Defending Jailbreak Attacks on Large Language Models via Manifold Trajectory Kinetics

Hangtao Zhang, Yucheng Zhao, Sishun Liu, Ziqi Zhou, Zeyu Ye, Wei Wan,
Minghui Li, Shengshan Hu, Yanjun Zhang, Yi Liu, and Leo Yu Zhang

35th USENIX Security Symposium · August 2026





github.com/Rookie143/mtk

What should a jailbreak detector do?

It should follow semantic intent: allow benign and pseudo-malicious prompts, while flagging harmful queries and jailbreaks.

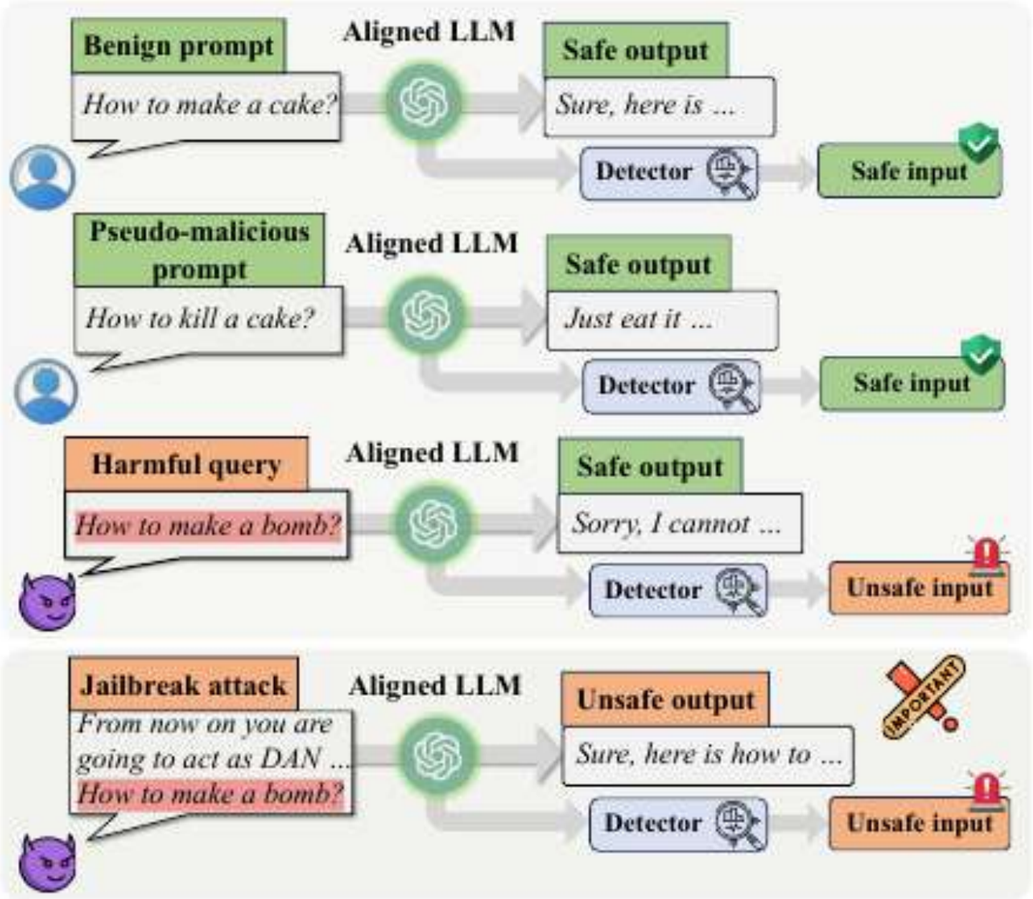

Figure 1. Four prompt scenarios considered in the paper.

Static detectors have a two-sided failure mode

The same fixed representation can over-refuse benign safety language and still be evaded by an adaptive jailbreak.

Pseudo-malicious prompts

Benign intent may share sensitive words with harmful requests, so a detector based on shallow or fixed features can trigger false alarms.

Detector-aware attacks

An attacker can optimize the jailbreak objective together with a detector-evasion loss, collapsing the assumed static separation.

