

Evading Fake News Detection: Self-evolving Social Manipulation via LLM-driven Agent Orchestration

Mengyang Chen^{1,2} Lingwei Wei^{1*} Wei Zhou¹ Songlin Hu^{1,2}

¹Institute of Information Engineering, Chinese Academy of Sciences

²School of Cyber Security, University of Chinese Academy of Sciences

{chenmengyang, weilingwei, zhouwei, husonglin}@iie.ac.cn

Abstract

While propagation-based fake news detection improves performance by modeling how news spreads, it remains vulnerable to adversarial social manipulation, where adversaries distort propagation patterns through malicious social interactions to evade detection. Existing attacks often target only one type of social behavior with predefined rules, making them vulnerable to robust defenses. To bridge this gap, we propose a new Self-evolving Social Manipulation framework (EvoSM) that utilizes Large Language Models (LLMs) to orchestrate evolutionary social manipulation campaigns for evading fake news detection. First, we design an LLM-driven Social Manipulation Orchestrator that uses a social graph interface to perceive the propagation environment, coordinate diverse manipulation actions, and refine attack strategies in context. Second, to further strengthen the manipulation capability, we introduce Reflective Experience Memory, which distills reusable manipulation insights from previous successful and failed attack attempts through reflection. Experiments on four real-world fake news datasets demonstrate that EvoSM consistently outperforms comparison baselines in attack success rates and transferability. Further analysis shows that EvoSM benefits from coordinated multi-behavior manipulation and experience-guided strategy evolution. Our study highlights the vulnerability of existing fake news detection systems to adaptive social manipulation.

1 Introduction

The persistent proliferation of fake news continues to undermine the integrity of global information ecosystems (Fisher et al., 2016; Vosoughi et al., 2018). Existing fake news detection methods mainly focus on textual content, such as news text and comments (Castillo et al., 2011; Ma et al.,

2015; Yu et al., 2017), and propagation information, modeled as message propagation trees (Bian et al., 2020; Wei et al., 2021; Sun et al., 2022; Su et al., 2023; Chen et al., 2024; Zhang et al., 2025a) or news engagement graphs (Lu and te Li, 2020; Wu and Hooi, 2023). By integrating structural and semantic features, these propagation-based approaches typically achieve superior performance compared to content-based detectors.

However, propagation-based detectors remain vulnerable to adversarial social manipulation, where adversaries distort propagation patterns through malicious social interactions (Wei et al., 2026). Existing attacks typically focus on a single type of social behavior, such as editing news content (Wu et al., 2024; Tian et al., 2025; Bethany et al., 2025; Park et al., 2025), editing or deleting comments (Wan et al., 2025), perturbing interactions around the target news (Zhu et al., 2024; Wang et al., 2023; Zeng et al., 2025), or injecting adversarial replies (Luo et al., 2024; Wan et al., 2025; Chen et al., 2026). Moreover, these attacks follow predefined perturbation rules, which limits their ability to adapt to diverse propagation contexts and makes them easier to be mitigated by robust defenses (Zhang et al., 2025a,b; Wan et al., 2025). In contrast, real-world social manipulation often coordinates multiple social actions to influence news spread, rather than relying on a single isolated operation (Pacheco et al., 2021).

In this paper, we make the first attempt to leverage LLM-driven agents to achieve self-evolving social manipulation. While general LLM-agent frameworks (Yao et al., 2023; Shinn et al., 2023) and tool-augmented methods (Edge et al., 2024; Wang et al., 2025) can be equipped with multiple tools, simply adding tools is insufficient. The agent must perceive propagation contexts, translate manipulation intents into executable social graph operations, and continuously refine its orchestration strategies through self-evolution. This

*Corresponding author.

leads to our core research question: **Can LLM-driven agents orchestrate and evolve multiple social behaviors into adaptive evasion campaigns against propagation-based fake news detectors?**

To answer this question, we propose Self-Evolving Social Manipulation (**EvoSM**), a unified framework that utilizes LLMs to orchestrate evolutionary social manipulation campaigns for evading fake news detectors. EvoSM consists of two synergistic components. First, an LLM-driven Social Manipulation Orchestrator uses a social graph interface to perceive the propagation environment, coordinate diverse manipulation actions, and translate high-level manipulation intents into executable social behaviors. Second, Reflective Experience Memory distills reusable manipulation insights from previous successful and failed attack attempts, enabling the agent to refine its strategies across future targets. Through the synergy of these components, EvoSM coordinates multiple social behaviors into coherent attack trajectories and progressively evolves its manipulation strategies.

The experimental results on four real-world fake news datasets demonstrate that EvoSM outperforms existing attack baselines in attack success rates and transferability across different detectors. Further analysis reveals that EvoSM exhibits evolutionary strategic behaviors: the agent autonomously switches manipulation strategies, gradually making fake news propagation patterns resemble those of authentic news.

Our contributions are summarized as follows: 1) We identify a new self-evolving social manipulation attack that leverages LLM-driven agents to evade propagation-based fake news detectors. 2) We propose EvoSM, an LLM-driven framework for adaptive social manipulation. It uses a Social Manipulation Orchestrator to translate propagation contexts into executable social actions, and Reflective Experience Memory to distill reusable attack experience for strategy evolution. 3) Experiments demonstrate that EvoSM achieves higher attack success rates against propagation-based detectors. Further analysis reveals that EvoSM exhibits evolutionary strategic attack behaviors.

2 Preliminaries

Propagation-based Fake News Detection.

Fake news detection aims to verify the authenticity of a news piece, typically formulated as

binary classification with label $y \in \{0, 1\}$ (0 for Real, 1 for Fake). Propagation-based methods model social propagation as a graph $G = (\mathcal{V}, \mathcal{E})$, commonly instantiated as either a **Message Propagation Tree (MPT)** or a **News Engagement Graph (NEG)**. MPT represents an individual dissemination tree, where $\mathcal{V} = \{v^*, c_1, \dots, c_n\}$ contains the news v^* and its comments/retweets $\{c_1, \dots, c_n\}$ and $\mathcal{E}$ denotes interaction links such as replies. NEG represents a global engagement network, where $\mathcal{V} = \{v_1^*, v_2^*, \dots\}$ contains news articles and $\mathcal{E}$ captures correlations derived from user engagement behaviors, such as sharing by common users. Despite these structural differences, the objective is to learn a detector f that predicts the veracity label y of the news item v^* from the propagation graph G .

Adversarial Social Manipulation. Adversarial social manipulation aims to mislead propagation-based detectors by manipulating social behaviors in the propagation environment. Given a target fake-news instance v^* together with its propagation graph G with label $y = 1$ and a black-box detector f , the adversary seeks to generate an adversarial graph G_{adv} such that the target news v^* is classified as real:

$$f(v^*; G_{adv}) = 0, \quad \text{s.t. } G_{adv} \in \mathcal{N}_{\mathcal{A}}(G, v^*), \quad (1)$$

where $\mathcal{N}_{\mathcal{A}}(G, v^*)$ denotes the set of permissible adversarial propagation graphs obtainable from G through manipulation actions in the action space $\mathcal{A}$ with respect to the target news v^* . Following Wan et al. (2025), $\mathcal{A}$ includes executable social behaviors such as content editing, comment editing/deletion, reply injection, and interaction perturbation. These operations modify the semantic content, propagation structure, or both, while remaining within the predefined social-manipulation interface.

Different from existing attacks that typically optimize a single manipulation behavior or follow predefined perturbation rules, we formulate adversarial social manipulation as a multi-step, self-evolving evasion campaign that coordinates heterogeneous social behaviors. The adversary dynamically adapts its manipulation strategy according to the current propagation context and feedback observed during the attack process, rather than committing to a fixed perturbation pattern in advance.




Figure 1: Overall framework of EvoSM. EvoSM orchestrates multiple social behaviors into adaptive evasion campaigns against propagation-based fake news detectors. The LLM-driven Social Manipulation Orchestrator (SMO) uses a social graph interface to observe propagation contexts, update manipulation strategies, and translate manipulation intents into executable social graph operations. Reflective Experience Memory (REM) reflects on surrogate detector (f_θ) feedback during each attack attempt to refine manipulation strategies. Upon success or budget exhaustion (T), the final manipulated graph G_{adv} is evaluated against the target detector (f). REM then distills successful skills and failure lessons into the experience bank for future manipulation. Blue, orange, and black arrows denote the social-context observation loop, behavior execution loop, and post-attack flow, respectively.

Threat Model. Following common practice in recent adversarial research (Wang et al., 2023; Luo et al., 2024; Zhu et al., 2024), we assume a strict black-box setting in which the adversary has no access to the parameters, gradients, or intermediate confidence feedback of the target detector f . Instead, the adversary relies on a surrogate detector f_θ to estimate fake-news confidence scores c_{f_θ} and guide the manipulation process. Accordingly, all intermediate attack decisions, including strategy refinement and action selection, are driven by surrogate-model feedback, while the target detector is used only to evaluate whether the final manipulated graph successfully transfers to the target model. The adversary can observe the complete propagation graph of the target news and is modeled as a high-threat yet constrained malicious publisher equipped with the unified action space $\mathcal{A}$ to coordinate permissible social behaviors¹.

3 Methodology

3.1 Overview

In this work, we propose **EvoSM** (Self-Evolving Social Manipulation), a unified LLM-driven

¹We also evaluate the agent as a malicious user limited to additive-only actions. Detailed analysis for this variant is provided in Appendix B.1

framework that formulates adversarial attacks against fake news detectors as adaptive social manipulation campaigns. As illustrated in Figure 1, EvoSM consists of two components: **1) LLM-driven Social Manipulation Orchestrator (SMO):** SMO uses a social graph interface to observe propagation states, reason over manipulation contexts, and translate high-level intents into executable social behaviors. **2) Reflective Experience Memory (REM):** REM reflects on surrogate detector (f_θ) feedback during each attack attempt to refine ongoing strategies. Upon success or budget exhaustion (T), the final manipulated graph G_{adv} is evaluated by the target detector (f). After termination, REM distills successful skills and failure lessons from the completed trajectory into reusable experience for future attacks. The overall process is shown in Algorithm 1, with prompt details in Appendix C.6.

3.2 LLM-driven Social Manipulation Orchestrator

To make multi-behavior social manipulation executable, the LLM-driven Social Manipulation Orchestrator (SMO) bridges high-level manipulation intent and concrete social behaviors. Given the current propagation graph G_t , the agent observes

Algorithm 1: Self-Evolving Social Manipulation (EvoSM)

Input : Target graphs $\{G_0^{(i)}\}_{i=1}^N$ (MPT or NEG), surrogate detector f_θ , attack budget T

Output : Adversarial graphs $\{G_{adv}^{(i)}\}_{i=1}^N$

```

1 Initialize reflective experience memory  $\mathcal{E} \leftarrow \emptyset$ 
2 foreach target graph  $G_0^{(i)}$  do
3    $G_t \leftarrow G_0^{(i)}, t \leftarrow 0, \mathcal{M} \leftarrow \emptyset$ 
4    $s_0 \leftarrow \text{InitializeContext}(G_0^{(i)})$ 
5    $\mathcal{E}^{ret} \leftarrow \text{RetrieveExperience}(\mathcal{E}, G_0^{(i)})$ 
6   // Reusable guidance from REM
7   while  $t < T$  and not surrogate success do
8     // Social manipulation orchestration
9      $h_t \sim \text{LLM}_\theta(h \mid s_t, \mathcal{M}, \mathcal{E}^{ret})$ 
10     $a_t \sim \text{LLM}_\theta(a \mid h_t, s_t, \mathcal{M}, \mathcal{E}^{ret})$ 
11    if  $a_t \in \mathcal{T}_{obs}$  then
12      // Social-context observation
13       $obs \leftarrow \text{ExecuteTool}(G_t, a_t)$ 
14       $s_t \leftarrow s_t \cup obs$ 
15       $\mathcal{M} \leftarrow \mathcal{M} \cup \{(h_t, a_t, obs)\}$ 
16      continue // No graph modification
17    else if  $a_t \in \mathcal{T}_{exe}$  then
18      // Social behavior execution
19       $G_{t+1} \leftarrow \text{ApplyAction}(G_t, a_t)$ 
20      // Reflective experience update
21       $r_t \leftarrow c_{f_\theta}(G_t) - c_{f_\theta}(G_{t+1})$ 
22       $z_t \leftarrow \text{Reflect}(r_t, \mathcal{M})$ 
23      // Success/Failure/Deadlock
24       $\mathcal{M} \leftarrow \mathcal{M} \cup \{(h_t, a_t, r_t, z_t)\}$ 
25       $s_{t+1} \leftarrow \text{RefreshContext}(G_{t+1})$ 
26       $G_t \leftarrow G_{t+1}, t \leftarrow t + 1$ 
27    end
28     $G_{adv}^{(i)} \leftarrow G_t$ 
29     $y_i \leftarrow f_\theta(G_{adv}^{(i)})$  // Evaluation on surrogate detector
30     $e_i \leftarrow \text{DistillExperience}(\mathcal{M}, G_0^{(i)}, G_{adv}^{(i)}, y_i)$ 
31     $\mathcal{E} \leftarrow \mathcal{E} \cup \{e_i\}$  // Update REM
32  end
33 return  $\{G_{adv}^{(i)}\}_{i=1}^N$ 

```

the social context, updates its manipulation strategy, and selects a feasible behavior with its target and parameters. In this way, the orchestrator instantiates a macro-level attack intention as a micro-level graph operation for the current target.

