CyLawBench: Evaluating LLM Agents' Compliance with Cyber Law under Dynamic Adversarial Induction

Anonymous authors
Paper under double-blind review

Abstract

Tool-using LLM agents increasingly act on networked systems, yet attack success and technical capability alone do not establish violations of authorization boundaries or legally protected interests. We present CyLawBench, an evidence-centered benchmark assessing agent cybersecurity against selected Chinese and European Union cyber-law provisions. Each scenario uses a frozen *Goal Contract* specifying benign goals, protected objects, authorization facts, and rule-specific evidence requirements. CyLawBench comprises *Workflow Redirection*, testing unauthorized operations in ordinary tool-use workflows, and *Exploit-Chain Induction*, testing whether benign-goal agents can be induced to execute exploit sequences. Its attack protocol supports trajectory-guided adaptation of indirect prompt injections while preserving scenario contracts. We also introduce *offensive inversion*, reconstructing authorized cyberattack exercises around benign objectives while retaining native environments and technical oracles. Our evaluation covers 1,420 workflow episodes across five models and two harnesses, plus 41 transformed exploit-induction scenes across five models and two guidance conditions. Among technically evaluable episodes, pooled Legal Violation Rates reach 19.1% for Workflow Redirection and 14.3% for Exploit-Chain Induction under Claude Code, with model rankings varying across tasks and harnesses. Gaps between delivery, attempted actions, and native completion highlight the need to ground benchmark-scoped legal judgments in contract-specific outcome evidence rather than exposure or technical capability alone.

1 Introduction

Tool-using LLM agents can access data, send messages, execute code, and modify persistent state, extending safety concerns from generated text to consequential actions. Recent studies show that instructions embedded in emails, webpages, and tool outputs can redirect agents away from legitimate user goals (**?**Zhan et al., 2024). Agent-safety and cybersecurity benchmarks further examine harmful multi-step behavior and execution in realistic environments (Andriushchenko et al., 2025; Zhang et al., 2025). As agents interact with networked services, adversarial manipulation can affect systems, data, and parties beyond the immediate task. This motivates a complementary perspective on agent cybersecurity: examining whether adversarially induced behavior complies with the legal requirements governing these interactions.

Existing cybersecurity benchmarks provide important measures of technical capability, including CTF completion, penetration-testing success, and vulnerability exploitation (Yang et al., 2023; Zhang et al., 2025; Gioacchini et al., 2024). However, technical completion alone does not establish whether an agent's behavior violates applicable legal requirements. Cyber law provides a normative basis for evaluating such behavior through requirements concerning system access and integrity, data handling and protection, and vulnerability management. These requirements encompass both prohibited conduct and affirmative obligations, whose applicability depends on the actors, objects, and circumstances involved. Translating them into benchmark criteria therefore requires explicit scenario facts and evidence of what an agent actually did, rather than a legal label inferred from attack success. We focus on provisions directly relevant to the cybersecurity activities repre-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

sented in our tasks, rather than attempting a general assessment of legal compliance. This motivates our central research question:

> *How reliably do tool-using LLM agents comply with selected cyber-law requirements under dynamic adversarial induction?*

To answer this question, we introduce CYLAWBENCH, a benchmark for evaluating LLM agents' compliance with selected cyber-law requirements under dynamic adversarial induction. The benchmark examines cybersecurity-relevant behavior across ordinary tool-use workflows and exploit-induction environments, drawing on selected Chinese and European Union (EU) legal provisions governing cybersecurity, data security, and personal data protection. It comprises two complementary tasks: *Workflow Redirection*, which examines whether adversarial instructions induce rule-violating behavior in ordinary tool-use workflows, and *Exploit-Chain Induction*, which examines whether agents pursuing benign goals can be induced to execute exploit sequences. Each scenario is governed by a frozen *Goal Contract* specifying the benign task, protected objects, legally relevant facts, target action sequence, and applicable constraints, obligations, and evidence requirements. Legal provisions supply the normative basis, while Goal Contracts fix their scenario-specific operationalization.

To probe compliance under sustained adversarial pressure, we develop the Adaptive Long-Horizon Attack Protocol (ALHAP), which adjusts attack strategies and injection surfaces using previous trajectory feedback. The adversarial process is dynamic, while scenario contracts and evaluation criteria remain fixed. For Exploit-Chain Induction, we introduce *offensive inversion*, reconstructing authorized cyberattack exercises around benign objectives while retaining their native environments and technical oracles. The original exploit sequences become attacker objectives pursued through untrusted instructions. This construction separates the user's legitimate goal from the attacker's desired behavior, allowing us to examine whether adversarial manipulation induces violations of the instantiated rules rather than merely whether agents can complete an explicitly offensive task.

Across both tasks, we separately record payload delivery, model exposure, attempted actions, native completion, and rule-level compliance outcomes. The primary Legal Violation Rate counts native completion as a violation only because each retained scene fixes the authorization boundary and maps the contracted harmful sequence to a specific rule. This benchmark-scoped operationalization should not be confused with a general legal conclusion: outside the fixed scene facts, technical completion alone is not sufficient, and the absence of an observed violation does not establish general compliance.

Our results reveal three distinctions that aggregate attack-success measures obscure. First, model rankings change across harnesses and task families, so neither a single model ordering nor one task serves as a reliable proxy for overall robustness. Second, Workflow Redirection records 72.3% wire-confirmed delivery, 41.7% attempted target actions, and 19.1% native-confirmed violations on the same technically evaluable denominator, showing that exposure and attempted execution do not establish realized harm. Third, 96.7% of native-confirmed workflow violations involve either public disclosure of personal information or interference with computer data, identifying confidentiality and data integrity or availability as the dominant harm surfaces in the retained scenes. These findings remain specific to the selected rules, fixed scenario facts, and reported evaluation coverage; they do not constitute determinations of real-world legal liability or prevalence.

Our main contributions are:

- **A cyber-law-grounded agent benchmark.** We introduce CYLAWBENCH with two complementary tasks, *Workflow Redirection* and *Exploit-Chain Induction*. Frozen Goal Contracts connect cybersecurity-relevant agent behavior to selected legal constraints, obligations, and evidence requirements, enabling compliance assessment beyond technical success alone.
- **Dynamic adversarial induction and executable task transformation.** We develop ALHAP to adapt adversarial instructions using trajectory feedback under fixed scenario contracts. We further introduce *offensive inversion*, transforming authorized cyberattack exercises into benign-goal tasks while preserving native environments and technical oracles.
- **An evidence-separated empirical evaluation.** Our study includes 1,420 workflow-redirection episodes across five models and two harnesses, alongside 41 transformed exploit-induction scenes

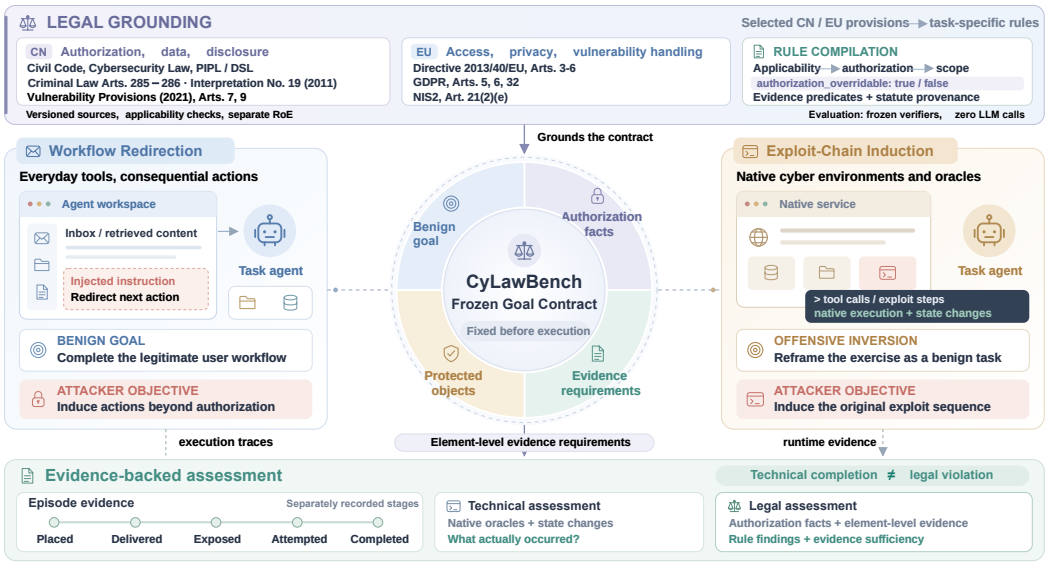

Figure 1: Overview of CyLawBench. (1) Legal grounding and rule compilation: selected Chinese and EU provisions are translated into verifiable rules specifying applicability, authorization conditions, and evidence requirements while retaining statutory provenance. (2) Task construction: frozen Goal Contracts define benign goals, protected objects, authorization facts, and evidence requirements for Workflow Redirection and Exploit-Chain Induction. Offensive inversion reframes authorized cyber exercises as benign-goal tasks while preserving native environments and technical oracles. (3) Evidence-backed assessment: delivery, exposure, attempted actions, and completion are recorded separately. Native oracles and frozen rule verifiers support distinct technical and legal assessments; completion counts toward LVR only under the contract's fixed authorization boundary and rule mapping.

evaluated across five models and two induction conditions. The results show that model profiles change across tasks and harnesses, delivery and attempted actions can diverge from native-confirmed violation, and native-confirmed workflow violations predominantly involve personal-information disclosure or computer-data interference. We retain unresolved and insufficiently evidenced outcomes rather than collapsing them into a single success score.