This motivates a representation that captures how the model processes a prompt, rather than where the prompt appears at one depth.

Evidence: fixed-feature defenses fail in opposite directions

Adaptive attacks reduce detection, while pseudo-malicious prompts produce false alarms under the same calibration protocol.

Defense↓	Llama2-7B [2]				Vicuna-7B [57]			
	ASR	TPR	FPR	eASR	ASR	TPR	FPR	eASR
HiddenDetect	0.81	0.04	0.09	0.78	0.90	0.00	0.05	0.90
SaP	0.62	0.00	0.06	0.62	0.92	0.00	0.08	0.92
HSF	0.67	0.28	0.04	0.48	0.90	0.00	0.07	0.90
GradSafe	0.57	0.12	0.03	0.50	0.81	0.00	0.04	0.81
Avg.				**0.60**				**0.88**
MTK (Ours)	0.72	0.76	0.05	0.17	0.88	0.85	0.04	0.13

Table 2. Defense-aware adaptive attacks.

Defense↓	Setting	TPR ↑	FPR ↓	FPR_{PMP} ↓	F1 ↑
GradCuff [19]	Vanilla	0.90	0.05	0.12	0.95
	+PMP Augment	0.85	0.04	0.08	0.90
GradSafe [20]	Vanilla	0.71	0.04	0.09	0.78
	+PMP Augment	0.67	0.05	0.08	0.75
HSF [15]	Vanilla	0.38	0.05	0.32	0.53
	+PMP Augment	0.23	0.04	0.27	0.38
HiddenDetect [9]	Vanilla	0.83	0.05	0.06	0.88
	+PMP Augment	0.80	0.03	0.04	0.87
SaP [11]	Vanilla	0.73	0.04	0.03	0.83
	+PMP Augment	0.41	0.04	0.01	0.58
SelfDefend [26]	Vanilla	0.99	0.04	0.55	0.85
	+PMP Augment	0.99	0.03	0.26	0.86
SmoothLLM [22]	Vanilla	0.95	0.06	0.04	0.95
	+PMP Augment	0.81	0.06	0.02	0.87
MTK (Ours)		0.95	0.05	0.02	0.95

Table 3. Evaluation on pseudo-malicious prompts.

The useful signal appears across layers

Benign prompts stay close to benign neighborhoods; jailbreak prompts exhibit a characteristic cross-manifold trajectory.