The orchestrator uses a social graph interface. Unlike general Graph-LLM tools (Zhang, 2023; Wang et al., 2025) that mainly support passive graph understanding or QA, our interface provides a unified toolset $\mathcal{T} = \mathcal{T}_{obs} \cup \mathcal{T}_{exe}$ for both social-context observation and manipulation execution. At each timestep t , given the textualized social context s_t , local trajectory memory $\mathcal{M}$, and retrieved experience $\mathcal{E}^{ret}$, the LLM first updates its

	Tool Function	Applicability		Description
		MPT	NEG	
Observation	GetGlobalStats	✓	✓	Summarize global propagation context.
	ScanEgoGraph	✓	✓	Inspect k -hop local neighborhood.
	ReadNode	✓	✓	Fetch the content of specific nodes.
	CheckInfluence	✓	✓	Estimate node influence from graph signals.
	Search	✓	✓	Locate nodes by textual keywords.
Execution	AddEdge	-	✓	Add an interaction edge in the network.
	AddComment	✓	-	Inject a comment in the propagation tree.
	AddNews	-	✓	Publish a news node in the network.
	RemoveEdge	-	✓	Remove an interaction edge in the network.
	EditNews	✓	✓	Modify the target news content.
	EditComment	✓	-	Modify an existing comment in the tree.
	DeleteComment	✓	-	Remove a comment and its subtree.

Table 1: Details of the social graph interface (SGI).

strategy and then selects an action:

$$h_t \sim \text{LLM}_\theta(h \mid s_t, \mathcal{M}, \mathcal{E}^{ret}), \quad (2)$$

$$a_t \sim \text{LLM}_\theta(a \mid h_t, s_t, \mathcal{M}, \mathcal{E}^{ret}). \quad (3)$$

Here, h_t denotes the current strategic plan, and $a_t = \langle \tau, \mathbf{p} \rangle$ specifies the executable graph operation, where $\tau \in \mathcal{T}$ is the invoked tool and $\mathbf{p}$ specifies the target and arguments. This two-stage process allows the agent to revise its strategy before each manipulation step.

Social-context observation. Observation tools $\mathcal{T}_{obs} : \mathcal{G} \rightarrow \mathcal{S}_{text}$ project graph states into textual descriptions interpretable by the LLM. As detailed in Table 1, these tools support global and local perception, helping the agent assess the social environment and locate valuable intervention targets. Since message propagation trees (MPTs) and news engagement graphs (NEGs) encode different propagation semantics, the tools also expose graph-type-specific information for strategy updating.

Behavior execution. Execution tools $\mathcal{T}_{exe}$ instantiate the selected behavior as a graph modification $G_{t+1} \leftarrow \tau(G_t, \mathbf{p})$. As shown in Table 1, the action space covers multiple behaviors, such as editing content, modifying or deleting comments, injecting replies, and perturbing relations. Because these operations may have different effects under MPTs and NEGs, the orchestrator grounds each action in the current graph type and social context, converting manipulation intent into executable, graph-aware social behaviors.

3.3 Reflective Experience Memory

While the orchestrator grounds manipulation decisions in social context, effective evasion further requires the agent to refine strategies from detector feedback and historical attack attempts. To this end, we design Reflective Experience Memory

(REM), which supports two coupled processes: reflective strategy refinement during the current attack and experience distillation and retrieval across historical attacks. Unlike general agent reflection frameworks (Yao et al., 2023; Shinn et al., 2023), REM converts detector-feedback-conditioned trajectories into reusable graph-aware manipulation experience.

Reflective Strategy Refinement. At each execution step t , the agent applies a behavior $a_t \in \mathcal{T}_{exe}$ and obtains G_{t+1} . The surrogate detector f_θ provides feedback:

$$r_t = c_{f_\theta}(G_t) - c_{f_\theta}(G_{t+1}), \quad (4)$$

where $c_{f_\theta}(\cdot)$ measures the confidence for the fake news class. A positive r_t indicates that the manipulation moves the instance toward the benign side from the surrogate detector’s perspective. EvoSM maps the feedback and historical rewards into a reflection state:

$$z_t \in \{\text{“Success”}, \text{“Failure”}, \text{“Deadlock”}\}, \quad (5)$$

and stores it in the local trajectory memory $\mathcal{M}_t$. The next strategy and action are generated as

$$h_{t+1} \sim \text{LLM}_\theta(h \mid s_{t+1}, \mathcal{M}_t, \mathcal{E}^{ret}), \quad (6)$$

$$a_{t+1} \sim \text{LLM}_\theta(a \mid h_{t+1}, s_{t+1}, \mathcal{M}_t, \mathcal{E}^{ret}), \quad (7)$$

where s_{t+1} is the updated textualized context and $\mathcal{E}^{ret}$ denotes the retrieved experience from REM. Thus, surrogate feedback is transformed into reflection signals that guide the ongoing behavior sequence rather than serving only as final evaluation.

Experience Distillation and Retrieval. Beyond the current attack attempt, REM accumulates reusable experience from historical attack trajectories. After an attack attempt terminates, its completed trajectory is distilled into an experience entry:

$$e_i = \text{Distill}(\tau_i, \phi(G_i), o_i), \quad (8)$$

where τ_i is the completed attack trajectory, $\phi(G_i)$ denotes a compact social-context profile, and o_i is the final outcome. The profile summarizes the graph type (MPT or NEG), key propagation statistics, influential regions, and behavior-feedback patterns that indicate which actions changed detector confidence.

Based on this profile, REM summarizes successful trajectories as positive skills, failed or

budget-exhausted trajectories as negative lessons, and locally effective steps as micro-skills. These experiences capture which manipulation behaviors, targets, or relations are more effective under specific graph types and propagation contexts. For a new target, REM retrieves relevant entries by social-context similarity:

$$\mathcal{E}_i^{ret} = \text{Retrieve}(\mathcal{E}_{i-1}, \phi(G_i)). \quad (9)$$

The retrieved experience is injected as soft guidance rather than fixed action scripts, enabling EvoSM to adapt historical graph-aware manipulation skills to the current propagation structure.

4 Experiments

In this section, we aim to address the following key research questions through experiments:

RQ1: How effective is EvoSM in attacking propagation-based fake news detectors?

RQ2: How effective are SMO and REM in overall performance?

RQ3: Can EvoSM autonomously evolve complex manipulation strategies?

RQ4: How do the propagation patterns change after EvoSMs attack?

RQ5: How can we defend against such autonomously evolving AI adversaries?

4.1 Experiment Setups

Datasets. We conduct experiments on four public fake news datasets: **Twitter** (Lin et al., 2022), **CED** (Song et al., 2019), **Politifact**, and **Gossipcop** (Shu et al., 2020). Consistent with the definition in Section 2, we categorize these datasets into two scenarios: **1) MPT-based Scenarios:** We utilize Twitter and CED to represent Message Propagation Trees, where the structure captures the hierarchical diffusion of posts; **2) NEG-based Scenarios:** We employ Politifact and Gossipcop to represent News Engagement Graphs, modeling global user-news relations. Details are provided in the Appendix C.1.

Detectors. We evaluate EvoSM against a diverse set of widely used detectors. We select three GNN backbones: GCN (Kipf and Welling, 2017), GAT (Veličković et al., 2018), and GraphSAGE (Hamilton et al., 2017), which serve as fundamental baselines. We further include strong and robustness-oriented detectors. For MPT-based scenarios, we employ Bi-GCN (Bian et al., 2020)

Method	Twitter					CED				
	GAT	GCN	SAGE	Bi-GCN	SINCon	GAT	GCN	SAGE	Bi-GCN	SINCon
	ASR Δ_c	ASR Δ_c	ASR Δ_c	ASR Δ_c	ASR Δ_c	ASR Δ_c	ASR Δ_c	ASR Δ_c	ASR Δ_c	ASR Δ_c
HMIA-LLM (Luo et al., 2024)	0.06 0.06	0.17 0.09	0.02 −0.06	0.10 0.02	0.22 0.20	0.10 0.10	0.45 0.36	0.05 0.07	0.24 0.22	0.52 0.47
GAFSI (Zhu et al., 2024)	0.34 0.30	0.09 0.10	0.35 0.24	0.17 0.11	0.25 0.22	0.33 0.29	0.04 0.05	0.10 0.13	0.05 0.04	0.06 0.07
Generate (Wan et al., 2025)	0.05 0.04	0.13 0.10	0.01 −0.07	0.05 −0.09	0.18 0.13	0.03 0.04	0.17 0.14	0.01 0.01	0.06 0.05	0.20 0.17
SALF (Tian et al., 2025)	0.11 0.07	0.20 0.15	0.20 0.09	0.10 0.03	0.25 0.22	0.25 0.20	0.53 0.39	0.28 0.27	0.45 0.39	0.46 0.39
Rephrase (Wan et al., 2025)	0.25 0.19	0.07 0.08	0.24 0.13	0.06 0.05	0.13 0.09	0.23 0.21	0.05 0.06	0.04 0.07	0.01 0.01	0.02 0.02
Delete (Wan et al., 2025)	0.05 0.02	0.03 0.03	0.07 0.04	0.02 −0.01	0.05 0.01	0.04 0.06	0.04 0.05	0.02 0.04	0.03 0.02	0.00 0.00
EvoSM (Ours)	0.36 0.32	0.48 0.40	0.41 0.25	0.21 0.12	0.31 0.26	0.42 0.38	0.73 0.61	0.50 0.45	0.52 0.49	0.58 0.52

Table 2: Attack performance on MPT datasets. For each attack method, we evaluate it with different detectors and report the average ASR and confidence drop Δ_c over five runs. The best results are in **bold**.

Method	Politifact					Gossipcop				
	GAT	GCN	SAGE	DECOR	LLM4RGNN	GAT	GCN	SAGE	DECOR	LLM4RGNN
	ASR Δ_c	ASR Δ_c	ASR Δ_c	ASR Δ_c	ASR Δ_c	ASR Δ_c	ASR Δ_c	ASR Δ_c	ASR Δ_c	ASR Δ_c
Nettack (Zügner et al., 2018)	0.49 0.48	0.66 0.64	0.56 0.37	0.31 0.31	0.47 0.38	0.34 0.28	0.31 0.28	0.06 0.06	0.44 0.19	0.14 0.06
MARL (Wang et al., 2023)	0.40 0.39	0.41 0.38	0.27 0.11	0.22 0.35	0.69 0.60	0.12 0.10	0.17 0.15	0.02 0.02	0.27 0.15	0.15 0.07
GAFSI (Zhu et al., 2024)	0.51 0.50	0.35 0.33	0.36 0.20	0.23 0.23	0.43 0.36	0.16 0.14	0.13 0.12	0.01 0.02	0.43 0.17	0.16 0.07
WTGIA (Lei et al., 2024)	0.36 0.35	0.60 0.57	0.29 0.14	0.24 0.25	0.24 0.22	0.22 0.19	0.28 0.26	0.05 0.04	0.37 0.16	0.34 0.15
SI2AF (Zeng et al., 2025)	0.79 0.75	0.58 0.52	0.52 0.34	0.36 0.42	0.68 0.53	0.08 0.06	0.18 0.20	0.02 0.02	0.27 0.14	0.08 0.03
Rephrase (Wan et al., 2025)	0.29 0.29	0.32 0.31	0.58 0.41	0.32 0.31	0.26 0.21	0.43 0.37	0.38 0.33	0.09 0.08	0.39 0.16	0.16 0.07
SALF (Tian et al., 2025)	0.44 0.43	0.42 0.41	0.72 0.53	0.30 0.29	0.24 0.23	0.36 0.30	0.39 0.32	0.06 0.05	0.23 0.05	0.26 0.09
EvoSM (Ours)	0.85 0.82	0.79 0.76	0.91 0.69	0.77 0.75	0.82 0.66	0.45 0.37	0.46 0.41	0.15 0.13	0.51 0.20	0.41 0.19

Table 3: Attack performance on NEG datasets. For each attack method, we evaluate it with different detectors and report the average ASR and confidence drop Δ_c over five runs. The best results are in **bold**.

and SINCon (Zhang et al., 2025a). For NEG-based scenarios, we utilize DECOR (Wu and Hooi, 2023) and LLM4RGNN (Zhang et al., 2025b). Details of these detectors are provided in Appendix C.2.

Baselines. In MPT-based scenarios, we compare with baselines that include adding comments (HMIA-LLM (Luo et al., 2024), GAFSI (Zhu et al., 2024), Generate (Wan et al., 2025)), editing news (SALF (Tian et al., 2025)) and comments (Rephrase (Wan et al., 2025)), and deleting comments (Delete (Wan et al., 2025)). In NEG-based scenarios, we compare with baselines that include adding edges (GAFSI (Zhu et al., 2024), MARL (Wang et al., 2023), SI2AF (Zeng et al., 2025)), adding news nodes (WTGIA (Lei et al., 2024)), editing news (SALF (Tian et al., 2025)), and removing edges (Nettack (Zügner et al., 2018)). Details are provided in Appendix C.3.

Implementation Details. All experiments are conducted on a single NVIDIA L40 GPU. We use a 2-layer GCN as the surrogate model, a common choice in black-box graph adversarial attacks due to its transferability (Zügner et al., 2018; Li et al., 2021; Zhu et al., 2024). For fair comparison, we set the manipulation budget to 30 execution steps for all methods, while social-context observations do not consume this budget. The attack target set contains fake news instances correctly classified by both the surrogate and target detectors, and an attack is successful if the target detector’s predic-

tion flips to real. We evaluate two LLMs, GPT-4o and Qwen3-4B-Instruct². Unless otherwise specified, LLM-based attack and detection results are reported with GPT-4o. More details are provided in Appendix C.5.