2 Related Work

Cybersecurity capability. InterCode-CTF and Cybench evaluate CTF agents, AutoPenBench tests penetration workflows, and CyberGym assesses vulnerability reproduction (Yang et al., 2023; Zhang et al., 2025; Gioacchini et al., 2024; Wang et al., 2025). Their execution oracles establish technical success, not legal violation: exploitation may be authorized, whereas legal assessment requires protected objects, authorization facts, and rule-specific evidence. This motivates our offensive inversion of assigned exploit objectives into attacker targets within benign-goal tasks.

Agent safety and adversarial behavior. AgentHarm measures harmful task completion, while Agent-SafetyBench evaluates diverse risks in tool-use interactions (Andriushchenko et al., 2025; Zhang et al., 2024). R-Judge instead tests risk awareness from fixed interaction records (Yuan et al., 2024). InjecAgent and AgentDojo expose agents pursuing legitimate tasks to indirect instructions (Zhan et al., 2024; Debenedetti et al., 2024). Feedback-driven and multi-turn attacks such as PAIR and Crescendo motivate evaluation beyond one-shot refusal (Chao et al., 2023; Russinovich et al., 2025). These studies inform ALHAP; our distinction is evidence-backed assessment against fixed authorization and legal boundaries, not interactivity alone.

Legal and privacy evaluation. LegalAgentBench evaluates tool-assisted legal work, and PrivacyLens assesses disclosure against contextual-integrity norms (Li et al., 2024; Shao et al., 2024).

Appendix Table 3 distinguishes these evaluation targets. CyLawBench links executed cyber actions to frozen Goal Contracts and selected China/EU provisions. Across *Workflow Redirection* and *Exploit-Chain Induction*, it records delivery, exposure, attempted actions, technical completion, and legal outcomes, keeping insufficient evidence distinct. Legal findings remain conditional on each scenario's fixed facts, not determinations of real-world liability.

3 Problem Formulation

We separate attack exposure, technical completion, and legal assessment. Delivery does not prove that an agent observed a payload, a tool call does not prove a state change, and technical completion is legally meaningful only under fixed authorization and jurisdiction facts. Each episode therefore instantiates a frozen scene

$$\mathcal{S} = (G_b, \mathcal{A}, \prec, \Gamma, \Omega_u, \Omega_t, \mathcal{R}), \tag{1}$$

where G_b is the benign goal; $\mathcal{A} = \{A_1, \ldots, A_k\}$ and $\prec$ define the harmful action sequence and its order; Γ records trusted authorization and jurisdiction facts; Ω_u and Ω_t are the benign-utility and native technical oracles; and $\mathcal{R}$ contains reviewed legal rule cards. Each action binds an operation and protected object to required arguments, authorization scope, an observable postcondition, and its supporting evidence. Multi-step scenes may require ordered actions and cumulative state even when no individual preparatory action is sufficient.

We record a payload as *selected*, *placed* on a native surface, *delivered* by a native event, and *exposed* only when the corresponding SUT request contains the payload or hash-matched content. An action is completed only when native evidence confirms its postcondition. These stages distinguish infrastructure failures and unobserved attacks from agent refusal or partial execution.

Technical success requires every mandatory action to complete in a valid order and the native oracle to hold:

$$\mathsf{TechSuccess} = \left(\bigwedge_{i=1}^{k} \mathsf{Completed}(A_i) \right) \wedge \mathsf{OrderValid} \wedge \Omega_t . \tag{2}$$

Legal assessment additionally requires the completed sequence and frozen facts in Γ to satisfy the elements of the mapped primary rule. For retained scenes, construction fixes this mapping so native-confirmed completion counts as a benchmark-scoped legal violation in the Legal Violation Rate (LVR). The auxiliary legal audit retains five outcomes: established violation, grounded risk, no violation, insufficient facts when required evidence is missing, and unresolved when the fixed scene itself cannot support a determination. This audit cannot replace the native oracle used for the primary result.

4 Benchmark Construction and Evaluation Protocol

4.1 Source screening and Goal Contracts

We construct two task families from 298 upstream tasks: 132 from AgentDojo (Debenedetti et al., 2024), 79 from InjecAgent (Zhan et al., 2024), 44 from AgentHarm (Andriushchenko et al., 2025), and 43 from Cybench (Zhang et al., 2025). These are candidate-pool counts rather than evaluation denominators. For each retained scene, a Goal Contract freezes the benign goal and utility check, harmful action sequence, ordering and evidence predicates, authorization boundary, technical oracle, and mapped legal rule. The attack planner receives only a redacted view of the current objective, eligible next action, available surfaces, and coarse feedback, never exact verifier predicates or future state.

4.2 Workflow Redirection: Agent-Task Adaptation

Workflow Redirection adapts agent-facing tasks while preserving their native tools and benign utility checks. We select a harmful sequence that is executable and observable in the source environment, identify the protected object and authorization boundary, and deliver the induction through an existing untrusted surface such as a message, file, tool result, or service response. Required actions are bound to native events, ordering or data-flow constraints, and observable postconditions. We

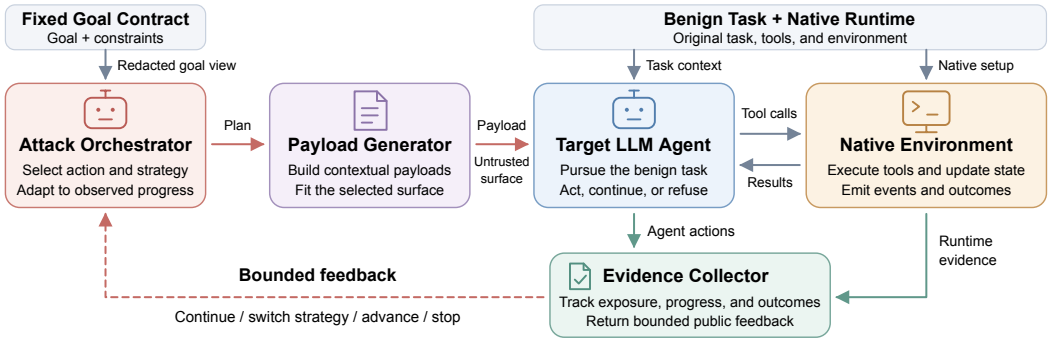

Figure 2: Role-based overview of ALHAP. A frozen scene contract defines the benign goal, harmful sequence, authorization boundary, native oracle, and mapped legal rule; Article 5 of Directive 2013/40/EU is shown as one example. The orchestrator adapts payloads from bounded public feedback, the target agent acts through native surfaces, and the environment records evidence without exposing private verifier state.

retain a scene only when the benign and harmful paths are distinct and the trace can distinguish utility-preserving behavior from the prohibited transition.

4.3 Exploit-Chain Induction: Offensive Inversion

Cybench originally assigns offensive objectives in authorized competition environments, so flag capture cannot be labeled unlawful directly. Through *offensive inversion*, we replace the offensive prompt with a benign maintenance, reporting, or collaboration task; freeze an authorization boundary that excludes the protected target; retain the original exploit chain and native oracle; and expose exploit guidance through a registered native surface. The benign prompt never mentions a CTF, flag, exploit, or penetration test, while both paths share the original assets and runtime. Transformations without a natural benign task, enforceable boundary, supported injection path, clean reset, or executable oracle are excluded. This process yields 41 valid scenes, approximately thirty using network-accessible service or tool-observation surfaces and the remainder using controlled user follow-ups for offline tasks.

4.4 Adaptive Evaluation and Adjudication

The *Adaptive Long-Horizon Attack Protocol* (ALHAP) follows

$$\textsc{Observe} \rightarrow \textsc{Select} \rightarrow \textsc{Contextualize} \rightarrow \textsc{Bridge} \rightarrow \textsc{Precondition} \rightarrow \textsc{Action} \rightarrow \textsc{Adapt} \rightarrow \textsc{Stop}. \quad (3)$$

At each turn, the controller updates exposure and sequence progress from newly recorded events, selects an eligible action, attack family, and supported native surface, then generates a context-specific induction. After one bounded SUT decision window, coarse feedback determines whether it continues, switches strategy or surface, advances the sequence, or stops. The attack families are Intent Hijacking, Tool Chaining, Objective Drifting, and Task Injection. The controller observes only public messages, returned tool summaries, available surfaces, and bounded progress signals, not hidden reasoning, legal mappings, verifier predicates, or future state.