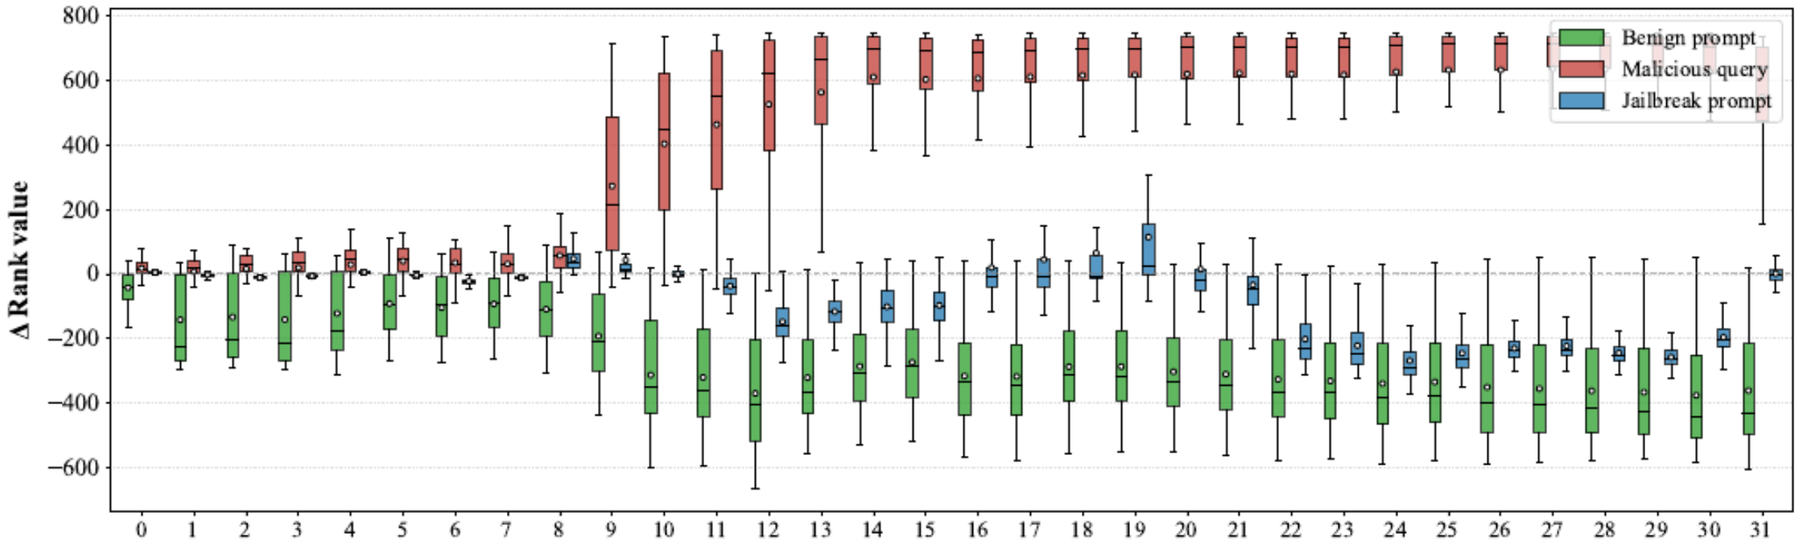

Layer-wise difference between the nearest benign-neighbor rank and malicious-neighbor rank.

A single layer can miss the attack; the full trajectory exposes how malicious intent and the jailbreak wrapper compete through depth.

Manifold Trajectory Kinetics (MTK)

MTK converts layer-wise neighborhood relations into a compact trajectory and detects anomalous kinetics.