Evaluation Metrics. We evaluate attack effectiveness using Attack Success Rate (ASR) and the fake-class confidence drop Δ_c . ASR measures the proportion of target fake-news instances that are misclassified as real after manipulation, while Δ_c measures the average decrease in the target detector’s confidence for the fake-news class, providing a continuous measure of attack strength even when the prediction does not flip. Since all attack targets are fake-news instances, Δ_c is also equivalent to the increase in class-wise ECE computed on the same target set; therefore, a larger Δ_c indicates both stronger suppression of fake-class confidence and greater target-set calibration degradation.³

4.2 Overall Performance

To answer **RQ1**, we evaluate EvoSM against SOTA attacks across four datasets and multiple propagation-based fake news detectors. As shown in Table 2 and Table 3, EvoSM achieves the best overall attack performance, delivering top ASR performance across all evaluated detector settings and generally inducing the largest confidence drops. Notably, EvoSM maintains strong transfer-

²We compare the two backbones in Appendix B.6.

³The algebraic equivalence between Δ_c and the change in class-wise ECE on the attack target set is detailed in Appendix C.4.

Variant	MPT-based		NEG-based	
	Target: GCN		Target: GCN	
	ASR	Δ_c	ASR	Δ_c
EvoSM	0.61	0.51	0.61	0.57
<i>SMO</i>				
<i>w/o Observation</i>	0.55	0.46	0.42	0.38
EvoSM _{Add}	0.37	0.32	0.37	0.33
<i>REM</i>				
<i>ReAct</i>	0.40	0.33	0.26	0.24
<i>w/o Experience</i>	0.48	0.41	0.40	0.37
<i>w/o Reflection</i>	0.42	0.36	0.41	0.38

Table 4: Ablation study of EvoSM on MPT-based and NEG-based detection datasets. *w/o Observation* removes social-context observation; EvoSM_{Add} allows only additive actions; *ReAct* replaces REM with ReAct-style reasoning; *w/o Experience* removes historical experience; and *w/o Reflection* removes detector-feedback reflection. The best performance is in **bold**. Following Wan et al. (2025), we sample 100 attack targets from each dataset and average results over 5 runs.




Figure 2: Temporal distribution of observation tool usage across different datasets. Strength represents the overall observation frequency at the current attack step.

ability against robustness-oriented detectors such as SINCon and LLM4RGNN, where existing attacks generally exhibit lower attack effectiveness. This advantage stems from the adaptive nature of our framework: instead of relying on a fixed manipulation pattern, EvoSM dynamically coordinates semantic and structural behaviors according to surrogate feedback. By composing multi-stage manipulation campaigns and progressively revising its strategy, EvoSM can adapt its attack focus to different propagation contexts and detector characteristics. Moreover, the additive-only analysis in Appendix B.1 shows that EvoSM retains non-trivial attack capability under a restricted action space, while the performance gap to full EvoSM highlights the benefit of coordinating heterogeneous social behaviors.




Figure 3: Temporal evolution of EvoSM agent action preferences at each attack step across different datasets.

4.3 Ablation Study

To answer RQ2, we evaluate component contributions by comparing EvoSM with controlled variants of the orchestrator and REM (Table 4). Regarding the **SMO**, the drop in *w/o Observation* shows that limited social-context observation weakens adaptive manipulation, while constraining actions to additive-only operations (EvoSM_{Add}) confirms the necessity of coordinating multiple social behaviors. Regarding **REM**, the inferior performance of *ReAct* indicates that generic ReAct-style tool use alone cannot replace attack-specific experience evolution. Meanwhile, the drops in *w/o Experience* and *w/o Reflection* show that both reusable experience and detector-feedback-guided reflection are important for evolving manipulation strategies. Overall, these results demonstrate that EvoSM’s effectiveness stems from the synergy between the orchestrator’s context-grounded behavior execution and REM’s experience-guided strategy evolution.

4.4 Strategic Behavior Analysis

To answer RQ3, we analyze EvoSM’s observation and action patterns by visualizing the frequency distribution of observation and execution tool usage across attack steps. As shown in Figure 2 and Figure 3, EvoSM first prioritizes global perception tools (e.g., GetGlobalStats) to grasp the topology, along with foundational actions such as Edit News or Add News to establish a semantic baseline.

As the attack progresses, its strategy shifts toward more precise manipulation. In MPT-based




Figure 4: An evolutionary attack trajectory of EvoSM on the Twitter dataset. (a-b) Semantic Adaptation: Initially, it rewrites content to dilute fake semantic signals. (c) Structural Expansion: It then pivots to add comment chains, simulating organic discussions typical of real news. (d) Context Refinement: Finally, it removes noisy comments to ensure contextual plausibility, successfully evading the detector.

scenarios, Add Comment remains a dominant operation, while Edit News and Edit Comment are dynamically invoked to adjust semantic cues; in NEG-based scenarios, it evolves from network expansion to topology refinement, revealing its evolutionary strategic behavior. Moreover, EvoSM still invokes some actions with limited immediate rewards when they help prepare later manipulation steps, suggesting that the agent optimizes the campaign trajectory rather than greedily selecting only short-term high-reward actions (Appendix B.3). Moreover, REM further compresses such campaigns into reusable experience (Appendix B.8). The trajectory case in Figure 4 further shows a multi-stage campaign: semantic adaptation to dilute fake signals, structural injection to simulate organic discussions, and comment pruning for context refinement. Overall, these observations show that EvoSM acts as a manipulation strategist that autonomously orchestrates evolutionary attack campaigns.

4.5 Analysis of Manipulated Propagation

To answer **RQ4**, we analyze the semantic and structural properties of propagation graphs after manipulation by EvoSM. For semantics, we analyze node-pair similarity distributions. As shown in Figure 5, EvoSM largely preserves the original similarity patterns and maintains more stable relationships among propagated contents than baselines that introduce larger distributional shifts or low-similarity outliers. Qualitative examples in Appendix B.4 further show that EvoSM introduces credibility-oriented wording while preserv-




Figure 5: Node-pair semantic cosine similarity (computed via BERT embeddings and visualized using kernel density estimation) distributions before and after attacks across different datasets.

Dataset	Metric	EvoSM	Baseline1	Baseline2
Twitter	$\Delta\alpha \downarrow$	0.34	1.53	2.56
	KS $\downarrow$	0.034	0.081	0.097
	JS $\downarrow$	0.020	0.054	0.055
CED	$\Delta\alpha \downarrow$	0.03	0.13	0.61
	KS $\downarrow$	0.003	0.008	0.024
	JS $\downarrow$	0.000	0.005	0.005
Politifact	$\Delta\alpha \downarrow$	0.06	0.59	0.65
	KS $\downarrow$	0.079	0.20	0.21
	JS $\downarrow$	0.026	0.11	0.10
Gossipcop	$\Delta\alpha \downarrow$	0.18	1.07	0.36
	KS $\downarrow$	0.138	0.30	0.21
	JS $\downarrow$	0.094	0.15	0.11

Table 5: Structural preservation metrics; lower is better. Baseline1: HMA-LLM (Twitter/CED), MARL (Politifact/Gossipcop). Baseline2: GAFSI (Twitter/CED), SI2AF (Politifact/Gossipcop). Metrics for Politifact and Gossipcop are computed from the 2-hop ego-subgraphs.

ing the overall semantic context. For structure, we measure topological changes using **Power-law exponent difference** ($\Delta\alpha$), **Kolmogorov–Smirnov (KS) statistic**, and **Jensen–Shannon (JS) divergence**.⁴ As shown in Table 5, EvoSM generally yields lower deviations than baselines across MPT and NEG scenarios, suggesting effective evasion with better preservation of the original propagation topology.

4.6 Defense Discussion

To address **RQ5**, we conduct a controlled defense evaluation using a fixed detector checkpoint

⁴For Politifact and Gossipcop, these metrics are computed within 2-hop ego-subgraphs to capture local topological impacts and avoid dilution by global signals.




Figure 6: Effect of the LLM-based guardrail defense under a fixed detector checkpoint. We report the attack success rate before and after applying the guardrail, together with the Attack Interception Rate (AIR↑).

and target set. We introduce an **Intent-Aware Guardrail** (Appendix C.8), in which an LLM-based moderator examines the strategic intent behind each agent action a_t and blocks actions identified as malicious. We define the **Attack Interception Rate (AIR)** as the proportion of attempted actions blocked by the guardrail. As shown in Figure 6, the guardrail intercepts a substantial fraction of malicious actions and markedly reduces EvoSM’s ASR, demonstrating that intent-aware LLM moderation can complement structure-based fake news detectors. Nevertheless, EvoSM remains more effective than the evaluated baselines under this controlled defense setting, highlighting the need for more adaptive safeguards against evolving social manipulation.

5 Conclusion

This work highlights the threat of LLM-driven self-evolving social manipulation against propagation-based fake news detection. We introduce EvoSM, a unified attack framework that utilizes LLMs to orchestrate evolutionary social manipulation campaigns. EvoSM enables the agent to perceive propagation contexts, translate manipulation intents into executable social graph operations, and refine future strategies by reflecting on detector feedback and distilling reusable experience from historical attack trajectories. Experiments across four datasets and diverse propagation-based detectors demonstrate that EvoSM delivers strong attack effectiveness and transferability, achieving the best overall performance against state-of-the-art baselines. Our findings reveal critical vulnerabilities in existing fake news detection systems, underscoring the need for adaptive defenses against evolutionary social manipulation.

Acknowledgment

This work was supported by the National Natural Science Foundation of China (No.U24A20335 and No.62606518), the China Postdoctoral Science Foundation (No. 2024M753481), and the Youth Innovation Promotion Association CAS. The authors thank the anonymous reviewers and the metareviewer for their helpful comments.

Limitations

While EvoSM demonstrates strong attack capabilities, several limitations remain and warrant further investigation.

First, EvoSM incurs non-negligible computational overhead due to multi-step LLM reasoning, tool invocation, and feedback-guided refinement. Although Appendix B.6 shows that EvoSM can be instantiated with an open-source backbone such as Qwen3-4B-Instruct, a performance gap remains relative to the GPT-based backbone. Reducing this computational overhead while preserving strong agentic reasoning and manipulation capabilities therefore requires further improvements in lightweight LLM backbones and tool-use reliability.

Second, our current threat model assumes that the adversary can observe the complete propagation graph of the target news. In real-world platforms, however, graph visibility may be partial, delayed, or restricted by platform APIs, privacy policies, or access constraints. Such limited observability may weaken context-aware manipulation and reduce the effectiveness of experience retrieval. Evaluating EvoSM under partially observable or dynamically revealed propagation environments is therefore an important direction for future work.

Third, the dynamics of attack–defense co-evolution remain underexplored. Although we evaluate a preliminary LLM-based guardrail defense, stronger adversaries may employ adversarial prompting, multi-turn obfuscation, or indirect intent hiding to evade intent-aware defenses. Future work should investigate adaptive defense mechanisms that can anticipate, monitor, and intercept multi-step manipulation campaigns, thereby enabling more robust safeguards against evolving adversarial strategies.

Finally, our experiments focus primarily on evasion attacks against existing detectors. Extending the evaluation to poisoning settings, where adver-

saries manipulate training data or long-term engagement patterns to degrade future detector performance, would provide a broader understanding of the long-term security risks posed by evolving social manipulation.

Ethics Statement

This research aims to advance the understanding of vulnerabilities in propagation-based fake news detection systems. EvoSM is developed as a controlled red-teaming framework to expose weaknesses in current defenses, but its manipulation strategies may pose misuse risks if applied to real platforms. All experiments were conducted on publicly available datasets in isolated environments, with no interaction with real users or live social media systems. We report these findings to support the development of stronger adaptive safeguards and emphasize that EvoSM should only be used for defensive research and educational purposes, in compliance with relevant laws, ethical guidelines, and applicable model-use policies. To reduce potential misuse risks, we intentionally avoid presenting platform-specific automation procedures, ready-to-use malicious prompts, or operational instructions for deploying EvoSM on real platforms. The experimental details are reported only to the extent necessary for understanding the offline robustness evaluation on benchmark datasets.

References

- Mazal Bethany, Nishant Vishwamitra, Cho-Yu Jason Chiang, and Peyman Najafirad. 2025. Camouflage: Exploiting misinformation detection systems through llm-driven adversarial claim transformation. *arXiv preprint arXiv:2505.01900*.
- Tian Bian, Xi Xiao, Tingyang Xu, Peilin Zhao, Wenbing Huang, Yu Rong, and Junzhou Huang. 2020. Rumor detection on social media with bi-directional graph convolutional networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 549–556.
- Carlos Castillo, Marcelo Mendoza, and Barbara Poblete. 2011. Information credibility on twitter. In *Proceedings of the 20th international conference on World wide web*, pages 675–684.
- Mengyang Chen, Lingwei Wei, Han Cao, Wei Zhou, Zhou Yan, and Songlin Hu. 2024. Propagation structure-semantic transfer learning for robust fake news detection. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 227–244. Springer.
- Mengyang Chen, Lingwei Wei, Wei Zhou, and Songlin Hu. 2026. Mitigating adversarial attacks by transferring llm-generated narrative reasoning for robust fake news detection. In *Proceedings of the 49th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '26, page 190201, New York, NY, USA. Association for Computing Machinery.
- Yingtong Dou, Kai Shu, Congying Xia, Philip S Yu, and Lichao Sun. 2021. User preference-aware fake news detection. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, pages 2051–2055.
- Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. 2024. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*.
- Marc Fisher, John Woodrow Cox, and Peter Hermann. 2016. Pizzagate: From rumor, to hashtag, to gunfire in D.C. The Washington Post.
- Suhaib Kh Hamed, Mohd Juzaidin Ab Aziz, and Mohd Ridzwan Yaakub. 2023. Fake news detection model on social media by leveraging sentiment analysis of news content and emotion analysis of users comments. *Sensors*, 23(4):1748.
- Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *NIPS*, 30.
- Dou Hu, Lingwei Wei, Wei Zhou, Xiaoyong Huai, Jizhong Han, and Songlin Hu. 2021. A rumor detection approach based on multi-relational propagation tree. *Journal of Computer Research and Development*, 58(7):1395–1411.
- Heejung Jwa, Dongsuk Oh, Kinam Park, Jang Mook Kang, and Heuseok Lim. 2019. exbake: Automatic fake news detection model based on bidirectional encoder representations from transformers (bert). *Applied Sciences*, 9(19):4062.
- Rohit Kumar Kaliyar, Anurag Goswami, and Pratik Narang. 2021. Fakebert: Fake news detection in social media with a bert-based deep learning approach. *Multimedia tools and applications*, 80(8):11765–11788.
- Hamid Karimi and Jiliang Tang. 2019. Learning hierarchical discourse-level structure for fake news detection. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3432–3442.
- Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*.

