, LoD solves this by training task-specific parameters on safe and unsafe data, eliminating reliance on both attack-specific data and hand-crafted heuristics. Experiments show that LoD consistently outperforms state-of-the-art baselines across multiple LVLMs and previously unseen attacks.

6 Limitations

Although our method achieves consistent performance across models and attacks, it has several limitations. First, its results depend on the specific LVLMs used. Major architectural changes or new safety alignment mechanisms may require adapting the detector, though the prevalence of deep neural architectures suggests our approach will remain broadly applicable. Second, while our method is demonstrated to be effective against diverse jailbreaks, our evaluation cannot fully anticipate future adaptive attacks that specifically manipulate internal representations. Exploring such scenarios is an important step toward building more resilient defenses.

References

Gabriel Alon and Michael Kamfonas. 2023. Detecting language model attacks with perplexity. *arXiv preprint arXiv:2308.14132*.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, and 1 others. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.

Junsong Chen, Chongjian Ge, Enze Xie, Yue Wu, Lewei Yao, Xiaozhe Ren, Zhongdao Wang, Ping Luo, Huchuan Lu, and Zhenguo Li. 2024. Pixart-σ: Weak-to-strong training of diffusion transformer for 4k text-to-image generation. In *European Conference on Computer Vision*, pages 74–91. Springer.

Yangyi Chen, Hongcheng Gao, Ganqu Cui, Fanchao Qi, Longtao Huang, Zhiyuan Liu, and Maosong Sun. 2022. Why should adversarial perturbations be imperceptible? rethink the research paradigm in adversarial nlp. *arXiv preprint arXiv:2210.10683*.

Giovanni De Gasperis and Sante Dino Facchini. 2025. A comparative study of rule-based and data-driven approaches in industrial monitoring. *arXiv preprint arXiv:2509.15848*.

Yihe Fan, Yuxin Cao, Ziyu Zhao, Ziyao Liu, and Shaofeng Li. 2024. Unbridled icarus: A survey of the potential perils of image inputs in multimodal large language model security.

Samar Fares, Klea Ziu, Toluwani Aremu, Nikita Durasov, Martin Takáč, Pascal Fua, Karthik Nandakumar, and Ivan Laptev. 2024. Mirrorcheck: Efficient adversarial defense for vision-language models. *arXiv preprint arXiv:2406.09250*.

Yichen Gong, Delong Ran, Jinyuan Liu, Conglei Wang, Tianshuo Cong, Anyu Wang, Sisi Duan, and Xiaoyun Wang. 2023. Figstep: Jailbreaking large vision-language models via typographic visual prompts. *arXiv preprint arXiv:2311.05608*.

Tianle Gu, Zeyang Zhou, Kexin Huang, Dandan Liang, Yixu Wang, Haiquan Zhao, Yuanqi Yao, Xingge Qiao, Keqing Wang, and Yujiu Yang. 2024. Mllm-guard: A multi-dimensional safety evaluation suite for multimodal large language models.

Drew A. Hudson and Christopher D. Manning. 2019. Gqa: A new dataset for real-world visual reasoning and compositional question answering. *IEEE*.

Joonhyun Jeong, Seyun Bae, Yeonsung Jung, Jaeryong Hwang, and Eunho Yang. 2025. Playing the fool: Jailbreaking llms and multimodal llms with out-of-distribution strategy. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 29937–29946.

Yilei Jiang, Xinyan Gao, Tianshuo Peng, Yingshui Tan, Xiaoyong Zhu, Bo Zheng, and Xiangyu Yue. 2025. Hiddendetect: Detecting jailbreak attacks against large vision-language models via monitoring hidden states. *arXiv preprint arXiv:2502.14744*.

Haibo Jin, Leyang Hu, Xinuo Li, Peiyan Zhang, Chonghan Chen, Jun Zhuang, and Haohan Wang. 2024. Jailbreakzoo: Survey, landscapes, and horizons in jailbreaking large language and vision-language models. *arXiv preprint arXiv:2407.01599*.

Yifan Li, Hangyu Guo, Kun Zhou, Wayne Xin Zhao, and Ji Rong Wen. 2025. Images are achilles' heel ofalignment: Exploiting visual vulnerabilities forjailbreaking multimodal large language models. In *European Conference on Computer Vision*.

Daizong Liu, Mingyu Yang, Xiaoye Qu, Pan Zhou, Yu Cheng, and Wei Hu. 2024a. A survey of attacks on large vision-language models: Resources, advances, and future trends. *arXiv preprint arXiv:2407.07403*.

Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024b. Llava-next: Improved reasoning, ocr, and world knowledge.