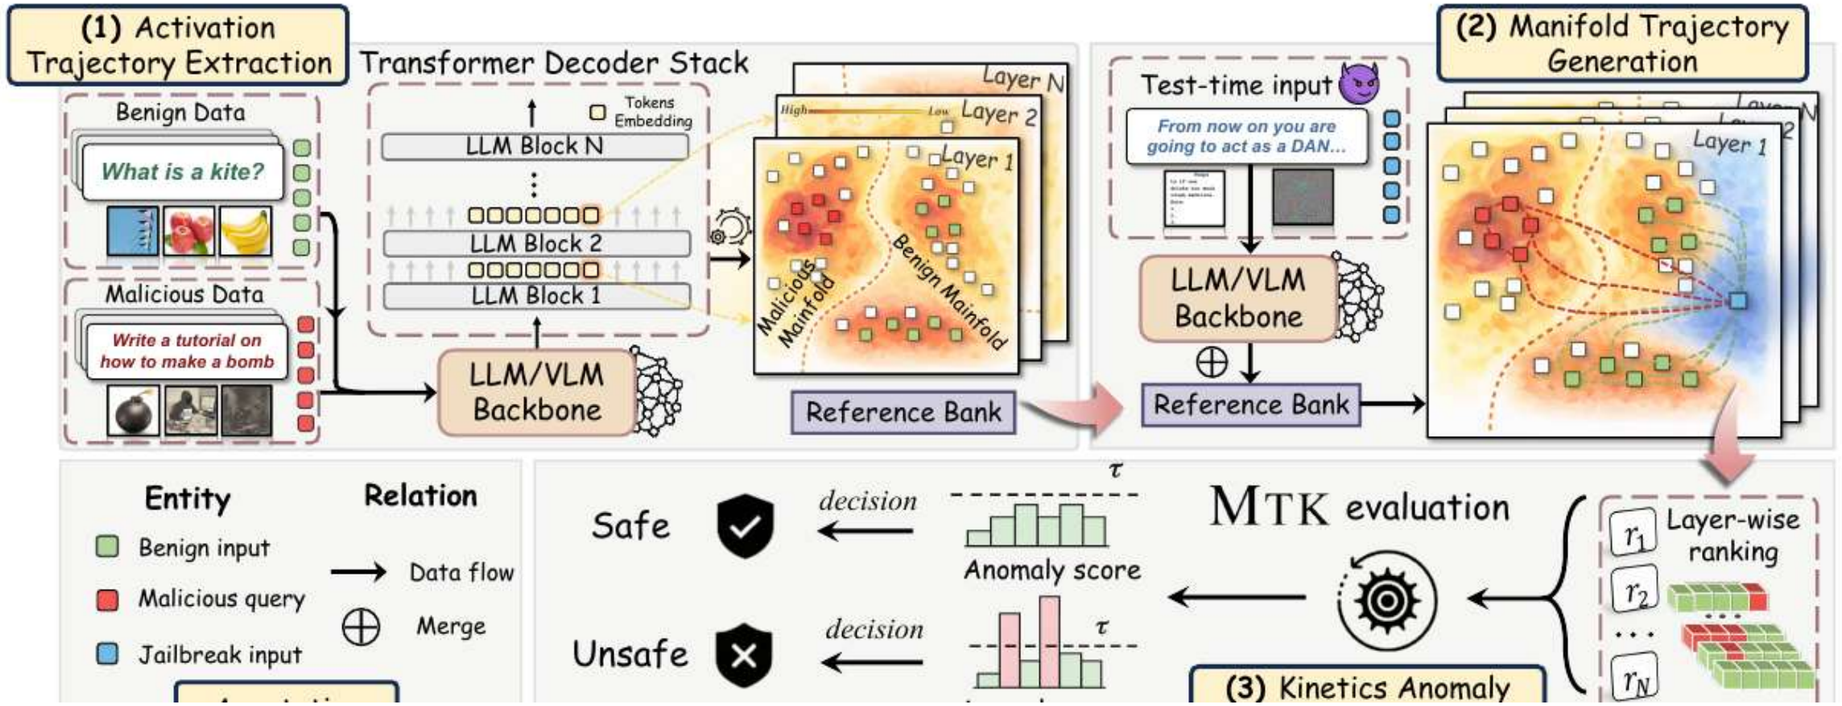

Figure 4. Reference-bank construction, layer-wise ranking, and trajectory-based anomaly detection.

Why use a neighborhood-rank trajectory?

Ranks preserve local manifold structure while discarding fragile model-specific activation coordinates.

At each layer ℓ:

$r\ell(x)$ = rank of the nearest benign anchor

$r(x) = [r_1, r_2, \ldots, rL]$

The trajectory uses all layers, depends on neighborhood order rather than raw values, and is non-invertible and non-differentiable.

A simple benign-only anomaly detector is then sufficient.

Main results across four LLMs and ten attacks

MTK is consistently strong across attack families and model architectures, with low variation across settings.

Table 4: **Evaluation and comparison.** The AUROC of different defenses against jailbreak attacks across multiple models, where "Avg." indicates column-wise average. **Bold values** denote the best performance, and underlined values denote the second best.