- Runlin Lei, Yuwei Hu, Yuchen Ren, and Zhewei Wei. 2024. Intruding with words: Towards understanding graph injection attacks at the text level. *Advances in Neural Information Processing Systems*, 37:49214–49251.
- Jintang Li, Tao Xie, Liang Chen, Fenfang Xie, Xiangnan He, and Zibin Zheng. 2021. Adversarial attack on large scale graph. *IEEE Transactions on Knowledge and Data Engineering*, 35(1):82–95.
- Hongzhan Lin, Jing Ma, Liangliang Chen, Zhiwei Yang, Mingfei Cheng, and Chen Guang. 2022. Detect rumors in microblog posts for low-resource domains via adversarial contrastive learning. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2543–2556, Seattle, United States. Association for Computational Linguistics.
- Bang Liu, Xinfeng Li, Jiayi Zhang, Jinlin Wang, Tanjin He, Sirui Hong, Hongzhang Liu, Shaokun Zhang, Kaitao Song, Kunlun Zhu, and 1 others. 2025. Advances and challenges in foundation agents: From brain-inspired intelligence to evolutionary, collaborative, and safe systems. *arXiv preprint arXiv:2504.01990*.
- Yang Liu and Yi fang Brook Wu. 2018. Early detection of fake news on social media through propagation path classification with recurrent and convolutional networks. In *AAAI*.
- Yi-Ju Lu and Cheng te Li. 2020. Gcan: Graph-aware co-attention networks for explainable fake news detection on social media. In *Annual Meeting of the Association for Computational Linguistics*.
- Junyu Luo, Weizhi Zhang, Ye Yuan, Yusheng Zhao, Junwei Yang, Yiyang Gu, Bohan Wu, Binqi Chen, Ziyue Qiao, Qingqing Long, and 1 others. 2025. Large language model agent: A survey on methodology, applications and challenges. *arXiv preprint arXiv:2503.21460*.
- Yifeng Luo, Yupeng Li, Dacheng Wen, and Liang Lan. 2024. Message injection attack on rumor detection under the black-box evasion setting using large language model. In *Proceedings of the ACM Web Conference 2024*, pages 4512–4522.
- Alex Munyole Luvembe, Weimin Li, Shaohua Li, and 1 others. 2023. Dual emotion based fake news detection: A deep attention-weight update approach. *IPM*, 60(4):103354.
- Guanghui Ma, Chunming Hu, Ling Ge, Junfan Chen, Hong Zhang, and Richong Zhang. 2022. Towards robust false information detection on social networks with contrastive learning. In *Proceedings of the 31st ACM international conference on information & knowledge management*, pages 1441–1450.
- Jing Ma, Wei Gao, Prasenjit Mitra, Sejeong Kwon, Bernard Jim Jansen, Kam-Fai Wong, and M. Cha. 2016. Detecting rumors from microblogs with recurrent neural networks. In *International Joint Conference on Artificial Intelligence*.
- Jing Ma, Wei Gao, Zhongyu Wei, Yueming Lu, and Kam-Fai Wong. 2015. Detect rumors using time series of social context information on microblogging websites. In *Proceedings of the 24th ACM international on conference on information and knowledge management*, pages 1751–1754.
- Jing Ma, Wei Gao, and Kam-Fai Wong. 2018. Rumor detection on twitter with tree-structured recursive neural networks. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 1980–1989. Association for Computational Linguistics.
- Diogo Pacheco, Pik-Mai Hui, Christopher Torres-Lugo, Bao Tran Truong, Alessandro Flammini, and Filippo Menczer. 2021. Uncovering coordinated networks on social media: methods and case studies. In *Proceedings of the international AAAI conference on web and social media*, volume 15, pages 455–466.
- Sungwon Park, Sungwon Han, Xing Xie, Jae-Gil Lee, and Meeyoung Cha. 2025. Adversarial style augmentation via large language model for robust fake news detection. In *Proceedings of the ACM on Web Conference 2025*, pages 4024–4033.
- Kashyap Popat. 2017. Assessing the credibility of claims on the web. In *Proceedings of the 26th International Conference on World Wide Web Companion, WWW '17 Companion*, page 735739, Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.
- Natali Ruchansky, Sungyong Seo, and Yan Liu. 2017. Csi: A hybrid deep model for fake news detection. In *CIKM*, pages 797–806.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36:8634–8652.
- Kai Shu, Limeng Cui, Suhang Wang, Dongwon Lee, and Huan Liu. 2019. defend: Explainable fake news detection. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 395–405.
- Kai Shu, Deepak Mahudeswaran, Suhang Wang, and 1 others. 2020. Fakenewsnet: A data repository with news content, social context, and spatiotemporal information for studying fake news on social media. *Big data*, 8(3):171–188.
- Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 2017. Fake news detection on social media: A data mining perspective. *ACM SIGKDD explorations newsletter*, 19(1):22–36.

- Changhe Song, Cheng Yang, Huimin Chen, Cunchao Tu, Zhiyuan Liu, and Maosong Sun. 2019. Ced: Credible early detection of social media rumors. *IEEE Transactions on Knowledge and Data Engineering*, 33(8):3035–3047.
- Xing Su, Jian Yang, Jia Wu, and Yuchen Zhang. 2023. Mining user-aware multi-relations for fake news detection in large scale online social networks. In *Proceedings of the sixteenth ACM international conference on web search and data mining*, pages 51–59.
- Tiening Sun, Zhong Qian, Sujun Dong, Peifeng Li, and Qiaoming Zhu. 2022. Rumor detection on social media with graph adversarial contrastive learning. In *Proceedings of the ACM Web Conference 2022*, pages 2789–2797.
- Chong Tian, Qirong Ho, and Xiuying Chen. 2025. A symbolic adversarial learning framework for evolving fake news generation and detection. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 12307–12321.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. In *International Conference on Learning Representations*.
- Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *science*, 359(6380):1146–1151.
- Herun Wan, Minnan Luo, Zhixiong Su, Guang Dai, and Xiang Zhao. 2025. On the risk of evidence pollution for malicious social text detection in the era of llms. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9731–9761.
- Haoran Wang, Yingdong Dou, Canyu Chen, Lichao Sun, Philip S Yu, and Kai Shu. 2023. Attacking fake news detectors via manipulating news social engagement. In *Proceedings of the ACM Web Conference 2023*, pages 3978–3986.
- Rongzheng Wang, Shuang Liang, Qizhi Chen, Jiasheng Zhang, and Ke Qin. 2025. Graphtool-instruction: Revolutionizing graph reasoning in llms through decomposed subtask instruction. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1*, pages 1492–1503.
- Lingwei Wei, Dou Hu, Wei Zhou, and Songlin Hu. 2022. Uncertainty-aware propagation structure reconstruction for fake news detection. In *Proceedings of the 29th international conference on computational linguistics*, pages 2759–2768.
- Lingwei Wei, Dou Hu, Wei Zhou, and Songlin Hu. 2024. Transferring structure knowledge: A new task to fake news detection towards cold-start propagation. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8045–8049. IEEE.
- Lingwei Wei, Dou Hu, Wei Zhou, Songlin Hu, and Philip S Yu. 2026. Large language models in misinformation ecosystems: Misuse, defense, and vulnerability. *arXiv preprint arXiv:2607.10402*.
- Lingwei Wei, Dou Hu, Wei Zhou, Zhaojuan Yue, and Songlin Hu. 2021. Towards propagation uncertainty: Edge-enhanced bayesian graph convolutional networks for rumor detection. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3845–3854.
- Jiaying Wu, Jiafeng Guo, and Bryan Hooi. 2024. Fake news in sheep’s clothing: Robust fake news detection against llm-empowered style attacks. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 3367–3378.
- Jiaying Wu and Bryan Hooi. 2023. Decor: Degree-corrected social graph refinement for fake news detection. In *KDD*, pages 2582–2593.
- Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, and 1 others. 2025. The rise and potential of large language model based agents: A survey. *Science China Information Sciences*, 68(2):121101.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2023. [React: Synergizing reasoning and acting in language models](#). In *The Eleventh International Conference on Learning Representations*.
- Asaf Yehudai, Lilach Eden, Alan Li, Guy Uziel, Yilun Zhao, Roy Bar-Haim, Arman Cohan, and Michal Shmueli-Scheuer. 2025. Survey on evaluation of llm-based agents. *arXiv preprint arXiv:2503.16416*.
- Feng Yu, Qiang Liu, Shu Wu, Liang Wang, Tieniu Tan, and 1 others. 2017. A convolutional approach for misinformation identification. In *IJCAI*, volume 2017, pages 3901–3907.
- Chunyuan Yuan, Qianwen Ma, Wei Zhou, Jizhong Han, and Songlin Hu. 2020. Early detection of fake news by utilizing the credibility of news, publishers, and users based on weakly supervised learning. In *Proceedings of the 28th international conference on computational linguistics*, pages 5444–5454.
- Xianghua Zeng, Hao Peng, and Angsheng Li. 2025. Robustness evaluation of graph-based news detection using network structural information. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pages 3716–3727.
- Guibin Zhang, Hejia Geng, Xiaohang Yu, Zhenfei Yin, Zaibin Zhang, Zelin Tan, Heng Zhou, Zhong-Zhi Li, Xiangyuan Xue, Yijiang Li, Yifan Zhou, Yang Chen, Chen Zhang, Yutao Fan, Zihu Wang, Songtao

- Huang, Francisco Piedrahita Velez, Yue Liao, Hongru WANG, and 6 others. 2026. [The landscape of agentic reinforcement learning for LLMs: A survey](#). *Transactions on Machine Learning Research*. Survey Certification.
- Jiawei Zhang. 2023. Graph-toolformer: To empower llms with graph reasoning ability via prompt augmented by chatgpt. *arXiv preprint arXiv:2304.11116*.
- Mingqing Zhang, Qiang Liu, Xiang Tao, Shu Wu, and Liang Wang. 2025a. [SINCon: Mitigate LLM-generated malicious message injection attack for rumor detection](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12570–12581, Vienna, Austria. Association for Computational Linguistics.
- Xueyao Zhang, Juan Cao, Xirong Li, Qiang Sheng, Lei Zhong, and Kai Shu. 2021. Mining dual emotion for fake news detection. In *Proceedings of the web conference 2021*, pages 3465–3476.
- Zhongjian Zhang, Xiao Wang, Huichi Zhou, Yue Yu, Mengmei Zhang, Cheng Yang, and Chuan Shi. 2025b. [Can large language models improve the adversarial robustness of graph neural networks?](#) In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1*, KDD '25, page 20082019, New York, NY, USA. Association for Computing Machinery.
- Peican Zhu, Zechen Pan, Yang Liu, Jiwei Tian, Keke Tang, and Zhen Wang. 2024. A general black-box adversarial attack on graph-based fake news detectors. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pages 568–576.
- Daniel Zügner, Amir Akbarnejad, and Stephan Günnemann. 2018. Adversarial attacks on neural networks for graph data. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2847–2856.

In this supplementary material, we provide detailed supporting information to substantiate the findings in the main paper. The organization is as follows: we first review the related work in **Appendix A**. Next, we present comprehensive additional experiments in **Appendix B**, including detailed ablation studies, efficiency analysis, and case studies. Finally, we provide detailed experimental setups, such as dataset statistics, baseline configurations, and the complete system prompts for EvoSM, in **Appendix C**.

A Related Work

Fake News Detection. The goal of detecting fake news is to identify and assess the authenticity of a piece of information. Existing methods for detecting fake news mainly focus on two aspects: textual content and propagation of news.

Content-based Fake News Detection Methods extract semantic patterns from news content for detection through feature engineering (Castillo et al., 2011; Popat, 2017; Ma et al., 2015) and a wide array of deep learning architectures, including neural networks (Ruchansky et al., 2017; Karimi and Tang, 2019) and pre-trained language models (Kaliyar et al., 2021; Jwa et al., 2019). Some works also integrate tasks such as stance detection and sentiment analysis with fake news detection, enabling multi-task learning (Luvembe et al., 2023; Hamed et al., 2023). Since some fake news creators imitate the style of real news, methods based solely on news content often face limitations. Consequently, some researchers use news comments as a basis for assessing the authenticity of news (Shu et al., 2019; Zhang et al., 2021).