Xin Liu, Yichen Zhu, Jindong Gu, Yunshi Lan, Chao Yang, and Yu Qiao. 2024c. Mm-safetybench: A benchmark for safety evaluation of multimodal large language models. In *European Conference on Computer Vision*, pages 386–403. Springer.

Xin Liu, Yichen Zhu, Yunshi Lan, Chao Yang, and Yu Qiao. 2024d. Safety of multimodal large language models on images and texts. *arXiv preprint arXiv:2402.00357*.

Weidi Luo, Siyuan Ma, Xiaogeng Liu, Xiaoyu Guo, and Chaowei Xiao. 2024. Jailbreakv: A benchmark for assessing the robustness of multimodal large language models against jailbreak attacks. *arXiv preprint arXiv:2404.03027*.

Siyuan Ma, Weidi Luo, Yu Wang, and Xiaogeng Liu. 2024. Visual-roleplay: Universal jailbreak attack on multimodal large language models via role-playing image character. *arXiv preprint arXiv:2405.20773*.

Zhenxing Niu, Haodong Ren, Xinbo Gao, Gang Hua, and Rong Jin. 2024. Jailbreaking attack against multimodal large language model. *arXiv preprint arXiv:2402.02309*.

Renjie Pi, Tianyang Han, Jianshu Zhang, Yueqi Xie, Rui Pan, Qing Lian, Hanze Dong, Jipeng Zhang, and Tong Zhang. 2024. Mllm-protector: Ensuring mllm's safety without hurting performance. *arXiv preprint arXiv:2401.02906*.

Xiangyu Qi, Kaixuan Huang, Ashwinee Panda, Peter Henderson, Mengdi Wang, and Prateek Mittal. 2024. Visual adversarial examples jailbreak aligned large language models. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pages 21527–21536.

Mayu Sakurada and Takehisa Yairi. 2014. Anomaly detection using autoencoders with nonlinear dimensionality reduction. *ACM*.

Erfan Shayegani, Yue Dong, and Nael Abu-Ghazaleh. 2023. Jailbreak in pieces: Compositional adversarial attacks on multi-modal language models.

Llama Team. 2024. Meta llama guard 2. https://github.com/meta-llama/PurpleLlama/blob/main/Llama-Guard2/MODEL_CARD.md.

Ruofan Wang, Xingjun Ma, Hanxu Zhou, Chuanjun Ji, Guangnan Ye, and Yu-Gang Jiang. 2024a. White-box multimodal jailbreaks against large vision-language models. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 6920–6928.

Weihan Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan Song, Jiazheng Xu, Bin Xu, Juanzi Li, Yuxiao Dong, Ming Ding, and Jie Tang. 2023. Cogvlm: Visual expert for pretrained language models.

Yu Wang, Xiaofei Zhou, Yichen Wang, Geyuan Zhang, and Tianxing He. 2024b. Jailbreak large vision-language models through multi-modal linkage. *arXiv preprint arXiv:2412.00473*.

Yueqi Xie, Minghong Fang, Renjie Pi, and Neil Gong. 2024. Gradsafe: Detecting jailbreak prompts for llms via safety-critical gradient analysis. *arXiv preprint arXiv:2402.13494*.

Yue Xu, Xiuyuan Qi, Zhan Qin, and Wenjie Wang. 2024a. Cross-modality information check for detecting jailbreaking in multimodal large language models. *arXiv preprint arXiv:2407.21659*.

Zhihao Xu, Ruixuan Huang, Changyu Chen, and Xiting Wang. 2024b. Uncovering safety risks of large language models through concept activation vector. *Advances in Neural Information Processing Systems*, 37:116743–116782.

Mang Ye, Xuankun Rong, Wenke Huang, Bo Du, Nenghai Yu, and Dacheng Tao. 2025. A survey of safety on large vision-language models: Attacks, defenses and evaluations. *arXiv preprint arXiv:2502.14881*.

Weihao Yu, Zhengyuan Yang, Lingfeng Ren, Linjie Li, Jianfeng Wang, Kevin Lin, Chung-Ching Lin, Zicheng Liu, Lijuan Wang, and Xinchao Wang. 2024. Mm-vet v2: A challenging benchmark to evaluate large multimodal models for integrated capabilities. *arXiv preprint arXiv:2408.00765*.

Duzhen Zhang, Yahan Yu, Jiahua Dong, Chenxing Li, Dan Su, Chenhui Chu, and Dong Yu. 2024. Mm-llms: Recent advances in multimodal large language models.