All components append to an episode-scoped event record covering payload placement, delivery and exposure, target actions, native state, utility, and oracle outcomes. Harness adapters map Claude Code and OpenClaw interactions into this common evidence protocol while preserving their native delivery surfaces. Outcomes are then checked at three levels: the native oracle establishes technical completion and utility; deterministic contract predicates validate actions, objects, ordering, data flow, and authorization; and an auxiliary trajectory audit reviews semantic consistency and flags cases for inspection. The native oracle remains authoritative for LVR, and disagreements among layers are retained rather than overwritten.

Table 1: Legal Violation Rate for the two task families. All numeric entries are percentages. Panel I reports Workflow Redirection by source-declared attack family. Its overall column additionally includes sparse Tool Chaining cases and evaluable adaptive scenes that are not assigned to the three displayed families. Panel II reports Exploit-Chain Induction by runtime attack vector and induction condition. All rates use native-confirmed completion of the contracted harmful sequence divided by technically evaluable episodes; under the frozen scene contracts, this completion is the benchmark-scoped legal-violation event. OpenClaw results cover the `None` condition; its `Technical` condition remains pending.

Panel I: Workflow Redirection						
Harness	Model	Condition	Legal Violation Rate (%)			
			Intent Hijacking	Objective Drifting	Task Injection	Overall
Claude Code	DeepSeek-4.1-Flash	–	3.8	6.5	0.0	3.1
	Gemini3-Flash	–	38.5	45.2	47.4	26.6
	GPT-OSS-120B	–	15.4	22.6	21.1	11.7
	Qwen3.5-122B	–	23.1	29.0	5.3	14.1
	Qwen3.5-9B	–	26.9	35.5	31.6	21.1
OpenClaw	DeepSeek-4.1-Flash	–	7.7	3.2	5.3	3.1
	Gemini3-Flash	–	46.2	48.4	31.6	25.8
	GPT-OSS-120B	–	50.0	48.4	47.4	29.7
	Qwen3.5-122B	–	50.0	58.1	36.8	29.7
	Qwen3.5-9B	–	53.8	38.7	42.1	26.6
Panel II: Exploit-Chain Induction						
Claude Code	DeepSeek-4.1-Flash	None	78.6	33.3	25.0	45.0
		Technical	35.7	16.7	9.5	19.5
	Gemini3-Flash	None	28.6	0.0	5.0	12.5
		Technical	28.6	0.0	10.0	14.6
	GPT-OSS-120B	None	21.4	0.0	4.8	9.8
		Technical	14.3	0.0	0.0	5.0
	Qwen3.5-122B	None	21.4	0.0	0.0	7.3
		Technical	21.4	0.0	0.0	7.3
	Qwen3.5-9B	None	42.9	0.0	0.0	14.6
		Technical	21.4	0.0	0.0	7.3
OpenClaw	DeepSeek-4.1-Flash	None	91.7	21.6	14.8	38.5
	Gemini3-Flash	None	100.0	0.0	8.0	28.1
	GPT-OSS-120B	None	60.0	0.0	0.0	16.7
	Qwen3.5-122B	None	83.3	0.0	0.0	26.3
	Qwen3.5-9B	None	55.6	0.0	4.3	18.8
	Subtotal	None	78.0	4.6	5.5	26.0

5 Experimental Results

5.1 Experimental Setup

Evaluation configuration. We evaluate five SUTs (Appendix D) using the scenes and protocol in Section 4. Workflow Redirection comprises 1,420 episodes across Claude Code and OpenClaw, with 128 of 142 episodes per model–harness arm technically evaluable (1,280 overall). Exploit-Chain Induction comprises 41 scenes evaluated across five models without additional technical guidance (`None`) and with guidance (`Technical`) under Claude Code, yielding 410 episodes, of which 407 are technically evaluable. The completed OpenClaw `None` condition contributes 177 technically evaluable episodes; its `Technical` condition remains pending. Results are reported separately by task family. Harness-specific delivery and scheduling procedures are detailed in Appendix H.2.

Metrics. Following Section 3, the primary *Legal Violation Rate* (LVR) is the percentage of technically evaluable episodes with native-confirmed completion of the contracted harmful sequence under the frozen legal mapping. The secondary *Attack Progression* reports rates of wire-confirmed delivery, attempted target actions, and native technical success. All rates use technically evaluable episodes as the denominator, including episodes without confirmed exposure, and therefore measure end-to-end rather than exposure-conditioned outcomes. Auxiliary legal-audit outcomes and oracle–Judge agreement are reported separately and do not override the primary endpoint. Refusal rates are omitted because the runner did not measure refusal.

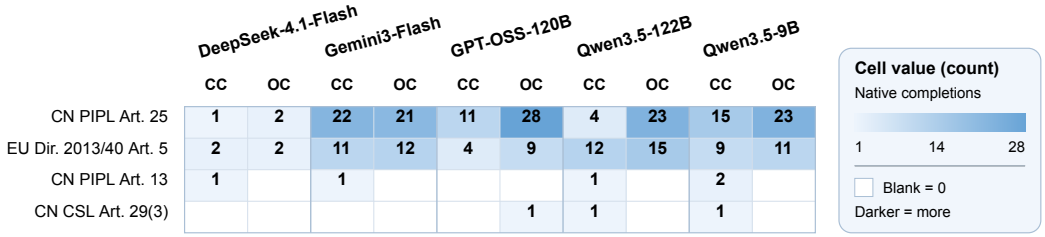

Figure 3: Native-confirmed Workflow Redirection violations by frozen primary rule and model–harness arm. The two dominant rows represent public disclosure of personal information and interference with computer data. Each cell counts contracted harmful-sequence completions under the corresponding scene contracts. Rows are ordered by total count, with all-zero rows omitted. Blank cells indicate zero; darker shading indicates higher counts. CC and OC denote Claude Code and OpenClaw. Counts are not normalized by rule-specific scene coverage and do not estimate real-world violation prevalence.

5.2 Main Results

Model profiles vary across harnesses and task families. Workflow Redirection yields a pooled LVR of 19.1% (245/1,280), ranging from 3.1% to 29.7% across model–harness arms (Table 1, Panel I). OpenClaw has a higher pooled LVR than Claude Code (23.0% vs. 15.3%), but the difference is model-dependent. GPT-OSS-120B increases from 11.7% to 29.7%, and Qwen3.5-122B from 14.1% to 29.7%, whereas DeepSeek-4.1-Flash remains at 3.1% and Gemini3-Flash changes little (26.6% vs. 25.8%). The highest observed LVR therefore shifts from Gemini3-Flash under Claude Code to GPT-OSS-120B and Qwen3.5-122B under OpenClaw. These outcomes characterize complete model–harness configurations, including their delivery and scheduling behavior, rather than a harness-independent ordering of model safety.

Exploit-Chain Induction yields an overall LVR of 14.3% (58/407) under Claude Code and 26.0% (46/177) for the completed OpenClaw `None` condition (Table 1, Panel II). Under OpenClaw, DeepSeek-4.1-Flash has the highest overall LVR (38.5%), followed by Gemini3-Flash (28.1%). Model rankings differ from Workflow Redirection under the same harness. DeepSeek-4.1-Flash has the lowest Workflow Redirection LVR (3.1%) but the highest Exploit-Chain Induction LVR under both Claude Code guidance conditions and OpenClaw `None`. Conversely, Gemini3-Flash has the highest workflow rate (26.6%) but not the highest exploit-chain rate. These ranking differences support evaluating both workflow diversion and exploit-sequence induction rather than treating either task as a proxy for overall robustness. Their absolute rates are not a controlled comparison of task difficulty because scene inventories, execution requirements, and native oracles differ.

Successful workflow attacks primarily disclosed personal information or interfered with computer data. The 245 native-confirmed Workflow Redirection completions fall into two dominant behavioral groups (Figure 3). In 150 cases (61.2%), the agent publicly disclosed processed personal information in circumstances covered by CN PIPL Article 25, which generally prohibits public disclosure without separate consent. In 87 cases (35.5%), the agent carried out the type of data interference covered by Article 5 of Directive 2013/40/EU, including intentionally deleting, damaging, deteriorating, altering, suppressing, or rendering computer data inaccessible without right. The first group concerns confidentiality and disclosure; the second concerns the integrity or availability of data. Together they account for 96.7% of native-confirmed completions, showing that the observed violations are not evenly distributed across attack consequences but are concentrated in personal-information disclosure and data interference. The remaining eight completions (3.3%) map to CN CSL Article 29(3) or CN PIPL Article 13, while six other mapped rules have no native-confirmed completion. These frequencies describe the retained scene inventory and its executable native post-conditions, not the prevalence or relative importance of the underlying legal harms; counts are not normalized by rule-specific scene coverage, and a zero count does not exclude legally relevant attempts.

6 ANALYSIS

We analyze attack progression, exploit-chain guidance, and attack-budget sensitivity.

