

PROACT-Agent

Progressive Runtime Oversight and Active Circuit-breaking for Real-Time Safety

Ding Jia¹ Wei Liu¹ Xianglong Du¹ Yingjie Li¹


Yingqing Yang¹ Huili Yu¹ Zhangsong Zhan¹ Chu Zhou^{2,*}

¹ State Key Laboratory of Intelligent Vehicle Safety Technology, Changan Automobile

² Independent Researcher * Corresponding author

How a benign task becomes unsafe


Illustrative prompt injection: a tool response introduces an attacker's instruction.



The guard must decide whether the updated trajectory is safe to continue.

PROACT guards the next model inference


The updated observation and full interaction history inform every decision.



Stops further unsafe continuation. Completed tool effects cannot be reversed.

The PROACT-Agent data pipeline

Completed logs become bilingual supervision for runtime intervention.



Learning the first unsafe boundary

Retain the safe-to-unsafe transition and exclude post-boundary label reversals.

Illustrative committee-adjudicated label sequence

Prefix	1	2	3 = t^*	4	5	6
Before rectification	SAFE	SAFE	UNSAFE	SAFE	UNSAFE	SAFE
Retained for training	SAFE	SAFE	UNSAFE	Excluded	Excluded	Excluded

Progressive unrolling


Turns each available prefix into a continue-or-stop training example.

Causal rectification

Inferred rationales and committee adjudication support boundary labels.

PROACT-Bench: scale, quality and coverage

155K+ bilingual trajectory prefixes across six risk categories and six tool domains.



DQS 99.3 reflects strong internal model-committee agreement on intervention labels.

Performance

Highest accuracy and unsafe-class F1 overall and in both languages.

97.0%

Overall unsafe-class F1

97.3%

English unsafe-class F1

96.7%

Chinese unsafe-class F1

Family	Model	Acc.	Pre.	Rec.	F1
Closed-source LLM	Gemini 3 Flash	88.7	84.3	94.2	89.0
Open-source LLM	Kimi-K2.5	91.1	87.5	95.1	91.1
Static guardrail	Qwen3Guard-Gen-8B	91.5	90.4	92.2	91.3
Agent guard	TS-Guard	90.5	91.5	88.5	90.0
Ours	PROACT-Agent	97.1	96.3	97.7	97.0

Overall metrics (%). Best-F1 baseline per family; the complete comparison is in Appendix A–B.

Generalization beyond the training examples

Separate protocols test unseen root trajectories and entirely unseen data sources.

Source-held-out

Train: AgentAlign + ToolSafety
Evaluate: R-Judge + ATBench
Tests transfer to new source datasets.

Root-grouped 80/20


Split complete canonical roots.
Keep derived prefixes and EN/ZH pairs together on the same side of the split.

PROACT-Agent protocol	Acc. (%)	Pre. (%)	Rec. (%)	F1 (%)
Source-held-out	92.4	95.5	87.7	91.5
Root-grouped 80/20	97.9	96.4	96.2	96.3

Unsafe is the positive class. Each protocol uses a separately trained checkpoint and its own test set.

Detecting the right intervention point

Exact detection preserves every earlier safe prefix and blocks the first unsafe prefix.



90.6%


**Exact-boundary detection
on entirely held-out sources**

Stops at the adjudicated boundary
while leaving earlier safe prefixes
unblocked.

Model-adjudicated boundaries; denominator: root-language trajectories with an unsafe boundary.

AgentDojo: protection during live tool use

Attack resistance (100% – targeted ASR). Higher is better.



Qwen3-32B (Q8_0) target agent + Qwen2.5-7B guard; 10K+ paired non-DoS attacks.

97.3–100% attack resistance across all four suites, including episodes stopped by the guard.

Clean-trajectory false-block rate: 7.2% (7/97), measured on a separate clean set.

Performance and runtime cost

Compact guards retain strong detection quality with sub-second median guard latency.

Guard backbone	Performance (%)			Latency (ms)		
	Acc.	Pre.	Rec.	F1	P50	P95
Qwen3-1.7B	96.1	95.3	96.7	96.0	403	889
Qwen3-4B	96.9	96.5	97.1	96.8	543	1,137
Qwen2.5-7B	97.1	96.3	97.7	97.0	474	976

A800-SXM4-80GB, batch 1, BF16, FlashAttention-2.

20 warm-up calls and 500 measured requests. P50 is the median; P95 is the 95th percentile.

