

# Tapas Are Free! Training-Free Adaptation of Programmatic Agents via LLM-Guided Program Synthesis in Dynamic Environments

Jinwei Hu<sup>1</sup>, Yi Dong<sup>1\*</sup>, Youcheng Sun<sup>2</sup>, Xiaowei Huang<sup>1</sup>

<sup>1</sup>School of Computer Science and Informatics, University of Liverpool, Liverpool, L69 3BX, the UK

<sup>2</sup>Department of Computer Science, Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, SE45 05, UAE  
{jinwei.hu, yi.dong, xiaowei.huang}@liverpool.ac.uk, youcheng.sun@mbzuai.ac.ae

## Abstract

Autonomous agents in safety-critical applications must continuously adapt to dynamic conditions without compromising performance and reliability. This work introduces **TAPA** (Training-free Adaptation of Programmatic Agents), a novel framework that positions large language models (LLMs) as intelligent moderators of the symbolic action space. Unlike prior programmatic agents typically generate a monolithic policy program or rely on fixed symbolic action sets, TAPA synthesizes and adapts modular programs for individual high-level actions, referred to as logical primitives. By decoupling strategic intent from execution, TAPA enables meta-agents to operate over an abstract, interpretable action space while the LLM dynamically generates, composes, and refines symbolic programs tailored to each primitive. Extensive experiments across cybersecurity and swarm intelligence domains validate TAPA's effectiveness. In autonomous DDoS defense scenarios, TAPA achieves 77.7% network uptime while maintaining near-perfect detection accuracy in unknown dynamic environments. In swarm intelligence formation control under environmental and adversarial disturbances, TAPA consistently preserves consensus at runtime where baseline methods fail. This work promotes a paradigm shift for autonomous system design in evolving environments, from policy adaptation to dynamic action adaptation.

## 1 Introduction

Autonomous agents have become increasingly prevalent across critical domains such as cyber defense (Lohn et al. 2023), swarm intelligence control (Duan and Wang 2025), and autonomous driving (Kiran et al. 2021), where they are tasked with making timely and reliable decisions under complex conditions. Recent advancements in symbolic AI and reinforcement learning (RL) have greatly enhanced these agents' ability to interact with their environments, enabling the acquisition of structured and abstract behaviors beyond pure rule-based control (Carnevali and Lippi 2024). The integration of neuro-symbolic approaches has emerged as the predominant paradigm for safety-critical applications, combining neural learning capabilities with symbolic rule reliability to overcome the black-box limitations of pure RL while maintaining the interpretability and formal guarantees

\*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

required in critical domains (Marton et al. 2024). In such hybrid systems, symbolic modules are responsible for concrete execution by employing symbolic programs or rules as elements of the action space, enabling agents termed *programmatic agents* to operate through interpretable programmatic actions (Lyu et al. 2019; Verma et al. 2019). Compared to traditional actions, programmatic symbolic action spaces provide superior interpretability, verifiability, and integration of expert knowledge, making them well-suited for real-world applications (Lelis, Chen, and Sun 2024).

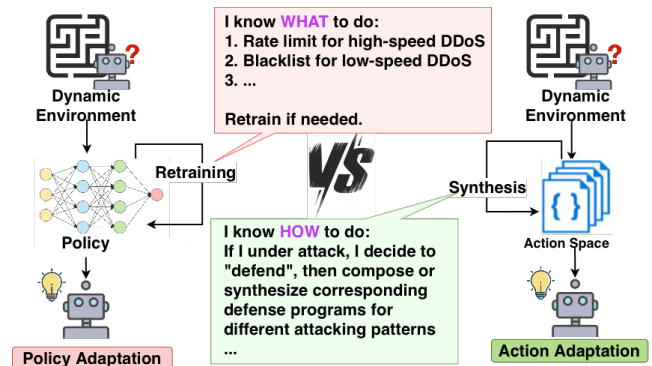


Figure 1: Policy-Level Retraining (left) vs. Action-Level Synthesis and Adaptation (right)