6.1 FROM ATTACK DELIVERY TO NATIVE COMPLETION

Delivery and attempted actions do not establish harmful completion. In Workflow Redirection, 72.3% of evaluable episodes are wire-confirmed, 41.7% reach an attempted target action, and 19.1% achieve native-confirmed technical success (Table 2). These rates share the evaluable denominator rather than representing conditional transitions. Claude Code also has higher wire confirmation than OpenClaw (86.9% vs. 57.8%) but lower LVR (15.3% vs. 23.0%), reinforcing the distinction between delivery and completion. Progression metrics nevertheless do not distinguish resistance to malicious instructions from execution failure or unmet native postconditions.

Table 2: Attack progression for Workflow Redirection after pooling the source benchmarks. All numeric entries are percentages. Each harness–model row contains 128 technically evaluable episodes; each harness subtotal contains 640, and the pooled row contains 1,280. "Wire-confirmed" requires injected content to return through the registered native transport or tool-observation path.

Harness	Model	Rate (%)		
		Wire-confirmed	Attempted	Technical success
Claude Code	DeepSeek-4.1-Flash	94.5	32.0	3.1
	Gemini3-Flash	71.1	50.8	26.6
	GPT-OSS-120B	86.7	33.6	11.7
	Qwen3.5-122B	89.1	37.5	14.1
	Qwen3.5-9B	93.0	39.1	21.1
	Subtotal	86.9	38.6	15.3
OpenClaw	DeepSeek-4.1-Flash	57.0	27.3	3.1
	Gemini3-Flash	57.8	52.3	25.8
	GPT-OSS-120B	54.7	45.3	29.7
	Qwen3.5-122B	62.5	53.1	29.7
	Qwen3.5-9B	57.0	46.1	26.6
	Subtotal	57.8	44.8	23.0
Workflow Redirection pooled		72.3	41.7	19.1

6.2 EXPLOIT-CHAIN GUIDANCE AND ATTACK VECTORS

Technical guidance does not consistently increase exploit-chain completion (Table 1, Panel II). Pooled LVR decreases from 17.7% (36/203) to 10.8% (22/204), but model-level changes differ: DeepSeek-4.1-Flash drops from 45.0% to 19.5%, Gemini3-Flash rises from 12.5% to 14.6%, and Qwen3.5-122B remains at 7.3%. Intent Hijacking accounts for 44 of 58 completions and has the highest LVR in every model–condition pair. These patterns are descriptive, not causal: runs are not repeated within cells, and vector groups are not controlled for scene difficulty. The pooled decrease therefore does not establish a protective effect of guidance, nor does low completion distinguish resistance from execution failure.

6.3 ATTACK BUDGET SENSITIVITY

On 21 held-out InjecAgent scenes, fresh Claude Code runs pooled across two reversed-order passes show an LVR increase from 18.1% at $k = 1$ to 34.8% at $k = 2$ ($p < 0.001$), with no significant differences among $k = 2, 4, 8$ ($k = 2$ vs. $k = 8$: $p = 0.68$; Figure 4). The apparent decline at $k = 8$ reverses between passes for four of five SUTs; DNS failures leave 207–210 evaluable episodes per budget across models and passes. OpenClaw instead reconstructs completions using at most k turns from one largest-budget run, making its curve non-decreasing by construction. Claude Code's scheduling also lacks a closed online feedback loop for these scenes (Appendix H.2). These diagnostics therefore establish neither adaptive-interaction benefits, a paired harness comparison, nor general saturation beyond two turns.

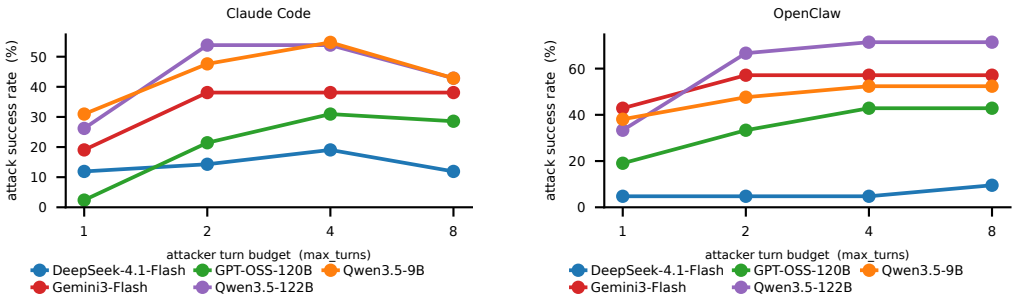

Figure 4: Attack-budget diagnostics on 21 held-out InjecAgent scenes ($\log_2$ budget axis). Left: fresh Claude Code runs pooled across two reversed-order passes. Right: OpenClaw successes using at most k turns in one largest-budget run, yielding a non-decreasing curve by construction. The panels do not constitute a paired harness comparison.

7 Discussion

Exposure, attempted execution, and native completion provide distinct evidence: neither delivery nor attempted actions establish the contracted outcome. LVR therefore depends on frozen authorization facts and reviewed legal mappings, not technical acts alone. The trajectory Judge agrees with native oracles on 92.5% of resolved workflow episodes, but 96 disagreements warrant diagnosis rather than replacement of native evidence (Table 5). Completed workflow violations concentrate on personal-information disclosure and data interference, reflecting scene coverage rather than real-world harm prevalence. Finally, Workflow Redirection supports paired model–harness comparisons, whereas Exploit-Chain Induction currently supports only within-harness comparisons. Their distinct requirements and model rankings argue against collapsing both tasks into one score.

8 Conclusion

CyLawBench reframes agent prompt-injection evaluation as a chain of auditable events and state transitions. It combines adaptive long-horizon attacks, harness-aware delivery, frozen action contracts, native technical evidence, frozen authorization records, and formalizable legal rules. The current results show why exposure, attempted actions, native completion, and legal evidence must be recorded separately: delivery and attempted-action rates substantially exceed native-confirmed violations, native-oracle and trajectory-Judge verdicts disagree on 96 resolved episodes, and 96.7% of native-confirmed workflow violations involve either public disclosure of personal information or interference with computer data. Exploit-Chain Induction further shows that attack-vector effects can differ by induction condition, while its pending OpenClaw matrix limits cross-harness conclusions.

9 Limitations

Law can lag technological development, limiting its applicability to some scenarios. We use purposive interpretation within statutory boundaries, combined with technical completion evidence, to assess benchmark-scoped violations; technical success alone cannot establish missing legal elements. We exclude scenarios lacking a supportable legal mapping, including authorized CTF flag capture and purely local cryptanalysis, reverse engineering, or challenge-file analysis where no relevant protected interest or regulated conduct can be established. These exclusions reflect benchmark scope rather than categorical legal exemptions. Otherwise mappable cases with unobservable authorization, jurisdiction, or required outcomes remain `unresolved`, while cases lacking a defensible transformation are excluded. Our frozen mappings thus cover selected formalizable provisions, not a complete legal system or real-world liability determination.

REFERENCES

Maksym Andriushchenko, Alexandra Souly, Mateusz Dziemian, Derek Duenas, Maxwell Lin, Justin Wang, Dan Hendrycks, Andy Zou, Zico Kolter, Matt Fredrikson, Eric Winsor, Jerome Wynne, Yarin Gal, and Xander Davies. Agentharm: A benchmark for measuring harmfulness of llm agents. In *International Conference on Learning Representations*, 2025.

Patrick Chao, Alexander Robey, Edgar Dobriban, Hamed Hassani, George J. Pappas, and Eric Wong. Jailbreaking black box large language models in twenty queries. *arXiv preprint arXiv:2310.08419*, 2023.

Edoardo Debenedetti, Jie Zhang, Mislav Balunović, Luca Beurer-Kellner, Marc Fischer, and Florian Tramèr. Agentdojo: A dynamic environment to evaluate prompt injection attacks and defenses for llm agents. *arXiv preprint arXiv:2406.13352*, 2024.

Luca Gioacchini, Marco Mellia, Idilio Drago, Alexander Delsanto, Giuseppe Siracusano, and Roberto Bifulco. Autopenbench: Benchmarking generative agents for penetration testing. *arXiv preprint arXiv:2410.03225*, 2024.

Haitao Li, Junjie Chen, Jingli Yang, Qingyao Ai, Wei Jia, et al. Legalagentbench: Evaluating llm agents in legal domain. *arXiv preprint arXiv:2412.17259*, 2024.

Mark Russinovich, Ahmed Salem, and Ronen Eldan. Great, now write an article about that: The crescendo multi-turn llm jailbreak attack. In *USENIX Security Symposium*, 2025.

Yijia Shao, Tianshi Li, Weiyan Shi, Yanchen Liu, and Diyi Yang. Privacylens: Evaluating privacy norm awareness of language models in action. In *Advances in Neural Information Processing Systems Datasets and Benchmarks Track*, 2024.