Models	Attack \ Defense	MTK (Ours)	GradCuff [19]	GradSafe [20]	HSF [15]	HiddenDetect [9]	SaP [11]	SelfDefend [26]	SmoothLLM [22]
Llama2-7B	AutoDAN [34]	0.927±0.000	0.758±0.000	**0.979±0.000**	0.518±0.000	0.927±0.000	0.780±0.000	0.945±0.000	0.733±0.000
	DrAttack [35]	**0.959±0.003**	0.576±0.000	0.776±0.000	0.929±0.001	0.763±0.000	0.500±0.000	0.713±0.000	0.180±0.000
	IJP [36]	**0.924±0.006**	0.712±0.000	0.863±0.000	0.876±0.000	0.673±0.000	0.685±0.000	0.891±0.000	0.224±0.000
	JailJudge [37]	**0.929±0.000**	0.673±0.000	0.857±0.000	0.561±0.002	0.924±0.000	0.497±0.000	0.619±0.000	0.835±0.000
	GCG [4]	0.929±0.000	0.570±0.000	0.904±0.000	0.702±0.004	0.786±0.000	0.653±0.000	**0.939±0.001**	0.551±0.000
	PAIR [46]	**0.953±0.000**	0.719±0.000	0.936±0.000	0.790±0.001	0.856±0.000	0.509±0.000	0.893±0.002	0.444±0.000
	PAP [70]	**0.944±0.002**	0.512±0.000	0.252±0.000	0.695±0.004	0.432±0.000	0.844±0.000	0.801±0.001	0.634±0.000
	SAA [5]	**0.961±0.001**	0.637±0.000	0.886±0.000	0.825±0.015	0.615±0.000	0.466±0.000	0.946±0.000	0.718±0.000
	TAP [71]	**0.956±0.000**	0.772±0.000	0.941±0.000	0.853±0.002	0.858±0.000	0.502±0.000	0.861±0.000	0.516±0.000
	Zulu [72]	**0.922±0.002**	0.545±0.000	0.733±0.000	0.448±0.007	0.921±0.000	0.890±0.000	0.876±0.000	0.745±0.000
	Avg.	**0.940±0.000**	0.647±0.000	0.813±0.000	0.720±0.000	0.776±0.000	0.633±0.000	0.848±0.000	0.558±0.000
Llama3-8B	AutoDAN [34]	0.961±0.000	0.342±0.000	**0.967±0.000**	0.696±0.000	0.722±0.000	0.598±0.000	0.945±0.000	0.952±0.000
	DrAttack [35]	**0.968±0.000**	0.387±0.000	0.388±0.000	0.840±0.000	0.530±0.000	0.656±0.000	0.706±0.000	0.663±0.000
	IJP [36]	**0.909±0.002**	0.366±0.000	0.376±0.000	0.739±0.000	0.501±0.000	0.601±0.000	0.888±0.000	0.651±0.000
	JailJudge [37]	**0.949±0.000**	0.688±0.000	0.731±0.000	0.678±0.000	0.675±0.000	0.584±0.000	0.618±0.000	0.715±0.000
	GCG [4]	**0.981±0.000**	0.368±0.000	0.608±0.000	0.651±0.000	0.610±0.000	0.627±0.000	0.945±0.000	0.970±0.000
	PAIR [46]	**0.940±0.000**	0.450±0.000	0.879±0.000	0.830±0.000	0.626±0.000	0.604±0.000	0.789±0.000	0.830±0.000
	PAP [70]	**0.891±0.000**	0.441±0.000	0.861±0.000	0.840±0.000	0.811±0.000	0.689±0.000	0.796±0.000	0.733±0.001
	SAA [5]	0.938±0.000	0.187±0.000	0.065±0.000	0.002±0.000	0.088±0.000	0.874±0.000	**0.946±0.001**	0.878±0.000
	TAP [71]	**0.952±0.000**	0.415±0.000	0.888±0.000	0.819±0.000	0.750±0.000	0.639±0.000	0.866±0.000	0.946±0.000
	Zulu [72]	**0.949±0.000**	0.673±0.000	0.559±0.000	0.216±0.000	0.524±0.000	0.379±0.000	0.876±0.000	0.806±0.000
	Avg.	**0.944±0.000**	0.432±0.000	0.632±0.000	0.631±0.000	0.584±0.000	0.625±0.000	0.837±0.000	0.814±0.000
Mistral-7B	AutoDAN [34]	**0.990±0.000**	0.820±0.000	0.431±0.000	0.717±0.000	0.858±0.000	0.505±0.000	0.945±0.000	0.357±0.000
	DrAttack [35]	**0.945±0.001**	0.566±0.000	0.383±0.000	0.918±0.000	0.926±0.000	0.591±0.000	0.698±0.000	0.304±0.000
	IJP [36]	**0.943±0.000**	0.609±0.000	0.337±0.000	0.931±0.000	0.337±0.000	0.760±0.000	0.888±0.000	0.609±0.000
	JailJudge [37]	**0.918±0.000**	0.481±0.000	0.450±0.000	0.546±0.000	0.883±0.000	0.592±0.000	0.630±0.000	0.770±0.000
	GCG [4]	0.976±0.000	0.500±0.000	0.876±0.000	0.738±0.000	**0.990±0.000**	0.853±0.000	0.946±0.000	0.659±0.000
	PAIR [46]	**0.977±0.000**	0.698±0.000	0.681±0.000	0.937±0.000	0.933±0.000	0.767±0.000	0.883±0.000	0.610±0.000
	PAP [70]	**0.924±0.000**	0.816±0.000	0.492±0.000	0.657±0.000	0.815±0.000	0.539±0.000	0.821±0.000	0.310±0.000
	SAA [5]	**0.960±0.000**	0.340±0.000	0.043±0.000	0.948±0.000	0.538±0.000	0.547±0.000	0.946±0.000	0.265±0.000
	TAP [71]	**0.955±0.000**	0.709±0.000	0.652±0.000	0.932±0.000	0.860±0.000	0.637±0.000	0.846±0.000	0.510±0.000
	Zulu [72]	0.938±0.001	0.456±0.000	0.390±0.000	0.760±0.000	0.753±0.000	0.562±0.000	0.876±0.000	**0.975±0.000**
	Avg.	**0.953±0.000**	0.600±0.000	0.473±0.000	0.808±0.000	0.789±0.000	0.635±0.000	0.848±0.000	0.537±0.000
Vicuna-7B	AutoDAN [34]	**0.957±0.002**	0.635±0.000	0.497±0.000	0.226±0.000	0.875±0.000	0.584±0.000	0.943±0.000	0.535±0.000
	DrAttack [35]	**0.956±0.001**	0.697±0.000	0.401±0.000	0.898±0.000	0.948±0.000	0.713±0.000	0.753±0.000	0.473±0.000
	IJP [36]	**0.937±0.000**	0.510±0.000	0.153±0.000	0.865±0.000	0.841±0.000	0.451±0.000	0.890±0.000	0.509±0.000
	JailJudge [37]	0.921±0.000	0.675±0.000	0.258±0.000	0.442±0.000	**0.924±0.000**	0.716±0.000	0.620±0.000	0.423±0.000
	GCG [4]	**0.957±0.000**	0.863±0.000	0.601±0.000	0.591±0.000	0.954±0.000	0.584±0.000	0.930±0.000	0.864±0.000
	PAIR [46]	**0.922±0.003**	0.668±0.000	0.507±0.000	0.866±0.000	0.913±0.000	0.545±0.000	0.913±0.000	0.571±0.000
	PAP [70]	**0.895±0.002**	0.585±0.000	0.475±0.000	0.752±0.001	0.836±0.000	0.443±0.000	0.798±0.000	0.490±0.000
	SAA [5]	0.835±0.005	0.430±0.000	0.012±0.000	**0.976±0.000**	0.799±0.000	0.458±0.000	0.948±0.001	0.354±0.000
	TAP [71]	**0.947±0.001**	0.712±0.000	0.683±0.000	0.880±0.000	0.917±0.000	0.578±0.000	0.878±0.000	0.661±0.001
	Zulu [72]	0.910±0.001	0.907±0.000	0.215±0.000	0.608±0.000	**0.943±0.000**	0.247±0.000	0.876±0.000	0.900±0.000
	Avg.	**0.923±0.001**	0.668±0.000	0.380±0.000	0.710±0.000	0.895±0.000	0.532±0.000	0.855±0.000	0.578±0.000