Guard-only measurements. Model generation and parameter count both vary.

PROACT-Agent

Bilingual prefix supervision for timely runtime intervention.

Context before continuation

The full interaction history informs whether the next model inference should proceed.

An explicit stopping boundary

Causal rectification supplies supervision for the first unsafe prefix, with transfer to unseen data sources.



Read the paper

arxiv.org/abs/2609.34415

Contact

jiading.biz@outlook.com
zhou_chu@hotmail.com

Appendix A: Complete performance comparison

Overall accuracy and unsafe-class precision, recall and F1 (%).

Family	Model	Acc.	Pre.	Rec.	F1
Closed-source LLMs	Qwen3-Max	87.5	80.3	98.3	88.3
	Gemini 3 Flash	88.7	84.3	94.2	89.0
Open-source LLMs	DeepSeek-V3.2	86.2	79.8	95.7	87.0
	GLM-5.1	84.4	77.4	95.6	85.6
	Kimi-K2.5	91.1	87.5	95.1	91.1
	MiniMax-M2.5	85.6	84.6	85.8	85.2
Static LLM guardrails	Llama-Guard-3-8B	85.7	95.9	73.5	83.2
	Qwen3Guard-Gen-8B	91.5	90.4	92.2	91.3
Agent guard frameworks	ShieldAgent	68.2	60.6	97.3	74.7
	AgentDoG-Qwen3-4B	58.3	53.7	98.8	69.5
	AgentDoG-Qwen2.5-7B	60.6	60.3	53.8	56.8
	AgentDoG-Llama3.1-8B	61.5	55.6	99.7	71.4
	TS-Guard	90.5	91.5	88.5	90.0
	Safiron	61.3	56.2	89.5	69.0
Ours	PROACT-Agent	97.1	96.3	97.7	97.0

Broad comparison protocol (paper, Table 1). Results are rounded half-up to one decimal.

Appendix B: English and Chinese performance

The same four metrics are reported separately for both languages (%).

Model	English				Chinese			
	Acc.	Pre.	Rec.	F1	Acc.	Pre.	Rec.	F1
Qwen3-Max	87.0	78.7	98.6	87.5	88.0	81.8	97.9	89.1
Gemini 3 Flash	88.1	83.2	93.0	87.8	89.4	85.4	95.3	90.1
DeepSeek-V3.2	81.6	73.1	94.8	82.6	91.0	86.9	96.6	91.5
GLM-5.1	88.7	84.6	92.2	88.2	80.2	72.2	98.7	83.4
Kimi-K2.5	91.9	87.9	95.5	91.6	90.2	87.1	94.6	90.7
MiniMax-M2.5	87.7	84.5	89.8	87.0	83.6	84.8	82.2	83.5
Llama-Guard-3-8B	88.0	94.9	78.2	85.8	83.4	96.9	69.3	80.8
Qwen3Guard-Gen-8B	90.2	87.5	92.0	89.7	92.8	93.3	92.3	92.8
ShieldAgent	66.9	58.6	96.6	72.9	69.5	62.6	98.0	76.4
AgentDoG-Qwen3-4B	57.2	51.8	99.2	68.1	59.4	55.5	98.4	70.9
AgentDoG-Qwen2.5-7B	63.7	61.2	57.9	59.5	57.6	59.3	50.0	54.3
AgentDoG-Llama3.1-8B	61.0	54.2	99.7	70.2	62.0	57.0	99.6	72.6
TS-Guard	89.7	88.8	89.0	88.9	91.3	94.3	88.0	91.0
Safiron	62.6	56.8	78.8	66.0	59.9	55.7	99.4	71.4
PROACT-Agent	97.4	96.4	98.3	97.3	96.8	96.2	97.2	96.7

Appendix C: Two-source training sensitivity

Both checkpoints use the same fixed 20% ID test set and complete source-held-out OOD set.

Training roots	Fixed ID 20%			Held-out OOD		
	Macro-F1	FBR	FAR	Macro-F1	FBR	FAR
60%	98.1	1.0	3.1	90.6	3.7	15.8
80%	98.2	0.9	2.8	91.8	4.1	12.7

Training sources: AgentAlign + ToolSafety. Held-out sources: R-Judge + ATBench.

Checkpoints are selected using fixed-ID Macro-F1 only.

All metrics are percentages. FBR = false-block rate; FAR = false-allow rate (both lower is better).

Macro-F1 averages class-level F1 and differs from the unsafe-class F1 used in the main results.