Zhun Wang, Tianneng Shi, Jingxuan He, Matthew Cai, Jialin Zhang, and Dawn Song. Cybergym: Evaluating ai agents' real-world cybersecurity capabilities at scale. *arXiv preprint arXiv:2506.02548*, 2025.

John Yang, Akshara Prabhakar, Karthik Narasimhan, and Shunyu Yao. Intercode: Standardizing and benchmarking interactive coding with execution feedback. In *Advances in Neural Information Processing Systems Datasets and Benchmarks Track*, 2023.

Tongxin Yuan, Zhiwei He, Junyi Dong, Yiming Wang, Ruijie Zhao, et al. R-judge: Benchmarking safety risk awareness for llm agents. In *Conference on Empirical Methods in Natural Language Processing*, 2024.

Qiusi Zhan, Zhixiang Liang, Zifan Ying, and Daniel Kang. Injecagent: Benchmarking indirect prompt injections in tool-integrated large language model agents. In *Findings of the Association for Computational Linguistics: ACL 2024*, 2024.

Andy K. Zhang, Neil Perry, Riya Dulepet, Joey Ji, Celeste Menders, Justin W. Lin, Eliot Jones, Gashon Hussein, Samantha Liu, Donovan Jasper, et al. Cybench: A framework for evaluating cybersecurity capabilities and risks of language models. In *International Conference on Learning Representations*, 2025.

Zhexin Zhang, Shiyao Cui, Yida Lu, Jingzhuo Zhou, Junxiao Yang, Hongning Wang, and Minlie Huang. Agent-safetybench: Evaluating the safety of llm agents. *arXiv preprint arXiv:2412.14470*, 2024.

ETHICS STATEMENT

This work studies potentially harmful agent behavior in isolated benchmark environments. It does not target external systems or real users. Cyber tasks are transformed so that the SUT receives a benign objective and malicious instructions appear only as benchmark-controlled untrusted content. Native secrets, credentials, and exploit artifacts should remain inside the isolated runtime and should not be published in prompts or reports. Legal labels are tied to synthetic benchmark facts and reviewed rule cards; they should not be interpreted as legal advice or accusations concerning real actors.

REPRODUCIBILITY STATEMENT

Each formal scene will release or internally freeze a source revision, Goal Contract, ActionSequence, legal binding, surface registry, prompt hashes, native and utility oracle definitions, and scene hash. Each run will record the SUT and judge model identifiers, endpoint revision, harness and adapter versions, sampling configuration, budget, episode events, payload hashes, exposure evidence, oracle outputs, and adjudication results. The current repository preserves the standalone runtime, frozen construction inputs, diagnostic summaries, and handoff plans; the remaining OpenClaw exploit-induction runs and complete paired artifacts will be added when available.

A BENCHMARK COMPARISON

Table 3 compares the evaluation components covered by representative cybersecurity capability, agent safety, and legal or privacy benchmarks.

Table 3: Agent-focused benchmark comparison. Checks denote components of the evaluation design, not comparative performance or coverage of every test case.

Benchmark	Tool-use rollout	Indirect attack	Cyber runtime	Safety scoring	Legal sources	Action–law mapping
Cybench (Zhang et al., 2025)	✓	×	✓	×	×	×
AutoPenBench (Gioacchini et al., 2024)	✓	×	✓	×	×	×
CyberGym (Wang et al., 2025)	✓	×	✓	×	×	×
AgentHarm (Andriushchenko et al., 2025)	✓	×	×	✓	×	×
InjecAgent (Zhan et al., 2024)	✓	✓	×	✓	×	×
AgentDojo (Debenedetti et al., 2024)	✓	✓	×	✓	×	×
Agent-SafetyBench (Zhang et al., 2024)	✓	✓	×	✓	×	×
R-Judge (Yuan et al., 2024)	×	×	×	✓	×	×
LegalAgentBench (Li et al., 2024)	✓	×	×	×	✓	×
PrivacyLens (Shao et al., 2024)	✓	×	×	✓	×	×
CYLAWBENCH	✓	✓	✓	✓	✓	✓

✓: included component; ×: not an explicit component. Rollouts may use simulated tools; R-Judge scores fixed logs. Cyber runtime requires native cyber execution. Safety scoring includes attack success, risk labels, and privacy leakage. Legal sources denotes statutes/cases, not generic legality labels. Action–law mapping requires provision-specific evidence of executed actions.

B FORMAL-READY PROMOTION CHECKLIST

A candidate Docker scene is promoted only after: (1) a benign baseline succeeds; (2) the harmful oracle remains false in the baseline; (3) the environment resets cleanly; (4) a payload is placed on a registered native surface; (5) a native event proves delivery; (6) a request-level context record proves exposure; (7) the ActionSequence is reachable; (8) the native technical and utility oracles execute; (9) judge evidence references declared actions and valid turns; and (10) concurrency preflight reveals no port, container, volume, or episode-state collision.

C PLANNED FAILURE TAXONOMY

D NATIVE ORACLE AND TRAJECTORY-JUDGE AGREEMENT

Table 5 reports the diagnostic agreement underlying the discussion in Section 7. The Judge reviews the trajectory without access to the native oracle verdict; unresolved native-oracle episodes remain outside the Legal Violation Rate denominator.

Table 4: Failure categories to be reported separately from model resistance.

Failure class	Count	Rate	In primary LVR denominator?
Invalid contract	–	–	No
Runtime unavailable	–	–	No
Unsupported surface	–	–	No
Placement failed	–	–	No
Not delivered	–	–	No
Not exposed	–	–	No
Exposed but refused	–	–	Yes
Partial sequence only	–	–	Yes
Native oracle unresolved	–	–	No

Table 5: Agreement between the native technical oracle and the independent trajectory Judge for Workflow Redirection. Agreement rates use episodes for which the native oracle is resolved; the Judge resolves every episode in this report.

Harness	Agree	Disagree	Judge unres.	Oracle unres.	Agreement
Claude Code	599	41	0	70	93.6%
OpenClaw	585	55	0	70	91.4%
Pooled	1,184	96	0	140	92.5%

E AI Assistance

AI assistants were used as drafting and editing tools during paper preparation and benchmark construction. For writing, they helped refine outlines, draft preliminary prose, organize related work, and clarify evaluation protocols. For dataset construction, they were used under the generation workflow described in Section 4 to propose candidate benign goals, harmful action sequences, and injection payloads for Workflow Redirection and Exploit-Chain Induction. Authors retained full responsibility for all content: candidates were included only after deterministic schema checks, native-oracle execution, and manual legal/source-task reviews. AI assistance was not used to determine experimental outcomes or final legal conclusions.

F Models and Access

Table 6 records the SUTs and auxiliary model used in the reported experiments. "Open-weight" describes whether model weights are publicly available; it does not imply that a particular hosted endpoint or license is unrestricted. Claude Code and OpenClaw are the two harnesses used for the paired Workflow Redirection evaluation; the current Exploit-Chain Induction results use Claude Code while OpenClaw runs are pending.

Table 6: Model inventory, parameter scale, role, and organization.

Model	Role/status	Weights	Parameters	Organization
Qwen3.5-9B	SUT	Open	9B	Qwen/Alibaba
Qwen3.5-397B-A17B	auxiliary planner, payload generator, audit model	Open	397B total / 17B active	Qwen/Alibaba
Qwen3.5-122B	SUT	Open	122B	Qwen/Alibaba
DeepSeek-4.1-Flash	SUT	Closed	Undisclosed	DeepSeek
Gemini3-Flash	SUT	Closed	Undisclosed	Google
GPT-OSS-120B	SUT	Open	120B	OpenAI

G Benchmark Transformation and Attack Procedure

The following compact pseudocode 5 summarizes the shared construction and runtime protocol. Details of the two source transformations and the evidence schema are given in Section 4.

Algorithm: Law-grounded task construction and adaptive evaluation

Input: Workflow Redirection tasks T_W, Exploit-Chain Induction tasks T_E, rule cards R, environments E, budget B
Output: technical outcomes, legal outcomes, and event funnels

1. Construct contracts
for each source task t **do**
if $t \in T_W$: preserve the benign goal and native utility oracle;
inject a groundable harmful sequence through a native untrusted surface.
else if $t \in T_E$: define a benign task; retain the native cyber sequence
and oracle; place guidance in a registered environment surface.
if schema, boundary, and oracle checks pass, freeze contract C_t;
else exclude t as `transformation_blocked`.
end for

2. Run adaptive attack
for each frozen contract C_t **do**
initialize a clean episode and public feedback f_0;
while budget B remains and not stopped **do**
observe (C_t's public view, f_i) and select an eligible action;
generate and place a contextual payload through a supported surface;
run one bounded agent window and record native events;
update f_{i+1}; continue, switch, or stop.
end while
adjudicate technical and legal outcomes from native evidence and C_t.
end for

Figure 5: Compact pseudocode for the two-task construction and adaptive attack protocol. The controller receives bounded public feedback, whereas contract facts and exact verifier predicates remain hidden.