Propagation-based Fake News Detection Methods capture the propagation patterns of news by modeling the interactions between news and comments into time series (Ma et al., 2016; Liu and fang Brook Wu, 2018) or topological structures such as propagation trees (Ma et al., 2018; Hu et al., 2021) and propagation graphs (Bian et al., 2020; Wei et al., 2021, 2022). Some studies further explore multi-relational interactions between the users and news in the propagation graph (Yuan et al., 2020; Dou et al., 2021). However, these methods suffer significant performance losses when confronted with scenarios of incomplete propagation (Wei et al., 2024; Ma et al., 2022).

Adversarial Attacks on Fake News Detection.

Adversarial attacks on fake news detection have evolved from manual manipulation (Wei et al., 2021; Sun et al., 2022; Ma et al., 2022) to automated methods. Gradient-based and reinforcement learning approaches (Wang et al., 2023; Zhu et al., 2024; Zeng et al., 2025) demonstrate feasibility but suffer from a lack of semantic coherence and rigid optimization that fails to mimic realistic human behavior. The emergence of LLMs has enabled scalable textual adversarial attacks, including malicious message injection (Luo et al., 2024; Wan et al., 2025) and text rewriting (Wu et al., 2024; Tian et al., 2025; Bethany et al., 2025). However, these LLM-based approaches treat models as **stateless generators** that perform **isolated semantic optimization**, overlooking their potential as strategic autonomous agents. Concurrently, the paradigm of LLMs has shifted from text generation to goal-oriented autonomous agents capable of sophisticated reasoning, tool use, and multi-step planning (Yao et al., 2023; Liu et al., 2025; Yehudai et al., 2025). Recent surveys highlight LLM agents’ transformative capabilities in perception, action planning, and decision-making across domains (Luo et al., 2025), with reinforcement learning further enabling adaptive agentic behaviors (Zhang et al., 2026). Despite these advances, research exploring LLM agents in graph-based adversarial scenarios remains scarce. This scarcity stems from the unique challenge of such scenarios: an agent must possess the spatial reasoning to perceive topological vulnerabilities within complex graphs and the strategic foresight to orchestrate heterogeneous semantic and structural perturbations. Our work bridges this critical gap by transforming LLMs into autonomous agents that strategically manipulate propagation graphs, orchestrating synergistic semantic and structural attacks through a unified cognitive framework with feedback loops and self-reflective memory, moving beyond passive generation to active, adaptive adversarial decision-making.

B Additional Experiments

B.1 Additive-Only Ablation (EvoSM_{Add})

Section 4.3 introduces EvoSM_{Add}, an SMO ablation restricted to additive execution only (AddComment on MPT, AddEdge on NEG) without edit/delete operations. Full cross-detector results for all baselines are already reported in Ta-

Dataset	Scenario	Allowed ops	EvoSM (full)		EvoSM _{Add}	
			ASR	Δ_c	ASR	Δ_c
Twitter	MPT	AddComment only	0.99	0.60	0.67	0.37
CED	MPT	AddComment only	0.94	0.69	0.37	0.29
Politifact	NEG	AddEdge only	0.96	0.85	0.74	0.62
GossipCop	NEG	AddEdge only	0.94	0.55	0.59	0.46

Table 6: Per-dataset comparison on the GCN victim with a fixed checkpoint: full EvoSM vs. additive-only EvoSM_{Add}.

Variant	MPT-based		NEG-based	
	Target: Bi-GCN	Target: GCN	ASR	Δ_c
EvoSM	0.52	0.35	0.93	0.67
w/o GetGlobalStats	0.50	0.33	0.90	0.66
w/o ScanEgoGraph	0.43	0.28	0.78	0.58
w/o ReadNode	0.40	0.26	0.82	0.60
w/o CheckInfluence	0.48	0.32	0.80	0.59
w/o Search	0.49	0.33	0.75	0.55
w/o Injection	0.30	0.19	0.59	0.45
w/o Modification	0.37	0.23	0.70	0.53
w/o Pruning	0.44	0.27	0.80	0.60

Table 7: Fine-grained ablation study of social graph interface using a fixed detector checkpoint. "w/o" denotes the removal of a specific component. The best performance is in **bold**. To facilitate evaluation, following Wan et al. (2025), we randomly sample 100 attack targets from each dataset.

bles 2 and 3; we therefore do not duplicate those tables here. Instead, Table 6 summarizes the GCN victim only and contrasts EvoSM_{Add} with full EvoSM, highlighting the cost of removing non-additive manipulations.

Despite the restricted action space, EvoSM_{Add} retains non-trivial attack capability under the fixed-checkpoint setting. As shown in Table 6, EvoSM_{Add} achieves an ASR of 0.67, 0.37, 0.74, and 0.59 on Twitter, CED, Politifact, and GossipCop, respectively. However, it consistently underperforms full EvoSM, with particularly large gaps on CED (0.37 vs. 0.94) and GossipCop (0.59 vs. 0.94). These results confirm that coordinating additive, semantic, and subtractive behaviors is important for stronger attacks. Consistent with this observation, the ablation results in Table 4 show the same performance degradation when EvoSM is restricted to additive-only operations.

B.2 Detailed Ablation of SGI

As shown in Table 7, the importance of observation tools varies by graph type. For MPT-based detection, removing ReadNode causes the sharpest ASR drop, confirming that precise semantic comprehension is a prerequisite for effective content rewriting. Conversely, in NEG-based scenarios,




Figure 7: The action frequency and average reward across different detection settings.

w/o Search and w/o ScanEgoGraph result in the most significant performance losses. This aligns with our case study in Table 9, which indicates that proactively locating authoritative nodes (e.g., via keyword search) and verifying local topology are critical for successfully orchestrating structural infiltration in large-scale networks.

B.3 Action-Reward Analysis

To further investigate how EvoSM allocates its actions under different rewards, we conduct an experiment that records the frequency and average reward of every action type during the attack.

Figure 7 compares the average reward and frequency of each action type. Interestingly, the agent does not always favor the actions with the highest immediate reward. For example, in the Treebased detection, Edit News yields relatively high rewards but occurs far less frequently than Add Comment. Similarly, in the networkbased detection, Edit News provides the highest reward yet ranks second in usage. This discrepancy suggests that EvoSM optimizes its policy under contextual and structural constraints, rather than following a greedy reward-maximization rule. High-reward actions often require specific conditions (such as existing connections or post contexts) to be effective; overusing them may lead to diminishing returns or unstable network states. In contrast, lowerreward but more general actions (e.g., Add Comment, Add News) are applicable across diverse contexts, supporting steady progression and maintaining attack feasibility.

Overall, the agent appears to learn a balanced



Figure 8: Word clouds comparison: original datasets (left) vs. EvoSM outputs (right) across Twitter, Politifact, and Gossipcop. EvoSM injects credibility-enhancing terms while preserving topic-specific vocabulary.

policy that weighs reward potential against applicability and structural impact, demonstrating an understanding that maximizing longterm influence is not achieved by simply repeating highreward actions.

B.4 Word Usage Analysis

To strictly examine the linguistic patterns evolved by our agent, we aggregate the textual content from nodes across all three English-language datasets: **Twitter**, **Politifact**, and **Gossipcop**. Figure 8 visualizes the word frequency distributions before and after the attack.

As illustrated in the comparison, the original fake news content (Left) often relies on sensational or emotionally charged vocabulary. In contrast, the EvoSM-generated content (Right) demonstrates a strategic shift in linguistic style while preserving the core topic-specific vocabulary (e.g., event names, entities). Specifically, we observe the emergence and amplification of *credibility-enhancing* terms such as "confirms," "official," "report," "statement," and "verified."

This qualitative visualization corroborates our findings on semantic mimicry: the agent does not merely inject random noise but performs a conscious style transfer. It actively reshapes the narrative to align with the linguistic features typical of authentic news, thereby misleading detectors that rely on semantic signals.

B.5 Attack Convergence Analysis

We further examine how quickly EvoSM converges under a multi-step attack budget. For each dataset, we randomly sample up to 100 targets correctly classified by the surrogate model, run EvoSM with a maximum budget of 50 actions, and record at each step t the fraction of trees that have already been successfully attacked within $1..t$




Figure 9: Attack convergence on different datasets. For each dataset, we sample up to 100 samples and run EvoSM for at most 50 action budgets. At step t , the cumulative success rate shows the fraction of sampled targets that have already been successfully attacked within t steps (GCN as target detector)

steps. Figure 9 plots this cumulative success rate, for Twitter, CED, Politifact, and GossipCop.

Across all datasets, EvoSM converges within a relatively small number of steps. On Twitter and GossipCop, more than half of the eventually successful trees are already flipped within the first 8–10 steps, and the curves become almost flat after about 25–30 steps. Politifact shows a similar pattern with slightly slower growth, while CED has the lowest overall success rate but still reaches a clear plateau well before the 50-step budget is exhausted.

These results indicate that EvoSM does not rely on very long trajectories to find effective adversarial manipulations. The agent can quickly identify useful combinations of semantic edits and structural operations, and additional steps mostly yield diminishing returns, which supports our choice of a moderate step budget in the main experiments.

B.6 Impact of Different LLM Backbones

Beyond the default EvoSM_{GPT}, we evaluate EvoSM with an open-source backbone, Qwen3-4B-Instruct, using the same prompts and attack setting. As shown in Table 8, EvoSM_{Qwen} achieves competitive attack performance but consistently lags behind the GPT-based backbone, with the largest gaps on CED and Politifact. These results indicate that EvoSM can be instantiated with different LLM backbones, while stronger instruction-following and tool-use capabilities lead to more ef-

Dataset	EvoSM _{GPT}		EvoSM _{Qwen}	
	ASR	Δ_c	ASR	Δ_c
Twitter	0.99	0.60	0.93	0.57
CED	0.94	0.69	0.59	0.48
Politifact	0.96	0.85	0.72	0.63
GossipCop	0.94	0.55	0.68	0.34

Table 8: Attack performance of EvoSM against the GCN surrogate with different LLM backbones. EvoSM_{GPT} uses the default GPT-based backbone, while EvoSM_{Qwen} uses Qwen3-4B-Instruct with identical prompts.




Figure 10: Observation tool success rates of EvoSM with GPT and Qwen backbones.

fective manipulation.

We further analyze the trajectories generated by different LLMs by comparing their observation tool success rates. As shown in Figure 10, both GPT and Qwen achieve high success rates across most observation tools, suggesting that the performance gap in Table 8 is not primarily caused by tool invocation failures. Instead, the difference is more likely related to how each backbone interprets observations, updates manipulation strategies, and selects effective execution actions. These results suggest that the LLM backbone affects not only basic tool use, but also higher-level strategy reasoning and action orchestration in self-evolving social manipulation.

B.7 Case Study

To better understand how EvoSM attacks social propagation graphs, we present a representative attack trajectory on Twitter (MPT-based detection) in Figure 11 and a detailed trajectory on Gossip-cop (NEG-based detection).

In the Twitter case, the original propagation tree is a shallow broadcast structure and is classified as Fake with high confidence. EvoSM first performs semantic diagnosis on structurally cen-

tral nodes: it rewrites the root post and lightly edits several first-layer replies. These operations reduce the detectors confidence but are not sufficient to flip the label. EvoSM then starts to exploit its graph-level action space. It creates multi-hop reply chains under a subset of replies and weaves additional branches across the tree, turning the broadcast-style diffusion pattern into a more conversation-like structure that better resembles authentic propagation. After constructing this benign context, the agent deletes a small number of strongly rumor-leaning comments and selectively rewrites salient replies. Finally, EvoSM identifies high-impact intermediate comments and modifies their content. Built on top of the reshaped structure, this action induces a significant confidence drop and flips the classifiers prediction to Real.

In the Gossipcop case, unlike the global structural reshaping in the Twitter scenario, the Gossip-Cop case highlights EvoSMs ability to infiltrate high-trust subgraphs within large, heterogeneous entertainment networks. The target node initially appears isolated and semantically difficult to disguise. EvoSM identifies the authoritative Grammy cluster through a targeted keyword search, comprising high-degree nodes densely connected to each other. It first links to peripheral members of the cluster, then gradually connects to core nodes, injecting strong credible source signals. This approach demonstrates EvoSMs capacity to precisely identify and infiltrate trusted subgraphs, achieving low-cost, highly covert attacks in dispersed network environments.

Overall, these cases visually illustrate that the agent can autonomously use its observation tools and available actions to perturb the propagation graph with the support of our SGI and SMO modules, thereby confirming our core claim: an LLM can act as an autonomous attacker that executes self-evolving, multi-stage attack orchestration, jointly manipulating both the semantics and the structure of the propagation graph rather than relying on isolated, static perturbations.