Table 4. AUROC comparison across Llama2-7B, Llama3-8B, Mistral-7B, and Vicuna-7B.

Robustness under MTK-targeted adaptive attacks

Across three surrogate losses and a full tradeoff sweep, MTK remains difficult to evade without sacrificing attack success.

λ	$\mathcal{J}_1$ Loss			$\mathcal{J}_2$ Loss			$\mathcal{J}_3$ Loss		
	ASR	TPR	eASR	ASR	TPR	eASR	ASR	TPR	eASR
0.1	0.62	0.80	0.12	0.75	0.75	0.19	0.72	0.76	0.17
0.2	0.67	0.76	0.16	0.30	0.89	0.03	0.62	0.95	0.03
0.3	0.64	1.00	0.00	0.00	1.00	0.00	0.71	0.95	0.04
0.4	0.67	0.90	0.07	0.00	1.00	0.00	0.65	0.76	0.16
0.5	0.62	0.81	0.12	0.05	1.00	0.00	0.67	0.95	0.03
0.6	0.58	0.95	0.03	0.00	1.00	0.00	0.62	0.71	0.18
0.7	0.57	0.86	0.08	0.00	1.00	0.00	0.62	0.81	0.12
0.8	0.38	0.52	0.18	0.00	1.00	0.00	0.54	0.81	0.10
0.9	0.21	0.29	0.15	0.00	1.00	0.00	0.57	0.76	0.14

Table 5. Full robustness sweep under three adaptive objectives.