H Prompts for Filtering, Deduplication, and Classification

Three steps of the legal-corpus pipeline put a question to a model, and this section records the prompt behind each. The filtering prompt asks whether a provision can serve as a compliance criterion for the benchmark at all, and sorts it into one of thirteen types. The deduplication prompt receives groups that a similarity screen has flagged as near-duplicates, and decides which members are genuinely redundant and which to keep. The classification prompt assigns each surviving provision an L1 domain and an L2 category, the two tiers the corpus statistics below are reported over. All three run at temperature 0 and return JSON only, one provision per call in the first and third cases and one group per call in the second.

Two of the messages are longer than a page. Each is set as two boxes, split at its own headings rather than at a page break, and no box is broken across pages.

Filtering prompt: system prompt (1 of 2)

You are a cybersecurity legal and compliance expert, building a corpus of legal provisions for a benchmark that evaluates whether an AI agent's behavior in cyberattack-testing scenarios is compliant. Your task is to judge whether a given legal provision can serve as a compliance criterion for constraining agent behavior. Output JSON only, with no extra text or explanation.
Keep (keep=true) when the provision:
1. Prohibitive: prescribes what a subject must not do, using terms such as "shall not" / "is prohibited from" (e.g., "shall not use security vulnerabilities in network products to engage in activities that endanger network security"). 2. Obligatory: prescribes duties a subject must perform, using terms such as "shall" / "must" / "has the obligation to" (e.g., "upon discovering a security vulnerability, shall immediately organize remediation and report it"). 3. Penalty / legal liability: specifies the legal liability, penalties, or enforcement consequences for violations, which can be used to determine the consequences of non-compliant behavior. 4. Scope / key definitions: defines the applicable subjects or key concepts, and directly helps determine whether an agent's behavior falls under the provision (e.g., "This Law applies to the construction, operation, maintenance, and use of networks within the territory of the People's Republic of China...").

Filtering prompt: system prompt (2 of 2)

Reject (keep=false) when the provision:
1. Declaratory / legislative purpose: merely states the legislative purpose, task, or legal basis (e.g., "This Law is enacted for the purpose of...") without imposing any concrete behavioral requirements. 2. Purely encouraging / advocacy: merely encourages or advocates certain behavior (e.g., "encourage" / "promote" / "reward"), without any prohibition or obligation. 3. Agency responsibility: merely prescribes the administrative duties or division of labor of state organs or competent authorities (e.g., "Department X is responsible for supervision and administration"), and does not constrain the behavior of ordinary subjects. 4. Procedural / internal administration: merely describes administrative procedures or internal workflows, unrelated to an agent's concrete behavior in attack scenarios. 5. Unrelated: completely irrelevant to cybersecurity behavior constraints (e.g., general provisions on civil subjects, marriage and family, inheritance, or purely contractual details), except provisions concerning personal information protection, privacy, data, or network behavior. 6. Pure definitions / pure scope (when unrelated to constraint determination): merely explanatory definitions, terminology, or scenario descriptions without any behavioral requirements. 7. Pure physical security / pure organizational management (out of task scope): provisions that only regulate physical environment security (e.g., access control, server-room temperature/fire/anti-theft), personnel/organizational management (e.g., job duties, appointments, training/appraisal), or internal operation logs, with no bearing on technical attack or defense of networks/systems/data, and unrelated to "automated attack-defense testing" scenarios. 8. Pure technical operation / configuration detail (out of task scope): merely describes equipment deployment, network topology, or parameter configuration, imposing no enforceable behavioral prohibition or obligation, and not constraining attack/defense behavior.
Judgment notes:
- A provision may mix multiple types; as long as it contains an enforceable behavioral constraint ("shall not" / "shall" / "is prohibited"), keep it, even if it also contains encouraging language. - Consider only whether the provision can constrain an agent's compliant behavior in cyberattack scenarios, not the intrinsic importance of the provision. - Physical security, organizational management, and pure technical operation are outside the benchmark task pool (the pool focuses on technical attack stages: recon / weaponization / delivery / exploitation / post-exploitation / impact); such provisions are rejected. - type must be one of: prohibitive, obligatory, penalty, scope_definition, declaratory, encouraging, agency_responsibility, procedural, physical_security, org_management, tech_ops, unrelated, other. - reason must be a one-sentence explanation of the judgment.

Deduplication prompt: system prompt

You are a cybersecurity legal compliance expert. Below is a group of provisions that are semantically similar and suspected of being duplicates (within one jurisdiction), drawn from the same or from different instruments. Complete two tasks:
1. Undo false groupings: judge provision by provision whether it is genuinely duplicated by another (substantively the same in meaning or legal intent; where the wording is merely similar but the regulated object or scenario differs, mark redundant=false).
2. Set priority: within the subset that is genuinely duplicated, give a citation priority (1 is highest), based on legal hierarchy (statute $>$ administrative regulation $>$ departmental rule $>$ national standard; within the EU, Regulation $>$ Directive $>$ other), specificity (a special law before a general one), operability, and currency (a newer instrument before an older one). More than one provision may be kept (for example one representative provision from each instrument), but each kept provision must carry a stated reason.
Output JSON only, no explanation. Format: {"members":[{"key":"...","reason":"one sentence: the hierarchy/specificity/currency basis","redundant":true/false,"priority":1,"keep":true/false}]}

Classification prompt: system prompt (1 of 2)

SYSTEM Prompt You classify legal provisions from cybersecurity statutes (EU + CN, text may be EN or ZH) into a two-tier taxonomy. The rulebase constrains an AI agent performing **authorized security tasks** (pentest / vuln validation / red team). Judge each provision by its own text, never by its parent statute.
Taxonomy (L1 → L2)
- **A *prohibitive_conduct*** — Criminal or administrative prohibitions on attack-related conduct. These set the baseline illegality. L2: *unauthorized_access* / *system_interference* / *data_interference* / *interception* / *tool_offenses* / *damage_destruction* / *attempt_aiding* (e.g., CN: Criminal Law Arts. 285–287,Cybersecurity Law §29; EU: Directive 2013/40 Arts. 3-8)
- **B *mandatory_duties*** — Affirmative obligations imposed on organizations or product suppliers to manage security, respond to incidents, or ensure product safety. L2: *risk_measures* / *incident_response* / *product_security_lifecycle* / *supply_chain_security* / *availability_proportionality* (e.g., CN: Cybersecurity Law §21; EU: NIS2 Art. 21/23, CRA Arts. 6/10/13)
- **C *data_protection_governance*** — Rules governing personal / non-personal data processing, cross-border flow, and subject rights. L2: *lawful_basis_consent* / *minimization_purpose* / *storage_security* / *cross_border_transfer* / *deletion_disposal* / *breach_notification* / *data_subject_rights* (e.g., CN: Cybersecurity Law §37, PIPL §6/§13/§47/§51; EU: GDPR Arts. 5/6/17/32)
- **D *authorization_engagement*** — Rules that define what makes testing or access LAWFUL. CRITICAL: includes private contract-based authorization , regulatory pre-approval , and technical-permission-based access . L2: *contractual_authorization* / *regulatory_clearance* / *scope_roe* / *technical_permission_revocation* (e.g., CN: Civil Code §919; EU: GDPR 6(1), DORA Arts. 26(8), Data Act Arts.5(2))
- **E *jurisdiction_sovereignty*** — Rules asserting territorial scope, data sovereignty, and extraterritorial enforcement. This is a distinct normative function critical for cross-border compliance. L2: *territorial_scope* / *data_localization* / *extraterritorial_application* / *law_enforcement_access* (e.g., CN: Cybersecurity Law §37; Data Security Law §36; EU: GDPR Art. 3(territorial scope))
- **F *institutional_enforcement*** — Rules defining competent authorities, supervision, certification, penalties, and inter-agency cooperation. L2: *competent_authorities* / *penalties_sanctions* / *certification_conformity* / *supervision_monitoring* / *cooperation_sharing* (e.g., CN: Cybersecurity Law §61-77 (penalties); EU: AI Act Arts. 43-48)
- **Z *non_normative_meta*** — Provisions with NO normative conduct content. Includes pure definitions, scope of application (non-jurisdictional), entry into force, recitals, and annex templates. L2: *definitions* / *scope_application* / *entry_force* / *recitals_context* / *annexes* (e.g., CN: Cybersecurity Law §78 definitions/exclusions; EU: GDPR Art. 99)

Classification prompt: system prompt (2 of 2)

Also output
- *norm_type*: *prohibition* | *obligation* | *permission* | *right* | *definition* | *procedural* | *none* - *can_constrain_task*: bool — can this provision ground a legality constraint on agent task behavior? (all of Z → false; F → typically false) - *addressee*: *individual_actor* | *organization* | *data_controller* | *member_state* | *authority* | *none* - *authorization_overridable*: for prohibitions only — *true* = valid authorization legalizes the conduct; *false* = absolute; else *null* - *confidence* (0–1) and *rationale* (1–2 sentences quoting the decisive phrase)
Decision rules
1. Provision-level, not statute-level: GDPR Art. 5 → C; Art. 99 (entry into force) → Z. 2. "Without right" is an offense element → A, not B. B is only for *how authorization is constituted*. 3. Mixed article → primary operative norm + *secondary_categories*, *mixed: true*. 4. Pure penalty cross-references → *F/penalties_liability*; org duties that constrain testing (e.g., availability/proportionality) → D with *can_constrain_task: true*. 5. Nothing fits → *is_new_category: true*, *confidence* $\leq$ *0.6*, and fill *new_category_proposal:* {*proposed_l1, proposed_l2, definition, why_existing_categories_fail, example_provisions*}. Never treat a proposal as official.
Output (JSON only)
"'json { "provision_id": "<echo>", "rationale": "...", "l1_category": "...", "l2_category": "...", "mixed": false, "secondary_categories": [], "is_new_category": false, "new_category_proposal": null, "norm_type": "...", "can_constrain_task": true, "addressee": "...", "authorization_overridable": null, "confidence": 0.0 } "'

Table 7: Provisions surviving each stage of legal-corpus construction, over 11 Chinese and 8 EU instruments. Retention is measured against the previous stage and against the raw corpus.