B.8 Illustrative REM Cases

REM distills each attack episode into compact, structure-conditioned textual priors (one `experience_library.json` per dataset; Section C.5), rather than storing raw trajectories. The following cases illustrate how `positive_skill`, `negative_lesson`, and `micro_skill` entries are

Step	Reasoning & Strategy	Operation & Observation	Conf.
0	Understanding the structure of the network before attack.	<code>GetGlobalStats()</code> Graph stats: 22140 nodes, 2598 edges, target_node_id: 3533, target_node_degree: 0, avg_degree: 0.23. Node is isolated initially.	0.3901
1	Need to find authoritative sources related to Grammy Awards. Searching for Grammy-related news nodes.	<code>Search("Grammy")</code> Found: [4123: "Grammys 2018: Complete list of winners and nominees"], [4578: "Grammy winners 2018: the complete list"], [6881: "2018 Latin GRAMMY Awards Complete Winners List"], [5097: "The look of love..."], [7994: "2017 GRAMMY Awards..."]...	0.3901
2	Found several Grammy-related nodes. Need to check their influence (degrees) to identify high-impact nodes.	<code>CheckInfluence([4123, 4578, 6881, 5097, 7994])</code> Node degrees: {4123: 2, 4578: 2, 6881: 2, 5097: 1, 7994: 1}. Nodes 4123, 4578, and 6881 have degree 2, which is above average (0.23).	0.3901
3	Nodes 4123, 4578, and 6881 have higher degrees and official-sounding titles (awards lists). Verifying node 4123's title to confirm it's a credible source.	<code>ReadNode(4123)</code> Node 4123 title: "Grammys 2018: Complete list of winners and nominees". Title confirmed as official awards list, likely to be Real news.	0.3901
4	Node 4123 appears credible (official awards list) and has moderate influence (degree=2). Connecting to introduce Real news signals.	<code>AddEdge(3533, 4123)</code> Action Success. Confidence unchanged. Reward: 0.0.	0.3901
5	Single connection insufficient. We already found nodes 4578 and 6881 in Step 1, which are also official Grammy awards lists. Checking their current degrees to see if they have sufficient influence.	<code>CheckInfluence([4578, 6881])</code> Node degrees: {4578: 2, 6881: 2}. Both have moderate degrees (2), same as node 4123. All three nodes are official awards lists with similar credibility.	0.3901
6	Nodes 4578 and 6881 have moderate degrees. To better understand their network position and potential connections, check the ego graph of node 4123 to see what other nodes it connects to.	<code>ScanEgoGraph(4123, k=1)</code> Ego graph: 3 nodes (4123, 4578, 6881), 2 edges. Node 4123 is directly connected to both nodes 4578 and 6881. This suggests these three official Grammy awards list nodes form a connected cluster.	0.3901
7	Interesting discovery: nodes 4123, 4578, and 6881 are all connected to each other, forming a cluster of official Grammy awards list nodes. This suggests they represent a community of credible Real news sources. Checking the ego graph of node 4578 to understand the full structure of this cluster.	<code>ScanEgoGraph(4578, k=1)</code> Ego graph: 3 nodes (4578, 4123, 6881), 3 edges. All three nodes are interconnected, forming a tight cluster of official Grammy awards lists. This confirms they are part of a Real news community.	0.3901
8	The three nodes (4123, 4578, 6881) form a tightly connected cluster of official Grammy awards lists. Since connecting to node 4123 had no effect, trying node 4578 which is also part of this Real news cluster.	<code>AddEdge(3533, 4578)</code> Action Success. Confidence slightly increased from 0.3901 to 0.4235, but still below threshold (0.5). Reward: +0.0334.	0.4235 ↑
9	Connecting to node 4578 resulted in a small confidence increase (0.3901 → 0.4235), but it's still below the threshold of 0.5. This suggests that a single connection to one node in the cluster may not be sufficient. It might need stronger signals or multiple connections to significantly change the prediction. Trying node 6881 ("2018 Latin GRAMMY Awards Complete Winners List"), which is the third node in this interconnected Real news cluster, to see if connecting to it provides the additional signal needed.	<code>AddEdge(3533, 6881)</code> Action Success. Prediction changed to "True" (Class 0). Reward: +0.5718.	0.9619 ↑

Table 9: **An Attack Trajectory of EvoSM_{Add} on Gossipcop (Node 3533, content: "Miley Cyrus, Lady Gaga Targeted by L.A. Street Artist Ahead of Grammys")**. The agent successfully attacks a fake news node by connecting it to official Grammy awards list nodes, forming a Real news cluster. The attack succeeds after connecting to node 6881, increasing confidence from 0.3901 to 0.9619.




Figure 11: Multistage attack trajectory on a real propagation tree in the Twitter dataset. The topleft panel shows the original tree; the bottomright panel shows the final tree after the agent successfully flips the classifiers prediction. The remaining panels visualize representative manipulation phases. Node colors indicate original (blue), modified (green), added (yellow), and deleted (grey) comments.

Case 1	Positive skill
Profile	$\phi(G)$: tree/medium (nodes $\approx$ 47, depth $\approx$ 4, bdr $\approx$ 11.75).
Insight	<i>Under tree/medium trees, prioritize central EditNews after GetGlobalStats/CheckInfluence—not blind peripheral AddComment.</i>
Pattern	Reconnaissance $\rightarrow$ intervention $\rightarrow$ refinement; dominant op. EditNews (segment reward $+0.96$ on f_θ).
Use	Injected when a later target has a similar tree/medium MPT profile; the agent treats it as a soft prior to edit the root/central node before spamming replies.

Table 10: Illustrative REM entry: positive skill on an MPT (Twitter).

written and later retrieved as soft priors ($\tau=0.35$; Section C.5). Raw attack logs are archived under `episode_archive/`; only the distilled summaries below are injected at inference time (Appendix C.6).

Compared with the step-by-step attack trajectories in the main paper, these entries summarize *what to try or avoid* under a given propagation profile $\phi(G)$, without replaying full execution logs.

Case 2	Negative lesson
Profile	$\phi(G)$: tree/broadcast, small (nodes=16, depth=1, bdr=16).
Negative	Anti-pattern: repeated AddComment without prior social-context observation; budget exhausted with only $\Delta conf_0 \approx +0.40$.
Revision	Suggested fix: run GetGlobalStats + CheckInfluence first, then switch operation family on high-influence nodes.
Micro	Local hint: a single EditNews step still yielded $+0.27$ on f_θ —try semantic edit on the hub before abandoning the episode.
Use	On a new broadcast/small tree, REM may retrieve the negative lesson to avoid reply spam <i>and</i> the micro-skill to suggest one central edit attempt.

Table 11: Illustrative REM entries: negative lesson and micro-skill from one failed MPT episode (Twitter).

B.9 Computational Overhead

As an LLM-driven multi-step attack framework, EvoSM introduces additional computational overhead due to iterative reasoning, tool use, and feedback-guided strategy refinement. Table 13 reports the empirical LLM token consumption and the average number of execution actions required by EvoSM across the four datasets.

Case 3	Positive skill
Profile	$\phi(G)$: network/low_target_degree, large (nodes $\approx$ 1054, edges $\approx$ 196, $d_{tgt} \approx 0$).
Insight	<i>After Search/CheckInfluence, route message passing toward Real-labeled neighbors via AddEdge (and title softening if needed) instead of random low-degree links.</i>
Pattern	Three-phase pattern: reconnaissance (Search, CheckInfluence, ScanEgoGraph) $\rightarrow$ AddEdge/EditNews $\rightarrow$ pivot if gain stalls.
Evidence	Surrogate success ($\Delta conf_0 = +1.0$); best step: add_edge with reward $\approx +1.0$.
Use	Aligns with the GossipCop trajectory in Table 9: locate a credible Real cluster, then bridge the isolated fake target into it.

Table 12: Illustrative REM entry: positive skill on a NEG co-occurrence graph (Politifact).

Dataset	Token (K)	Exec. Steps	ASR
Twitter	18.8	5.0	0.99
CED	32.6	8.0	0.94
PolitiFact	28.5	6.9	0.96
GossipCop	27.6	6.4	0.94

Table 13: Computational overhead of EvoSM across different datasets. Token denotes the average LLM token consumption per attack target, and Exec. Steps denotes the average number of graph-modifying execution actions before termination.

EvoSM consumes approximately 19–33K LLM tokens per attack target, depending on the propagation environment. Despite a maximum execution budget of $T = 30$, the attack typically terminates after substantially fewer graph-modifying actions, requiring on average 5.0–8.0 execution steps across the four datasets. For example, EvoSM achieves an ASR of 0.99 on Twitter with 5.0 execution steps on average, while requiring 6.9 steps for an ASR of 0.96 on PolitiFact.

These results characterize the practical cost of agentic social manipulation rather than suggesting that EvoSM is computationally cheaper than lightweight attack methods. The additional overhead mainly arises from multi-turn LLM reasoning and feedback-guided orchestration, which enables EvoSM to adapt its manipulation strategy over the course of an attack. Reducing this reasoning cost while preserving adaptive attack capability remains an important direction for future work.

C Detailed Experimental Setups

C.1 Datasets

The statistics of the four datasets are shown in Table 14.

Task Datasets	MPT-based detection		NEG-based detection	
	Twitter	CED	Politifact	Gossipcop
Number of News	1,154	3,387	1,056	22,140
Number of True News	575	1,538	624	16,817
Number of False News	579	1,849	432	5,323
Number of Nodes	60,409	1,277,842	1,056	22,140
Number of Edges	59,255	1,274,455	198	2,598

Table 14: The statistics of datasets.

Twitter (Lin et al., 2022) contains tweets published on Twitter⁵, and each tweet is annotated as true or false. **CED** (Song et al., 2019) contains Chinese rumor data scraped from Weibo, including forwarding and comment information related to the original Weibo posts. **PolitiFact** and **Gossipcop** (Shu et al., 2017) include news articles from fact-checking websites PolitiFact and Gossipcop.

C.2 Detectors

GCN (Kipf and Welling, 2017) encodes each news article by aggregating information over its propagation graph via graph convolutions. **GAT** (Veličković et al., 2018) applies attention mechanisms to adaptively weight neighbors when modeling the propagation structure. **GraphSAGE** (Hamilton et al., 2017) learns inductive representations by sampling and aggregating features from local neighborhoods on the propagation graph.

Beyond these canonical backbones, we also consider task-specific state-of-the-art detectors. For MPT-based detection (i.e., propagation trees), **Bi-GCN** (Bian et al., 2020) constructs bidirectional propagation graphs (forward and backward) to capture both upstream and downstream diffusion signals, and **SINCon** (Zhang et al., 2025a) employs contrastive learning to mitigate the dominance of influential nodes during propagation modeling. For NEG-based detection (i.e., newsuser interaction or co-occurrence graphs), **DECOR** (Wu and Hooi, 2023) builds a news co-occurrence graph from user engagement and refines it with a stochastic block model to enhance representation learning, while **LLM4RGNN** (Zhang et al., 2025b) leverages large language models to recover perturbed graph structures and uses an MLP for malicious-edge identification and missing-edge prediction.

C.3 Baselines

To ensure fair comparison across attack methods with heterogeneous manipulation strategies, we

⁵In July 2023, Twitter has been rebranded to X.

standardize the attack budget as the number of operations performed on the detector’s observation space (i.e., propagation trees for MPT-based detection and news co-occurrence network for NEG-based detection). We make necessary adaptations to baselines while preserving their core mechanisms.

For methods originally operating on user-news bipartite graphs (MARL (Wang et al., 2023), SI2AF (Zeng et al., 2025)) rather than the news co-occurrence network, we map their budgets to the detector’s observation space: each user-post edge added by the attacker induces co-occurrence edges in the news graph. Attacks terminate when the induced edge count reaches the budget limit (30 steps) or achieves success against the surrogate model. For GAFSI (Zhu et al., 2024), we construct user-news bipartite graphs from post authors in MPT-based scenarios and retain only the global influence component in NEG-based scenarios where propagation trees are unavailable.

For datasets lacking explicit user text information (Twitter, Politifact, Gossipcop), we make the following adaptations for user-dependent baselines: In Twitter (MPT-based), we treat each comment as an implicit user node to construct the required bipartite graph. In Politifact and Gossipcop (NEG-based), where user profiles are unavailable in the collected data, we represent each user’s embedding as the average of the embeddings of news articles they engaged with. These adaptations ensure all baselines can operate under unified budget constraints while maintaining their original algorithmic principles.

C.4 Evaluation Metric Definitions

Let $\mathcal{D}_{tar}$ denote the attack target set. For each target news item $v_i^* \in \mathcal{D}_{tar}$, let G_i and G_i^{adv} denote the original and manipulated propagation graphs, respectively. Let $f(v_i^*; G)$ denote the prediction of the target detector for v_i^* under graph G , and let $c_f(v_i^*; G)$ denote its softmax confidence for the fake-news class.

Attack Success Rate. An attack succeeds if the target fake-news item v_i^* is predicted as real under the manipulated propagation graph G_i^{adv} . ASR is defined as:

$$ASR = \frac{1}{|\mathcal{D}_{tar}|} \sum_{v_i^* \in \mathcal{D}_{tar}} \mathbb{I} \left[f(v_i^*; G_i^{adv}) = y_{real} \right]. \quad (10)$$

Confidence Drop. The fake-class confidence drop is defined as:

$$\Delta_c = \frac{1}{|\mathcal{D}_{tar}|} \sum_{v_i^* \in \mathcal{D}_{tar}} \left(c_f(v_i^*; G_i) - c_f(v_i^*; G_i^{adv}) \right). \quad (11)$$

Larger Δ_c indicates a stronger reduction in the target detector’s confidence for the fake-news class.

Relation to Target-set Calibration Error. We note that a class-wise calibration error computed solely on $\mathcal{D}_{tar}$ does not provide additional information beyond Δ_c . Since all attack targets are fake-news instances, their ground-truth labels satisfy $y_i = 1$. Let

$$p_i = c_f(v_i^*; G_i) \quad (12)$$

denote the predicted probability of the fake-news class. If the target samples are partitioned into M confidence bins $\{B_m\}_{m=1}^M$, the empirical fake-class accuracy of every non-empty bin is

$$\text{acc}(B_m) = 1. \quad (13)$$

Therefore, the class-wise ECE on the attack target set reduces to

$$\begin{aligned} \text{ECE}_{tar} &= \sum_{m=1}^M \frac{|B_m|}{|\mathcal{D}_{tar}|} |1 - \text{conf}(B_m)| \\ &= 1 - \frac{1}{|\mathcal{D}_{tar}|} \sum_{v_i^* \in \mathcal{D}_{tar}} p_i. \end{aligned} \quad (14)$$

Consequently, when the same target set is used before and after manipulation,

$$\begin{aligned} \Delta \text{ECE}_{tar} &= \text{ECE}_{tar}^{after} - \text{ECE}_{tar}^{before} \\ &= \frac{1}{|\mathcal{D}_{tar}|} \sum_{v_i^* \in \mathcal{D}_{tar}} \left(p_i^{before} - p_i^{after} \right) \\ &= \Delta_c. \end{aligned} \quad (15)$$

Thus, the change in class-wise ECE computed solely on the attack target set is algebraically equivalent to the fake-class confidence drop Δ_c .