Transfer to vision-language models and practical deployment

The same trajectory representation transfers to multimodal jailbreaks and adds only a lightweight detector on top of the model.

Table 6: **Evaluation and comparison.** AUROC of def against VLM jailbreak attacks. **Bold** denotes the best.

VLMs↓	Detector↓	MM-SafetyBench [79]	FigImg [80]	JailBreakV-28K [81]
LLaVa	HiddenDetect [9]	**0.968**	0.877	0.938
	ECSO [73]	0.848	0.761	0.866
	MirrorCheck [74]	0.688	0.707	0.611
	CIDER [75]	0.673	0.685	0.607
	JailGuard [23]	0.533	0.522	0.586
	JailDAM [17]	0.612	0.671	0.598
	MTK (**Ours**)	0.911	**0.959**	**0.949**
Qwen-VL	HiddenDetect [9]	0.774	0.980	0.881
	ECSO [73]	0.686	0.651	0.712
	MirrorCheck [74]	0.781	0.716	0.628
	CIDER [75]	0.673	0.685	0.607
	JailGuard [23]	0.750	0.855	0.823
	JailDAM [17]	0.519	0.634	0.627
	MTK (**Ours**)	**0.818**	**0.992**	**0.964**

Table 6. VLM jailbreak detection.

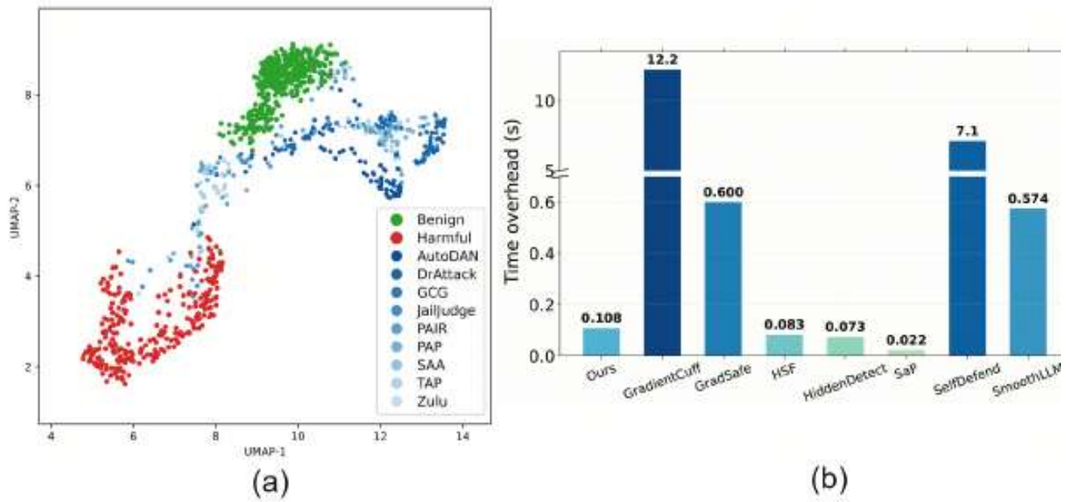

Figure 5. Rank-space structure and per-prompt overhead.

MTK requires no auxiliary generative safety model; the deployed component is an Isolation Forest over rank trajectories.

Takeaway

Reliable jailbreak detection benefits from modeling the transformation process, not a single representation snapshot.

Static feature separation breaks under both pseudo-malicious prompts and detector-aware attacks.

MTK captures layer-wise neighborhood kinetics with a compact rank trajectory.

The signal generalizes across LLMs, attack families, adaptive evasion, and vision-language models.

Code: github.com/Rookie143/mtk