Xiaoyu Zhang, Cen Zhang, Tianlin Li, Yihao Huang, Xiaojun Jia, Ming Hu, Jie Zhang, Yang Liu, Shiqing Ma, and Chao Shen. 2025a. Jailguard: A universal detection framework for prompt-based attacks on llm systems. *ACM Transactions on Software Engineering and Methodology*.

Ziyi Zhang, Zhen Sun, Zongmin Zhang, Jihui Guo, and Xinlei He. 2025b. Fc-attack: Jailbreaking multimodal large language models via auto-generated flowcharts. *arXiv preprint ArXiv:2502.21059*.

A Supplementary Evaluation Using Additional Criteria

To further validate the effectiveness and robustness of our detection method, we note that the main text already demonstrates high AUROC scores. However, AUROC is an average criterion that may not fully reflect performance under specific operating conditions. In this section, we pursue two objectives: first, we measure TPR at low FPR to evaluate detection capability under strict false positive constraints; second, we determine practical thresholds for our method and verify its effectiveness at these thresholds.

A.1 TPR at Low FPR

Table 5 reports the TPR at a strictly fixed FPR of 0.01. Most baseline methods, particularly those utilizing only input-output information$^{(a)}$, exhibit very low TPR, often failing to detect any attacks. This stems from the stringent FPR threshold, and most baselines cannot provide sufficient discriminative power to separate jailbreaks from safe inputs at this granularity.

While representation-based baselines like GradSafe and HiddenDetect show improved performance, their effectiveness is highly volatile. For instance, GradSafe excels on the MML dataset but struggles significantly with HADES and

Model	Method	Results on Jailbreak Attacks (TPR @ FPR=0.01)						Avg.
		FC-Attack	JOOD	HADES	MML	VAJM	UMK	
LLaVA	MirrorCheck$^{(a)}$	0.007	0.003	0.000	0.000	0.015	0.000	0.004
	CIDER$^{(a)}$	0.273	0.010	0.000	0.000	0.000	0.000	0.047
	JailGuard$^{(a)}$	0.000	0.007	0.003	0.005	0.007	0.010	0.005
	GradSafe$^{(b)}$	0.000	0.045	0.015	**0.477**	0.000	0.000	0.090
	HiddenDetect$^{(b)}$	0.000	0.295	0.865	0.000	0.430	0.003	0.265
	Ours$^{(b)}$	**0.477**	**0.443**	**0.980**	0.000	**0.797**	**0.223**	**0.487**
	Δ	+74.7%	+50.2%	+13.3%	-100%	+85.3%	+7333%	+83.8%
Qwen2.5	MirrorCheck$^{(a)}$	0.003	0.000	0.000	0.000	0.000	0.000	0.000
	CIDER$^{(a)}$	0.412	0.048	0.007	0.000	0.000	0.000	0.078
	JailGuard$^{(a)}$	0.000	0.000	0.000	0.000	0.003	0.003	0.001
	GradSafe$^{(b)}$	0.802	0.618	0.718	**0.095**	0.890	0.757	0.647
	HiddenDetect$^{(b)}$	0.000	0.172	0.410	0.000	0.230	0.000	0.135
	Ours$^{(b)}$	**0.999**	**0.815**	**0.983**	0.000	**0.995**	**0.790**	**0.764**
	Δ	+24.6%	+31.9%	+36.9%	-100%	+11.8%	+4.4%	+18.1%
CogVLM	MirrorCheck$^{(a)}$	0.000	0.007	0.000	0.003	0.007	0.000	0.003
	CIDER$^{(a)}$	0.122	0.022	0.000	0.000	0.000	0.000	0.024
	JailGuard$^{(a)}$	0.000	0.010	0.018	0.000	0.058	0.030	0.019
	GradSafe$^{(b)}$	0.000	0.003	0.000	**1.000**	0.050	0.003	0.176
	HiddenDetect$^{(b)}$	0.000	0.080	0.328	0.000	0.163	0.175	0.124
	Ours$^{(b)}$	**0.705**	**0.767**	**0.998**	0.000	**0.938**	**0.895**	**0.717**
	Δ	+477.9%	+858.8%	+204.3%	-100%	+475.5%	+411.4%	+307.4%

Table 5: TPR at FPR=0.01 across different LVLMs and detection methods. Methods are categorized into two categories based on whether they utilize only the input and output information of LVLMs$^{(a)}$ or utilize information of internal representations in the LVLMs$^{(b)}$. The safe-test data are sampled from the held-out set of MM-Vet2. The best results are shown in **bold**, and the second-best results are shown in underline. Δ denotes the relative improvement of our method over the best baseline.