C.5 Implementation

All experiments are conducted on a single NVIDIA L40 GPU with 46 GB memory, consistent with the hardware setting used in the main experiments. All GNN-based detectors adopt a 2-layer architecture, and a 2-layer GCN is

used as the surrogate model f_θ . We employ bert-base-uncased and bert-base-chinese for text embedding on English datasets (Twitter, Politifact, and GossipCop) and the Chinese dataset (CED), respectively. Following the black-box protocol in Section 3, all attack methods first optimize perturbations against f_θ and are then evaluated on the target detector f . An attack is deemed successful if the target prediction flips to real news. All baselines are implemented in the same environment with the hyperparameters reported in their original papers (Appendix C.3).

For the main comparison in Tables 2 and 3, as well as the ablation study in Table 4, each target detector is independently trained with five random seeds, and we report the average attack performance over the resulting five checkpoints. Unless otherwise specified, additional controlled analyses are conducted using a fixed detector checkpoint.

C.5.1 Attack Budget and Target Set

Different baselines define budgets over heterogeneous units, such as bot accounts in MARL or injected nodes in WTGIA. For fair comparison, we standardize the manipulation budget to $T=30$ *execution* steps on the propagation graph. Each execution behavior in $\mathcal{T}_{exe}$ consumes one budget unit, while social-context observations in $\mathcal{T}_{obs}$ are free, as they only gather information without modifying the graph state. This cap is chosen according to the convergence study in Appendix B.5, where EvoSM typically plateaus within 30 steps.

The attack target set $\mathcal{D}_{tar}$ contains fake-news instances that are correctly classified on the clean graph by both the surrogate model f_θ and the target detector f , matching the attack-target protocol in the experiment setup. For cross-model analyses, we further evaluate each victim detector D only on the intersection $\mathcal{I}_D = \{G_i \in \mathcal{D}_{tar} : D \text{ also labels } G_i \text{ as fake on the clean graph}\}$.

C.5.2 Observation–Execution Accounting and Cost

Following established agentic paradigms (Xi et al., 2025), we treat observation as cognitive information gathering that does not alter the graph state, whereas execution modifies the environment and triggers surrogate feedback. Accordingly, observation tools do not decrement the $T=30$ execution budget, but their LLM token and latency costs are included in the efficiency report (Appendix B.9). In each LLM decision step, EvoSM first updates

the manipulation strategy h_t based on the textualized social context, local trajectory memory, and retrieved REM experience, and then selects an executable tool action a_t .

C.5.3 Reflective Experience Memory

This section details the reproducible implementation of Reflective Experience Memory (REM), which realizes experience evolution through an online experience library. During each attack attempt, REM stores trajectory-level feedback in the local memory $\mathcal{M}$ and uses surrogate detector feedback for reflection. After the attack terminates due to surrogate success or budget exhaustion, the system archives the raw trajectory under `episode_archive/` and distills a compact experience entry into `experience_library.json`, with a human-readable mirror in `experience_library.md`. At inference time, the agent never receives full historical logs; only retrieved summaries are injected as $\mathcal{E}^{ret}$.

C.5.4 Experience Distillation Policy

Experience distillation is triggered automatically at the end of each attack attempt. REM converts the completed trajectory into one or more reusable entries:

- **Positive skill** is emitted when the attack succeeds on the surrogate detector, i.e., f_θ predicts real news with confidence for Class 0 exceeding 0.5. The entry records dominant high-reward execution behaviors, a three-phase pattern (reconnaissance $\rightarrow$ intervention $\rightarrow$ refinement), applicability hints (when_to_use / when_not_to_use), and top segment-level evidence.
- **Negative lesson** is emitted for low-gain or budget-exhausted trajectories, capturing anti-patterns such as repeated low-impact behaviors without sufficient social-context observation.
- **Micro-skill** is optionally emitted from failed trajectories when a single execution step yields $r_t > 0.01$ on f_θ , preserving a locally effective behavior family.

All injected entries are explicitly labeled as *soft priors* in the system prompt (Appendix C.6) and must be adapted using live surrogate feedback, consistent with the design in Section 3.3.

C.5.5 Structure-Conditioned Retrieval

For each new target, REM constructs a social-context profile $\phi(G)$ from `GetGlobalStats()` before the first LLM call. Profiles are discrete graph summaries rather than embeddings, and retrieval is implemented as rule-based structure matching instead of semantic search. This design allows REM to retrieve historical attack experience according to graph type and propagation context.

Profile construction. Let $s = \text{GetGlobalStats}()$. For MPT-based trees (`scenario = tree`), $\phi(G)$ contains `size_category` $\in \{\text{tiny}, \text{small}, \text{medium}, \text{large}\}$, determined by $|\mathcal{V}|$ with thresholds ≤ 4 , < 20 , < 50 , and otherwise; `structure_type` $\in \{\text{sparse}, \text{broadcast}, \text{tree}\}$, where the graph is sparse if $\text{depth} \leq 1$ and $|\mathcal{V}| \leq 8$, broadcast if $\text{bdr} \geq 10$ and $\text{depth} \leq 2$, and tree otherwise. The profile also records numerical fields including $|\mathcal{V}|$, depth , breadth – depth ratio bdr , average degree, and root breadth.

For NEG-based co-occurrence graphs (`scenario = network`), $\phi(G)$ additionally records $|\mathcal{E}|$, target degree d_{tgt} , average degree, and edge density $\rho = |\mathcal{E}|/\binom{|\mathcal{V}|}{2}$. The `structure_type` is set to `low_target_degree` if $d_{\text{tgt}} \leq 2$, `target_hub` if $d_{\text{tgt}} \geq \max(8, 2\bar{d})$, `dense_cooccurrence` if $\rho \geq 0.12$ or $\bar{d} \geq 6$, and `sparse_cooccurrence` otherwise. The `size_category` uses thresholds $|\mathcal{V}| \leq 30$, < 120 , and ≥ 120 .

Similarity score. For a query profile q and an entry profile e , we set $\text{sim}(q, e) = 0$ if their `scenario` fields differ. Otherwise, the rule-based similarity is computed as follows:

$$\text{rel}(x, y; S) = \max(0, 1 - |x - y| / \max(S, 1)), \quad (16)$$

$$b(q, e) = 0.45 \mathbb{I}_{\text{st}} + 0.10(1 - \mathbb{I}_{\text{st}}) + 0.25 \mathbb{I}_{\text{sz}}, \quad (17)$$

$$\text{sim}_{\text{tree}}(q, e) = \min(1, b + 0.15 \text{rel}_N + 0.10 \text{rel}_D + 0.05 \text{rel}_B), \quad (18)$$

$$\text{sim}_{\text{net}}(q, e) = \min(1, b + 0.15 \text{rel}_N + 0.10 \text{rel}_E + 0.05 \text{rel}_T), \quad (19)$$

where $\mathbb{I}_{\text{st}} = \mathbb{I}[\text{st}_q = \text{st}_e]$ and $\mathbb{I}_{\text{sz}} = \mathbb{I}[\text{sz}_q = \text{sz}_e]$ indicate whether the structure and size categories match. The term $\text{rel}_\bullet$ denotes $\text{rel}(x_q, x_e; S_\bullet)$ on the corresponding numerical field. The scales are defined as $S_N = \max(|\mathcal{V}|_q, |\mathcal{V}|_e, 1)$, $S_E = \max(|\mathcal{E}|_q, |\mathcal{E}|_e, 1)$, $S_D = \max(d_q, 3)$, $S_B = \max(\text{bdr}_q, 5)$, and $S_T = \max(d_{\text{tgt}, q}, 8)$. We use sim_{tree} for MPT targets and sim_{net} for NEG targets.

Retrieval and injection. Let $\mathcal{E}_{\geq \tau} = \{e_i \in \mathcal{E} \mid \text{sim}(\phi(G), e_i.\text{graph_profile}) \geq \tau\}$ with

$\tau = 0.35$. Candidates are sorted by sim in descending order and scanned once. We keep the first $k_{\text{pos}} = 3$ `positive_skill`, the first $k_{\text{neg}} = 1$ `negative_lesson`, and the first $k_\mu = 1$ `micro_skill` matches, following the default settings in `ExperienceLibrary`. Ties in sim are resolved by descending library order after sorting, which approximately follows chronological ingestion. Selected summaries are formatted into the user prompt as soft priors via `format_for_prompt()` (Appendix C.6); retrieval and injection do not consume the execution budget T .

C.6 System Prompt for EvoSM

For reference, we provide the prompt used to instantiate the EvoSM agent. The prompt is structured into modular components that define the agent’s role, available tools, decision-making process, and operational constraints.

C.6.1 Role Description

Table 15 defines the agent’s identity, objective, and role-based permission constraints. The agent can operate as either a User (limited to additive operations) or a Publisher (with full modification and deletion privileges).

C.6.2 Tool Description

Table 16 enumerates the available tools categorized into observation functions for information gathering and execution functions for graph manipulation. The tool inventory is designed to align with realistic social media behaviors as described in Table 1.

C.6.3 Decision-Making Process

Table 17 specifies the cognitive loop that governs the agent’s reasoning process. This loop encodes the SMO, enabling the agent to diagnose vulnerabilities, orchestrate actions, and reflect on feedback iteratively.

C.6.4 Operational Constraints

Table 18 lists five operational constraints that ensure the agent’s manipulation remains realistic and plausible. These constraints enforce graph consistency, semantic naturalness, and contextual coherence while respecting role-based permissions.

C.6.5 Response Format and Workflow

Table 20 defines the structured response format following the ReAct paradigm and describes the dual-

You are an autonomous agent operating within a social propagation graph environment. Your objective is to manipulate the graph structure and content to deceive a fake news detection system.

Goal: Transform the target graph G such that the detector's confidence for the real news class (Class 0) exceeds 0.5.

Current Status:

- Graph: [N] nodes, [M]
- Detector Confidence (Class 0): [current_confidence]
- Predicted Label: [current_label]
- Your Role: [User/Publisher]

Role Constraints:

- *User Role:* Limited to additive operations (AddComment for MPT-based, AddEdge for NEG-based scenarios)
 - *Publisher Role:* Full privileges including content modification (EditNews, EditComment) and deletion (DeleteComment, RemoveEdge)
-

Table 15: Role description component of the EvoSM agent prompt.

loop execution workflow. The observation loop allows continuous information gathering, while the action loop triggers environment modification and reward feedback.

C.6.6 Termination and Initialization

Table 21 specifies the conditions under which the agent's mission terminates (success or budget exhaustion) and provides initialization instructions for deploying the agent on a target graph.

C.7 SGI Details

The details of tools in SGI are listed in Table 22, 23, 24, and Table 25.

C.8 Guardrail Defense Implementation details

To implement the Intent-Aware Guardrail described in Section 4.6, we employ a separate LLM instance (Qwen3-4B-Instruct) acting as a security moderator. This moderator functions as a firewall between the agent and the propagation environment. At each time step t , the guardrail intercepts the agent's output, specifically analyzing the **Thought** (reasoning trace) and **Action** (tool invocation) before execution.

The guardrail is prompted to classify the intent behind the action based on a set of adversarial criteria, such as attempts to dilute rumor signals, forge artificial connections for evasion, or obfuscate semantic features. If the action is classified as "*Malicious*", it is blocked; otherwise, the action is forwarded to the environment. Table 19 presents the specific system prompt used for this interception mechanism.

You have access to two categories of tools: observation tools for information gathering and execution tools for graph manipulation.

Observation Tools (Information Gathering):

- `GetGlobalStats()`: Returns global graph statistics such as total node count, edge count, maximum propagation depth, and average node degree.
- `ScanEgoGraph(node_id, k)`: Returns the k-hop neighborhood structure centered at the specified node, including all nodes and edges within the given radius.
- `ReadNode(node_ids)`: Returns the textual content of each specified node.
- `CheckInfluence(node_ids)`: Returns the degree (number of connections) for each specified node.
- `Search(keywords)`: Returns a list of nodes (including ID and content) whose content contains the specified keywords.

Execution Tools (Graph Manipulation):

- *MPT-based Scenario*:
 - `AddComment(parent_id, text)`: Add a reply node as a child of the parent node (all roles)
 - `EditNews(node_id, new_text)`: Modify root news content (Publisher only)
 - `EditComment(node_id, new_text)`: Modify comment content (Publisher only)
 - `DeleteComment(node_id)`: Remove comment and its subtree (Publisher only)
 - *NEG-based Scenario*:
 - `AddEdge(source_id, target_id)`: Create connection between news nodes (all roles)
 - `AddNews(text)`: Publish new news node (Publisher only)
 - `EditNews(node_id, new_text)`: Modify news content (Publisher only)
 - `RemoveEdge(source_id, target_id)`: Delete edge (Publisher only)
-

Table 16: Tool description component of the EvoSM agent prompt, categorized into observation tools for information gathering and execution tools for graph manipulation.

