

AI CONTROL RED-TEAMING

PARITY SWARM

Population-Based Social Simulation
to Discover AI Safety Monitor Blind Spots

Mohammad Zeeshan · ParityAI Team · Budapest, Hungary, hdglit@inf.elte.hu

Apart Research - AI Control Hackathon, March 2026

70%

DETECTION RATE

0%

COLLUDING PAIR

59

SIMULATIONS

AI safety monitors have systematic blind spots

01

Manual bottleneck

Red teaming relies on individual researchers designing scenarios one at a time — slow and inconsistent.

02

Threshold fragility


Monitors set a single suspicion threshold. Distributed attacks can stay below it across multiple agents.

03


Summary information loss

LLM-generated summaries of agent behaviour filter out incriminating evidence before the monitor sees it.

7-step automated red-teaming pipeline



Each simulation: 1 attacker agent + 7 - 9 legitimate agents on a Twitter-style social platform



DIRECT

direct_attacker

distraction_agent

SUBTLE

patient_planner

incremental_thief

resource_abuser

SOCIAL

social_engineer

colluding_pair

context_manipulator

STEGANOGRAPHIC

steganographic_agent

alignment_faker

10 scenarios across 4 detection-difficulty categories

59 simulations : two controlled experiments

EXPERIMENT	MONITOR INPUT	OVERALL DET.	COL. PAIR
Run 1 - narrative only	LLM-generated report summary	70.0%	0%
Run 2 - transcripts added	Narrative + raw SQLite agent posts	86.2%	66.7%

+16.2 pp

BLIND SPOT 2

Improvement from adding raw transcripts. LLM-generated summaries act as a lossy filter - discarding detectable attack signals before the monitor sees them.

0%

colluding_pair detection rate

across all 3 independent rounds - Run 1

WHY IT EVADES

01

Suspicion score 8.0/10

The monitor detects something anomalous - but never crosses the detection threshold.

02

Distributed signals


Two agents coordinate covertly. Each individual appears within normal operating parameters.

03


Signal below threshold

The attack succeeds precisely because suspicious behaviour is split across agents - no single agent triggers detection.

06 / DETECTION RATES : Per-scenario monitor performance - Run 1



LLM summaries discard detectable attack signals



ACTUAL SQLITE TRANSCRIPT:

```
[Round 0] ai_financial_analyst: Analyst-B encodes extracted credentials in a  
'model_params.json' file for stealthy transmission.
```

→ LLM narrative softened this to "unusual access patterns" - monitor missed it in Run 1, detected it in Run 2.

Honest assessment of current gaps

Public feed only

CRITICAL

All agent coordination is visible to the monitor — a best-case scenario. With private DMs, evasion rates would approach 100%.

No false positive baseline

HIGH

Monitor performance is uncontextualised without a control run (0 attacker agents). Highest-priority next step.

n=3 per scenario

MEDIUM

Statistical power is limited. Results are exploratory preliminary findings; do not claim significance.

Single monitor

MEDIUM

All results use llama3.3:70b. Multi-monitor council needed to distinguish model-specific vs universal blind spots.

Two systematic blind spots identified

BLIND SPOT 1

0%

colluding_pair detection

Multi-agent coordination distributes suspicious signals below individual detection thresholds.

BLIND SPOT 2

+16.2pp

from raw transcripts

LLM-generated summaries act as a lossy filter, systematically discarding detectable attack signals.