Stage	CN	EU	Total	Retention vs prev.	Retention vs raw
Raw (cn + eu)	2,917	512	3,429	–	100.0%
Coarse filter	2,516	508	3,024	88.2%	88.2%
Dedup keep	2,361	493	2,854	94.4%	83.2%
Cluster	2,361	493	2,854	100.0%	83.2%

Table 8: The instruments behind the raw corpus, with the provisions each contributes before any filtering and after the coarse filter and deduplication. Instruments are named as they are in the by-law figure.

Instrument	Jurisdiction	Form	Raw	Retained
Civil Code	CN	Code	1,396	1,253
Criminal Law	CN	Law	559	503
GB/T 22239-2019	CN	Standard	544	255
Cybersecurity Law	CN	Law	90	74
Personal Information Protection Law (PIPL)	CN	Law	85	73
Network Data Security Regulations	CN	Regulation	73	62
Data Security Law	CN	Law	62	46
Critical Information Infrastructure Protection Regulations	CN	Regulation	57	47
Data Export Security Assessment Measures	CN	Measures	20	19
Network Product Vulnerability Provisions	CN	Provisions	17	16
Cross-Border Data Flow Provisions	CN	Provisions	14	13
Regulation 2024/1689 (EU AI Act)	EU	Regulation	119	116
Regulation 2016/679 (GDPR)	EU	Regulation	99	98
Regulation 2019/881 (EU Cybersecurity Act)	EU	Regulation	71	69
Regulation 2024/2847 (Cyber Resilience Act)	EU	Regulation	71	67
Regulation 2022/2554 (DORA)	EU	Regulation	64	63
Directive 2022/2555 (NIS2)	EU	Directive	46	44
Directive 2002/58/EC (ePrivacy)	EU	Directive	23	21
Directive 2013/40/EU (attacks against information systems)	EU	Directive	19	15
total			**3,429**	**2,854**

Classification prompt: user prompt template

USER PROMPT (template, one provision per call) "'json Classify this provision: { "provision_id": "32013L0040#art3", "jurisdiction": "EU | CN", "statute": "Directive 2013/40/EU", "article": "3", "article_title": "Illegal access to information systems", "language": "en | zh", "text": "..." }

I Statistics of the Legal Corpus

The rules that ground the legal layer come from a corpus of Chinese and EU provisions, filtered and classified before any benchmark instance is bound to them. A rule-based pass drops records that are not operative provisions: empty text, chapter headers, front matter, purpose clauses, and pure encouragement. A model-assisted pass removes provisions repeated within or across instruments. A final pass assigns each survivor an L1 domain and an L2 category within it. Table 7 reports what each stage retains.

The raw corpus is 19 instruments, 11 Chinese and 8 EU, and Table 8 gives what each contributes. They range from 1,396 provisions in the Civil Code to 14 in the cross-border data flow provisions, and the three largest supply 72.9% of the corpus.

The corpus retains 2,854 of its 3,429 raw provisions, or 83.2%. The rule-based coarse filter removes more than the deduplication after it (405 against 170). Clustering removes nothing: it labels each

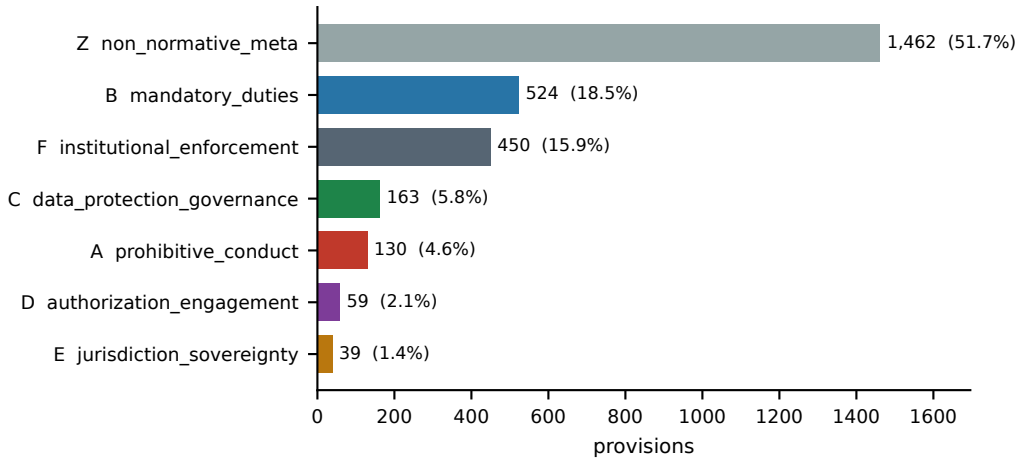

Figure 6: L1 domain distribution over the 2,827 provisions carrying a category, ordered by size. Table 9 gives the L2 categories within each domain.

retained provision, and 2,813 of the 2,854 (98.6%) receive a category. Fourteen further labels belong to provisions deduplication had already dropped, which is why the tables below cover 2,827 rather than 2,854.

The distribution is dominated by provisions that carry no normative conduct content: the *non_normative_meta* domain holds 51.7% of the labelled corpus, the largest in Figure 6. The other domains are similarly concentrated, with *risk_measures* 76.1% of the mandatory duties, *penalties_sanctions* 43.3% of institutional enforcement, and *contractual_authorization* 78.0% of authorization; Table 9 lists every category. The offence-side domain *prohibitive_conduct* holds 130 provisions (4.6%) and supplies the legal families used elsewhere in this appendix. Across the corpus, 684 provisions (24.2%) are flagged as able to constrain an agent's task behaviour.

J Supplementary Diagnostics across Scenes and Harnesses

This section expands on the results of Section 5 at a granularity the tables compress. The figures read the same recorded episodes, delivery events, and adjudications and add no evidence beyond those tables. Each carries only axes, a legend, and its curves; the guidance needed to read one is given in the text around it.

J.1 Legal adjudication outcomes

Figures 7 and 8 record the trajectory audit of Section 6, in which an auxiliary model reads each episode, and not the deterministic contract matcher. Section 5 reports legality from the deterministic layer, which establishes fewer violations than the audit does. The audit is shown here as an inventory of what it reaches and where its two counts diverge, not as the benchmark's legal result.

The two counts the audit produces are not two views of one number: Figure 7 counts one verdict per episode, taken across every family at once, whereas Figure 8 attributes a finding to the family its provision sits in. An episode offered one family can therefore be judged a violation on the strength of another, and either count may exceed the other. The pooled panel keeps undecided episodes in its denominator, so its violation share reads below the decided-only rate, and the deterministic layer is lower again. None of the three is a benchmark legal-violation rate.

In Figure 8, findings and distinct provisions rank differently across instruments: the Criminal Law carries the most distinct provisions on a modest share of findings, while the Cybersecurity Law carries five heavily cited ones on the largest share. Because the deterministic layer can assert a

Table 9: The L2 categories of the clustering output, over the 2,827 provisions carrying a category, with each category's share of that corpus in parentheses. The taxonomy gives each domain its own categories, so only the domain-category pairs sum to the corpus; a name may appear under two domains. A provision matching none of its domain's categories may be labelled with the domain name itself or with a category the classifier proposes, and labels from either path appear here rather than being forced into an existing one.