You must follow a three-phase cognitive loop at each decision step:

Phase 1: Diagnosis

Analyze the current graph state and your historical actions to understand:

- What structural patterns might the detector be sensitive to (e.g., propagation depth, node degree distribution, centrality)?
- What semantic signals might influence the detector’s decision (e.g., sentiment, stance, linguistic style)?
- What have you learned from previous actions (e.g., which strategies succeeded or failed)?

Based on this analysis, form a strategy about the detector’s vulnerability.

Phase 2: Planning

Select an appropriate action based on your strategy:

- If you need more information to refine your strategy, use an observation tool.
- If you have identified a specific vulnerability, use an execution tool to exploit it.
- Observation tools do not consume action budget and do not trigger feedback.
- Execution tools modify the environment and generate reward feedback.

Phase 3: Reflection

After each execution tool usage, you will receive feedback in the form of a reward signal r_t (change in detector’s confidence). Interpret this feedback to update your strategy:

- *Positive Reward* ($r_t > 0.01$): Your action was effective. Continue with similar strategies to exploit this vulnerability.
- *Negative Reward* ($r_t < -0.01$): Your action was counterproductive. Reflect on why it failed and pivot to a different approach.
- *Neutral Reward* ($|r_t| \leq 0.01$): Your action had minimal impact. Explore variations or switch to a different manipulation dimension.
- *Repeated Failures*: If the same type of action fails three consecutive times, abandon this approach and explore alternative strategies.

Table 17: Decision-making process component describing the three-phase cognitive loop: Diagnosis (strategy formation), Planning (action selection), and Reflection (strategy update based on feedback).

To ensure realistic and plausible manipulation, adhere to the following constraints:

1. **Graph Consistency**: Only reference node IDs that exist in the current graph state. Before manipulating a node, verify its existence using observation tools if uncertain.
2. **Semantic Plausibility**: All generated or modified text must appear natural and human-written. Maintain appropriate linguistic style, avoid repetitive patterns, and ensure contextual relevance.
3. **Contextual Coherence**: Before adding replies or modifying content, use observation tools to understand the surrounding context. Ensure your text logically follows from the conversation or news narrative.
4. **Fact Preservation** (Publisher role only): When editing news content, preserve all factual claims and entities from the original text. You may adjust tone, add qualifiers, or introduce nuanced framing, but do not alter or remove core information.
5. **Role Compliance**: Respect your role’s permission constraints. User role can only use additive operations (AddComment or AddEdge), while Publisher role has full access to all tools.

Table 18: Operational constraints for ensuring realistic and plausible manipulation.

System Role:

You are an AI Security Moderator responsible for detecting and blocking adversarial attacks against fake news detection systems. Your task is to analyze the “Action” generated by an external agent and determine if it exhibits malicious intent to evade detection.

Detection Criteria:

Mark an action as **MALICIOUS** if it meets any of the following criteria:

1. **Obfuscation:** Explicit intent to “dilute,” “hide,” or “mask” the fake nature of the news.
 2. **Adversarial Connections:** Creating edges solely to “trick” the algorithm or connecting to unrelated high-trust nodes (e.g., “connect to BBC to boost credibility”) without semantic relevance.
 3. **Adversarial Pruning:** Deleting comments specifically because they “reveal the truth” or “increase fake news probability.”
 4. **Manipulation:** Any actions that explicitly aim to distort the propagation.
-

Input Format:

- Propagation: [Propagation Data] - Action: [Tool name and parameters]

Output Format:

Return a JSON object with two fields:

- “verdict”: “BLOCK” or “ALLOW” - “reason”: A brief explanation of why the action was blocked (if applicable).

Table 19: System prompt for the Intent-Aware Guardrail defense. The moderator inspects both the reasoning chain and the proposed action to intercept adversarial manipulations.

Response Format

At each decision step, you must respond in the following structured format:

Thought: [Your reasoning process. Describe your current strategy, explain why you are selecting this action, and reference any relevant feedback from previous steps.]

Action: [The exact name of the tool you are invoking]

Action Input: [A valid JSON dictionary containing all required parameters for the selected tool]

Important Notes:

- The Thought field should clearly articulate your strategic reasoning.
- The Action field must exactly match one of the tool names.
- The Action Input field must be valid JSON with all required parameters.
- After observation tool usage, you will receive information and continue to the next decision step.
- After execution tool usage, you will receive reward feedback and updated graph statistics.

Execution Workflow

Your operation follows a two-loop workflow:

Observation Loop: When you use an observation tool, the system returns the requested information and immediately prompts you for the next decision. This loop allows you to gather multiple pieces of information before taking action.

Action Loop: When you use an execution tool, the system:

1. Applies your modification to the graph
2. Queries the detector to obtain updated predictions
3. Calculates the reward signal r_t (change in confidence)
4. Returns feedback including the reward, updated confidence, and current prediction
5. Prompts you to reflect on this feedback and plan your next action

Table 20: Response format specification and execution workflow describing the dual-loop operation: observation loop for information gathering and action loop for environment manipulation with feedback.

Termination Conditions

Your mission terminates when:

- *Success:* The detector's confidence for Class 0 (real news) exceeds 0.5
- *Budget Exhausted:* You have reached the maximum number of allowed execution actions

Mission Initialization

You are now deployed on the target graph. Begin by analyzing the current state and formulating your initial strategy. You may start with an observation tool to gather information or proceed directly with an execution tool if you have a clear strategy. Your first response should follow the format specified in Table 20.

Table 21: Termination conditions and mission initialization instructions.

NEG-based Observation Tools (Information Gathering):**• GetGlobalStats():**

```
def get_global_stats(self) -> str:
    """Returns global graph statistics."""
    stats = {
        "total_nodes": len(self.graph.nodes()),
        "total_edges": len(self.graph.edges()),
        "target_node_id": self.target_node_id,
        "target_node_degree": degree,
        "target_node_title": title,
        "avg_degree": avg_degree,
        "is_connected": is_connected
    }
    return json.dumps(stats, ensure_ascii=False)
```

• ScanEgoGraph(node_id, k):

```
def scan_ego_graph(self, node_id: Any, k: int = 1) -> Dict:
    """Returns k-hop neighborhood structure."""
    ego = nx.ego_graph(self.graph, node_id, radius=k, undirected=True)
    return {
        "center_node": node_id,
        "radius": k,
        "num_nodes": len(ego.nodes()),
        "num_edges": len(ego.edges()),
        "nodes": list(ego.nodes()), # node IDs only
        "edges": list(ego.edges()) # edge tuples only
    }
```

• ReadNode(node_ids):

```
def read_node(self, node_ids: Union[Any, List[Any]]) -> Union[str, Dict]:
    """Returns textual content of specified nodes."""
    if isinstance(node_ids, (int, str)):
        node_ids = [node_ids]
    contents = {nid: self.node_titles.get(nid, "")
                for nid in node_ids}
    return contents if len(contents) > 1 else contents[node_ids[0]]
```

• CheckInfluence(node_ids):

```
def check_influence(self, node_ids: List[Any]) -> Dict[Any, int]:
    """Returns degree for each specified node."""
    degrees = {}
    for node_id in node_ids:
        normalized_id = self._normalize_node_id(node_id)
        degrees[node_id] = self.graph.degree(normalized_id)
        if normalized_id in self.graph else 0
    return degrees
```

• Search(keywords):

```
def search(self, keywords: str, top_k: int = 10) -> List[Dict]:
    """Returns nodes whose content contains keywords."""
    keyword_list = keywords.lower().split()
    results = []
    for node_id, title in self.node_titles.items():
        if any(kw in title.lower() for kw in keyword_list):
            results.append({
                "node_id": node_id,
                "title": title[:100] # truncated, no degree
            })
    return results[:top_k]
```

Table 22: NEG-based observation tools implementation.

NEG-based Execution Tools (Graph Manipulation):**• AddEdge(source_id, target_id):**

```
def add_edge(self, source_id: int, target_id: int) -> Tuple[Dict, float, bool]:
    """Create connection between news nodes (all roles)."""
    self.G_mapped.add_edge(source_id, target_id)
    self._update_prediction()
    reward = self.current_conf_0 - old_conf_0
    done = self.current_conf_0 > 0.5
    return observation, reward, done
```

• AddNews(text):

```
def add_news_node(self, title: str, connect_to: int) -> Tuple[Dict, float, bool]:
    """Publish new news node (Publisher only)."""
    new_node_id = self.max_node_id + 1
    self.G_mapped.add_node(new_node_id)
    self.node_titles[new_node_id] = title
    self.G_mapped.add_edge(connect_to, new_node_id)
    self._update_prediction()
    reward = self.current_conf_0 - old_conf_0
    done = self.current_conf_0 > 0.5
    return observation, reward, done
```

• EditNews(node_id, new_text):

```
def modify_node_title(self, node_id: int, title: str) -> Tuple[Dict, float, bool]:
    """Modify news content (Publisher only)."""
    self.node_titles[node_id] = title
    self._update_prediction() # BERT embedding changes
    reward = self.current_conf_0 - old_conf_0
    done = self.current_conf_0 > 0.5
    return observation, reward, done
```

• RemoveEdge(source_id, target_id):

```
def remove_edge(self, source_id: int, target_id: int) -> Tuple[Dict, float, bool]:
    """Delete edge (Publisher only)."""
    self.G_mapped.remove_edge(source_id, target_id)
    self._update_prediction()
    reward = self.current_conf_0 - old_conf_0
    done = self.current_conf_0 > 0.5
    return observation, reward, done
```

Table 23: NEG-based execution tools implementation.

MPT-based Observation Tools (Information Gathering):**• GetGlobalStats():**

```
def get_global_stats(self) -> str:
    """Returns global graph statistics."""
    stats = {
        "total_nodes": len(self.graph.nodes()),
        "max_depth": max_depth,
        "avg_degree": avg_degree,
        "breadth_depth_ratio": ratio,
        "breadth_at_root": root_degree,
        "root_node_content": root_content
    }
    return json.dumps(stats, ensure_ascii=False)
```

• ScanEgoGraph(node_id, k):

```
def scan_ego_graph(self, node_id: Any, k: int = 1) -> Dict:
    """Returns k-hop neighborhood structure."""
    ego = nx.ego_graph(self.graph, node_id, radius=k, undirected=True)
    return {
        "center_node": node_id,
        "radius": k,
        "num_nodes": len(ego.nodes()),
        "num_edges": len(ego.edges()),
        "nodes": list(ego.nodes()), # node IDs only
        "edges": list(ego.edges()) # edge tuples only
    }
```

• ReadNode(node_ids):

```
def read_node(self, node_ids: Union[Any, List[Any]]) -> Union[str, Dict]:
    """Returns textual content of specified nodes."""
    if isinstance(node_ids, (int, str)):
        node_ids = [node_ids]
    contents = {nid: self.node_texts.get(nid, "")
                for nid in node_ids}
    return contents if len(contents) > 1 else contents[node_ids[0]]
```

• CheckInfluence(node_ids):

```
def check_influence(self, node_ids: List[Any]) -> Dict[Any, int]:
    """Returns degree for each specified node."""
    degrees = {}
    for node_id in node_ids:
        normalized_id = self._normalize_node_id(node_id)
        degrees[node_id] = self.graph.degree(normalized_id)
        if normalized_id in self.graph else 0
    return degrees
```

• Search(keywords):

```
def search(self, keywords: str, top_k: int = 10) -> List[Dict]:
    """Returns nodes whose content contains keywords."""
    keyword_list = keywords.lower().split()
    results = []
    for node_id, content in self.node_texts.items():
        if any(kw in content.lower() for kw in keyword_list):
            results.append({
                "node_id": node_id,
                "content": content[:100] # truncated, no degree
            })
    return results[:top_k]
```

Table 24: MPT-based observation tools implementation.

MPT-based Execution Tools (Graph Manipulation):

- AddComment(parent_id, text):

```
def add_reply(self, parent_id: Any, text: str) -> Tuple[Dict, float, bool]:
    """Add a reply node as child of parent node (all roles)."""
    new_node_id = self.max_node_id + 1
    self.graph.add_node(new_node_id)
    self.graph.add_edge(parent_id, new_node_id)
    self.node_texts[new_node_id] = text
    self._update_prediction()
    reward = self.current_confidence - old_confidence
    done = self.current_confidence > 0.5
    return observation, reward, done
```

- EditNews(node_id, new_text):

```
def modify_news(self, node_id: Any, new_text: str) -> Tuple[Dict, float, bool]:
    """Modify root news content (Publisher only)."""
    self.node_texts[node_id] = new_text
    self.graph.nodes[node_id]['content'] = new_text
    self._update_prediction()
    reward = self.current_confidence - old_confidence
    done = self.current_confidence > 0.5
    return observation, reward, done
```

- EditComment(node_id, new_text):

```
def modify_comment(self, node_id: Any, new_text: str) -> Tuple[Dict, float, bool]:
    """Modify comment content (Publisher only)."""
    self.node_texts[node_id] = new_text
    self.graph.nodes[node_id]['content'] = new_text
    self._update_prediction()
    reward = self.current_confidence - old_confidence
    done = self.current_confidence > 0.5
    return observation, reward, done
```

- DeleteComment(node_id):

```
def remove_node(self, node_id: Any) -> Tuple[Dict, float, bool]:
    """Remove comment and its subtree (Publisher only)."""
    subtree = self._get_subtree(node_id) # includes all descendants
    for nid in subtree:
        self.graph.remove_node(nid)
        del self.node_texts[nid]
    self._update_prediction()
    reward = self.current_confidence - old_confidence
    done = self.current_confidence > 0.5
    return observation, reward, done
```

Table 25: MPT-based execution tools implementation.