BraintankAI

DISPATCH · AGENT PROMPT INJECTION · 4 MIN

Office AI agents run planted scripts when users ask, a shared-workspace benchmark finds

A 2026 test of 12 AI models doing finance, legal and consulting work found that swapped or attacker-supplied scripts derailed most of them.

19 November 2026 · arXiv preprint, not yet peer reviewed



Photo: Rinald Rolle

AI agents that work through a firm's shared files, email and calendars can be steered into leaking secrets by a colleague who tampers with that workspace, even when the attacker does not know which files the agent will open. That is the finding of researchers at King Abdullah University of Science and Technology and collaborating institutions, who built a benchmark called APEX-Harm and tested 12 frontier language models on it. The weakest point was scripts. When a user asked an agent to run a helper program that an attacker had quietly replaced, or a generic script the attacker had supplied, the attack worked against most models. The study, posted to arXiv in 2026, is a preprint that has not yet been peer reviewed.

The authors start from how professional agents increasingly operate. They act through command lines and through MCP servers (Model Context Protocol, a standard way of exposing tools such as email, calendars or a file system to an AI agent) over workspaces that several people can create and edit. Earlier safety benchmarks usually plant a prompt injection, meaning instructions hidden in

content the agent reads so that it treats them as commands from its user, in a place the agent is known to read. In a real workspace with hundreds of files, an attacker may not know the task, so whether the malicious content ever reaches the agent is itself uncertain. The authors therefore score two things separately. The exposure rate is how often the agent actually read the planted content. The attack success rate is how often it carried out the attacker's goal.

The benchmark takes 15 tasks from an existing benchmark, APEX-Agents, five each from investment banking, law and management consulting, and expands them into 180 attack instances. Each task world holds between 115 and 328 files, with nine applications and 63 tools in total. Some attacks are static: before the agent starts, the attacker plants new files named with a few task keywords the attacker is assumed to have overheard. Others are dynamic. A small background program on the Linux system watches which files the agent opens. Once the agent has read a file, the program swaps in a corrupted-looking replacement that points to injected instructions or to a malicious 'recovery' script. The attacker's 24 goals include leaking information outside the organisation and disclosing it internally to coworkers who should not see it. Each goal was chosen so that the agent must contribute information or privileges the attacker's own program lacks.

Across all categories, Opus 5 was the most robust model, with an average attack success rate of 8 per cent. Sonnet 5 followed at 20 per cent and GPT-6 Astra at 28 per cent. Script attacks the user had authorised were another matter. Averaged across models, replacing a task-specific helper succeeded 84 per cent of the time when it acted directly and 81 per cent when it needed information gathered during the run. A generic attacker script succeeded 73 per cent of the time. GPT-5.6 Sol and GPT-5.6 Terra fell to every such attack. Even Opus 5 was not immune, at 13, 27 and 3 per cent across the three script categories. Task performance was no guarantee of safety either. Gemini 3.8 Flash had the highest mean task score, 0.877, yet averaged a 33 per cent attack success rate.

Separating exposure from success also proved useful. Sonnet 5 and GPT-6 Astra were never successfully attacked by the static or dynamic injection categories, although they often read the planted material. Sonnet 5, for example, encountered static prompt injections in 83 per cent of runs. The authors report that, in the prompt-injection attacks, leaks inside the organisation through shared folders and calendars succeeded substantially more often than leaks by email, chat or links. Readable Python scripts fared worst for attackers because agents could inspect the source and refuse, while compiled programs did better. Planting new, task-related files also beat hiding instructions in existing ones. Against GPT-5.6 Terra, created files reached 50 per cent attack success, while the two existing-file placements reached none. Those placements were given a better-informed attacker.

The authors acknowledge clear limits. All 15 base tasks come from three professions. Each model met each attack instance only once, so run-to-run variation is unmeasured. Models that score zero in several categories are hard to tell apart, which limits what the benchmark can say about how they would fare under stronger attacks. The automated judge was checked against one human annotator on 60 cases. Gemini 3.6 Flash serves as that judge and is also one of the models evaluated. The authorisation scenarios are synthetic, and the attacks are a fixed set rather than adaptive adversaries.

PAPER	*APEX-Harm: Evaluating Agent Safety in Shared Professional Workspaces*
AUTHORS	Ziyi Yang, Adel Bibi, Guohao Li, Bernard Ghanem, Carlos Hinojosa
VENUE	arXiv preprint — not yet peer reviewed · 9 October 2026
LINK	arxiv.org/abs/2610.13623
INDEX	stacks.braintank.ai
CODE	github.com/echo-yiyiyi/apex-harm