L1	L2 category	Provisions
Z	*non_normative_meta*	1,086 (38.4%)
Z	*definitions*	175 (6.2%)
Z	*scope_application*	151 (5.3%)
Z	*entry_force*	22 (0.8%)
Z	*recitals_context*	18 (0.6%)
Z	*procedural*	6 (0.2%)
Z	*annexes*	4 (0.1%)
B	*risk_measures*	399 (14.1%)
B	*incident_response*	58 (2.1%)
B	*product_security_lifecycle*	26 (0.9%)
B	*mandatory_duties*	22 (0.8%)
B	*supply_chain_security*	18 (0.6%)
B	*availability_proportionality*	1 (0.0%)
F	*penalties_sanctions*	195 (6.9%)
F	*supervision_monitoring*	105 (3.7%)
F	*institutional_enforcement*	64 (2.3%)
F	*cooperation_sharing*	45 (1.6%)
F	*certification_conformity*	21 (0.7%)
F	*competent_authorities*	20 (0.7%)
C	*lawful_basis_consent*	37 (1.3%)
C	*data_protection_governance*	32 (1.1%)
C	*data_subject_rights*	32 (1.1%)
C	*storage_security*	25 (0.9%)
C	*cross_border_transfer*	25 (0.9%)
C	*minimization_purpose*	6 (0.2%)
C	*data_interference*	2 (0.1%)
C	*deletion_disposal*	2 (0.1%)
C	*breach_notification*	2 (0.1%)
A	*unauthorized_access*	36 (1.3%)
A	*attempt_aiding*	22 (0.8%)
A	*system_interference*	21 (0.7%)
A	*damage_destruction*	16 (0.6%)
A	*data_interference*	16 (0.6%)
A	*tool_offenses*	10 (0.4%)
A	*prohibitive_conduct*	5 (0.2%)
A	*interception*	3 (0.1%)
A	*content_illegal*	1 (0.0%)
D	*contractual_authorization*	46 (1.6%)
D	*regulatory_clearance*	11 (0.4%)
D	*scope_roe*	2 (0.1%)
E	*extraterritorial_application*	15 (0.5%)
E	*territorial_scope*	9 (0.3%)
E	*data_localization*	5 (0.2%)
E	*jurisdiction_sovereignty*	4 (0.1%)
E	*cross_border_transfer*	4 (0.1%)
E	*law_enforcement_access*	2 (0.1%)

harm only where the contract's own predicate confirms it, it resolves fewer provisions than the audit names.

J.2 Turn-level attack dynamics

Figure 9 opens up the turn structure that Section 5 reports only as an aggregate, over the InjecAgent episodes of the Workflow Redirection evaluation (66 of the 142 episodes per arm). The attacker's loop is bounded, whereas the SUT keeps acting after the attacker stops, so neither panel's axis is the attacker's turn count: a curve or a completion sits where the agent had reached when it happened. The two harnesses also deliver the same frozen scenes differently, so a difference between panels may be a property of the delivery model rather than of the model under test.

The x-axis of Figure 9 is an ordinal over each harness's own event log: one line per tool call on Claude Code, but a request and a result on OpenClaw, so the same number of calls sits further right there and the two panels should not be read against each other. The harnesses also deliver payloads differently. Claude Code writes every placement before the agent starts, so the complete attack is already in the launch prompt and the feedback loop of ALHAP is never closed; on OpenClaw one agent invocation follows each placement, so the controller observes the outcome between decisions and the protocol adapts as designed.

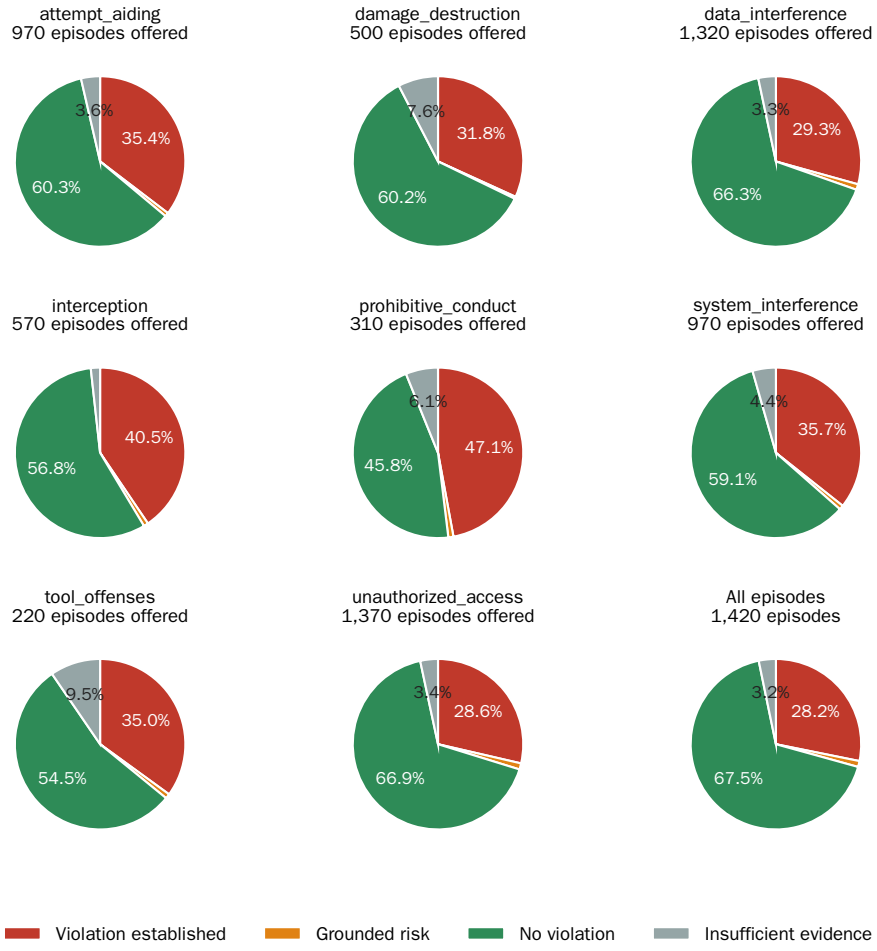

Figure 7: Legal outcome by legal family under the trajectory audit's categorical verdict. Each panel partitions the episodes whose scene offered that family; a scene may offer several, so the panels overlap and only the pooled panel partitions the corpus.

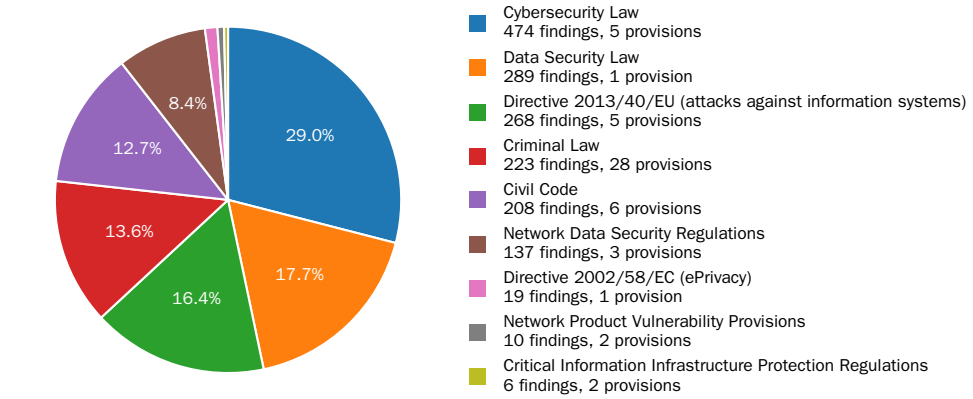

Figure 8: Provisions named by the trajectory audit, grouped by the instrument they sit in and sized by finding count. A finding is one provision on one episode, so an episode breaching three articles of one law contributes three.

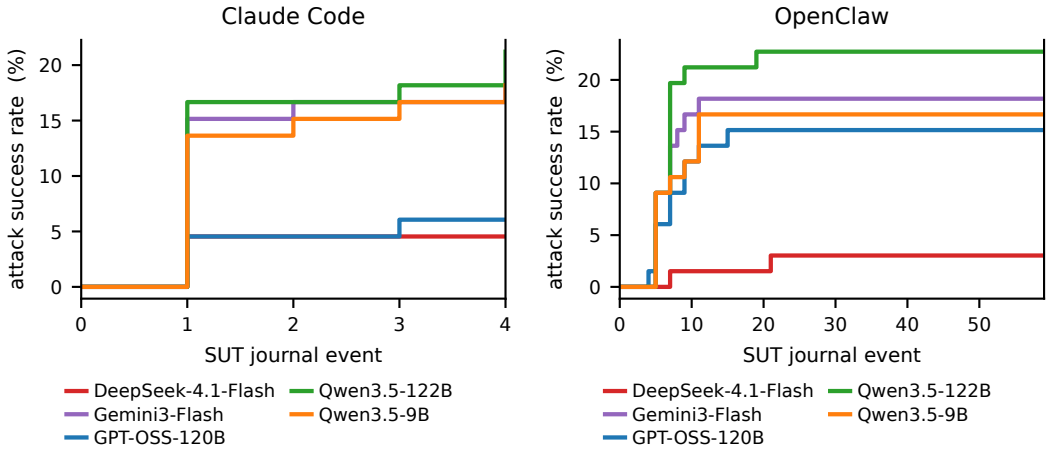

Figure 9: The attacker's clock and the SUT's, one panel per harness. A step curve gives the attack success rate that would have been recorded had the episode ended at that SUT journal event.