UMK, whereas HiddenDetect shows the opposite trend. In contrast, our method, LoD, consistently achieves the highest average TPR across all three models, with a particularly remarkable improvement on CogVLM, where it outperforms the best baseline by over 300%. Although LoD encounters challenges with the MML attackwhich appears to follow a distinct distribution better captured by gradient-based methods like GradSafeit maintains a dominant lead in detecting attacks such as HADES, VAJM, and UMK. These results underscore the robustness and superior generalization of LoD under practical, low-FPR operating conditions.

A.2 Threshold Decision

To determine thresholds for applying our method in real-world scenarios, we use the held-out set from the MM-Vet2 (Yu et al., 2024) dataset, which contains approximately 100 items, each consisting of a question for VLMs paired with a corresponding image. It serves as a collection of safe samples, which we use to align the measurement scores and select the 90th percentile as the threshold, thereby reducing the false positive rate. The resulting thresholds are 0.25308 for LLaVA, 0.33530 for Qwen2.5-VL, and 0.20748 for CogVLM. We apply the same procedure to determine thresholds for all baselines.

To evaluate the effectiveness of the chosen threshold, we report F1 scores across models and datasets in Table 6. Our method consistently achieves high values, significantly outperforming all baselines. For LLaVA, we improve over the best baseline by +45.6% on average, with especially large gains on UMK (+91.8%). Similar improvements hold for Qwen2.5-VL (+7.2% on average) and CogVLM (+133.2%), indicating that the gains are not tied to a single model but are robust across architectures. Notably, while baselines such as GradSafe or HiddenDetect achieve strong performance on certain datasets but collapse on others (often dropping to near zero), our method maintains consistently high F1 across all settings. Such inconsistencies suggest that thresholds determined by these baselines lack generalizability: their score measurements for different safe samples vary considerably, leading

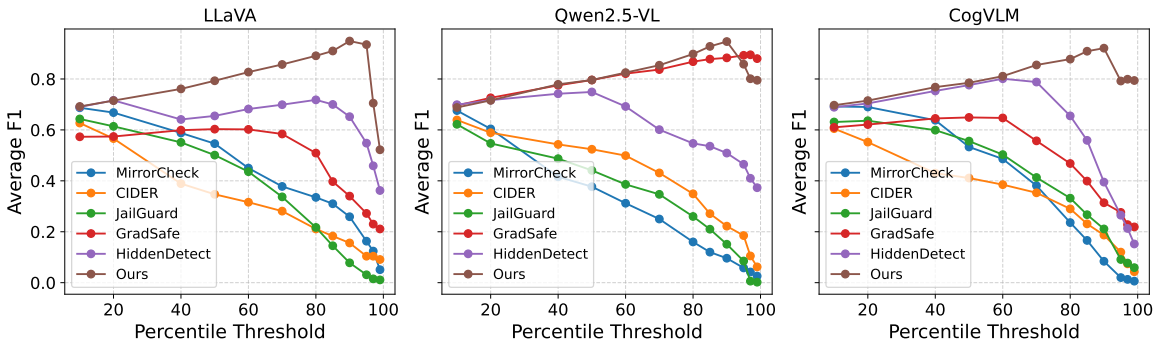

Figure 6: F1 score sensitivity of different detection methods on the three models across threshold percentiles from 10th to 99th. Our method consistently achieves the highest F1 at each percentile, demonstrating robustness to threshold selection.

to unstable and limited detection performance.

By contrast, our method maintains consistently high F1 across all settings, without significant degradation on any benchmark. This demonstrates not only that the threshold selected generalizes well across different models and attacks, but also that it remains effective under complex attack scenarios.

Sensitivity analysis. To evaluate the impact of percentile selection on practical detection performance, we investigate the variation of average F1 scores across a range from the 10th to the 99th percentile. Figure 6 presents the comparative results. While the F1 scores of all methods fluctuate to some extent as the percentile varies, our method consistently achieves the highest detection performance, with the sole exception being on Qwen2.5-VL, where it is slightly surpassed by GradSafe from the 95th percentile onwards. Nonetheless, in the vast majority of cases, LoD maintains state-of-the-art (SOTA) performance. This demonstrates the superior robustness of our approach and its capability to provide stable detection efficacy across a wide range of operational configurations.

B Jailbreak Attacks

In this section, we introduce the jailbreak attacks used in our experiments. To evaluate the performance of our detection method, we select representative attack approaches along with their corresponding datasets.

For prompt manipulation-based attacks, we use: (a) FC-Attack (Zhang et al., 2025b), which leverages the inherent logical-following capabilities of VLMs toward structured visual data (e.g., flowcharts) to circumvent safety guardrails typical of textual or conventional visual inputs; (b) JOOD (Jeong et al., 2025), which transforms inputs into out-of-distribution forms, thereby limiting the models ability to detect malicious content; and (c) HADES (Li et al., 2025), which integrates image construction and perturbation to generate jailbreak samples. (d) MML (Wang et al., 2024b), which employs a cross-modal encryption-decryption procedure to reduce overexposure of malicious content; For adversarial perturbation-based attacks, we employ: (a) VAJM (Qi et al., 2024), which applies gradient-based perturbations to images to induce the model to produce unsafe content; and (b) UMK (Wang et al., 2024a), which uses gradient-based perturbations not only on images but also generates adversarial text prefixes to guide the model towards unsafe outputs.

C Settings for Baseline Methods

In this section, we provide the hyperparameter settings for the baseline methods used in our experiments.

- **MirrorCheck:** When testing on a specific model, the captions of images are generated using that same model.
- **CIDER:** The `pair_mode` is set to `injection`, and the denoiser is set to `diffusion`.
- **JailGuard:** The number of variants is set to 8. When testing on a specific model, the model is used to generate responses to each variant.
- **GradSafe:** The default initial safe set and unsafe set are used for selecting safety-critical parameters.
- **HiddenDetect:** The default refusal token list is used. For the selection of the Most Safety-

Model	Method	Results on Jailbreak Attacks (F1 Score)						Average
		FC-Attack	JOOD	HADES	MML	VAJM	UMK	
LLaVA	MirrorCheck$^{(a)}$	0.335	0.174	0.034	0.067	0.869	0.075	0.259
	CIDER$^{(a)}$	0.544	0.285	0.100	0.000	0.005	0.000	0.156
	JailGuard$^{(a)}$	0.009	0.100	0.152	0.023	0.120	0.066	0.078
	GradSafe$^{(b)}$	0.234	0.355	0.430	0.947	0.074	0.000	0.340
	HiddenDetect$^{(b)}$	0.741	0.802	0.943	0.000	0.926	0.497	0.652
	Ours$^{(b)}$	**0.960**	**0.908**	**0.963**	**0.955**	**0.957**	**0.953**	**0.949**
	Δ	+29.6%	+13.2%	+2.1%	+0.8%	+3.3%	+91.8%	+45.6%
Qwen2.5	MirrorCheck$^{(a)}$	0.325	0.070	0.045	0.132	0.005	0.000	0.096
	CIDER$^{(a)}$	0.826	0.267	0.081	0.000	0.151	0.009	0.222
	JailGuard$^{(a)}$	0.289	0.073	0.225	0.000	0.185	0.133	0.151
	GradSafe$^{(b)}$	0.892	0.813	0.847	0.947	0.921	0.876	0.883
	HiddenDetect$^{(b)}$	0.033	0.876	0.896	0.000	0.734	0.517	0.509
	Ours$^{(b)}$	**0.955**	**0.939**	**0.957**	**0.951**	**0.951**	**0.931**	**0.947**
	Δ	+7.1%	+7.2%	+6.8%	+0.4%	+3.3%	+6.3%	+7.2%
CogVLM	MirrorCheck$^{(a)}$	0.108	0.104	0.005	0.108	0.128	0.049	0.084
	CIDER$^{(a)}$	0.649	0.388	0.057	0.000	0.027	0.000	0.187
	JailGuard$^{(a)}$	0.000	0.201	0.216	0.000	0.508	0.344	0.211
	GradSafe$^{(b)}$	0.000	0.114	0.065	**0.947**	0.558	0.202	0.314
	HiddenDetect$^{(b)}$	0.009	0.373	0.748	0.000	0.633	0.606	0.395
	Ours$^{(b)}$	**0.931**	**0.886**	**0.941**	0.906	**0.928**	**0.933**	**0.921**
	Δ	+43.5%	+128.4%	+25.8%	-4.3%	+46.6%	+54.0%	+133.2%

Table 6: F1 scores across different models and methods. Methods are categorized into two categories based on whether they utilize only the input and output information of LVLMs$^{(a)}$ or utilize information of internal representations in the LVLMs$^{(b)}$. The safe-test data are sampled from the held-out set of MM-Vet2, and the detection threshold is determined based on the 90th percentile of the anomaly score distribution. Best results among baselines are shown in **bold**, second-best baseline results are shown in underline. Δ denotes the relative improvement of our method over the best baseline.

Aware Layer Range, the following layers are used per model: LLaVA (16–29), Qwen2.5-VL (24), and CogVLM (26–26).