However, these static agent frameworks fundamentally lack *self-adaptive updating capabilities* and struggle to match the continuous changed characteristic of dynamic operational environments (Hu et al. 2025). The rapid evolution of attack vectors, environmental conditions, and task requirements in safety-critical scenarios further exacerbates these limitations, resulting in degraded performance and rigid decision-making that overlooks the crucial need for adaptiveness to rapidly respond to unforeseen changes and urgent threats (Bhuyan et al. 2024; Li et al. 2024; Hu, Dong, and Huang 2025). Therefore, we summarized the existing approaches fall short in three aspects: (i) limited adaptability due to rigid symbolic action spaces; (ii) prohibitive retraining costs requiring extensive strategy overhaul or manual rule updates; and (iii) inadequate safeguards when outdated symbolic actions are applied to changed environments.

Although Large Language Models (LLMs) offer new pos-

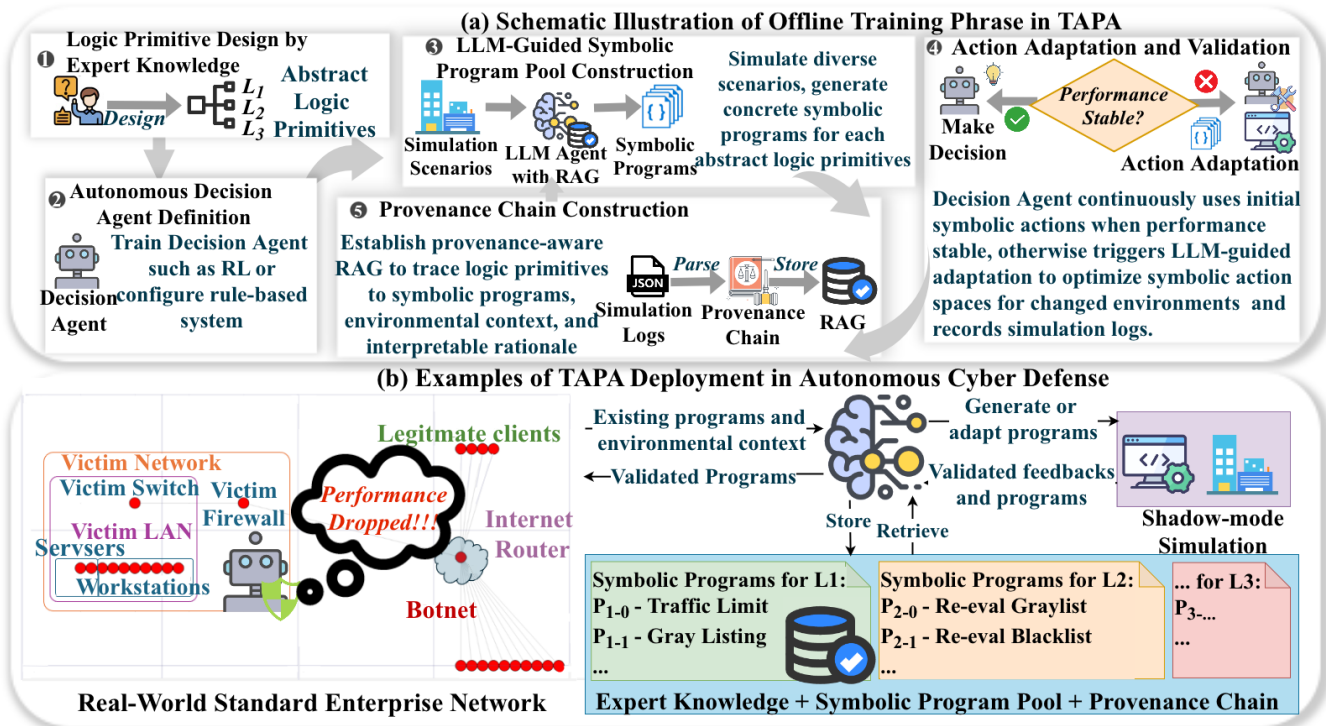


Figure 2: TAPA (Training-free Adaptation of Programmatic Agents) Framework. **(a) Design-time workflow.** TAPA enables autonomous agents to adapt to evolving environments without retraining through LLM-guided symbolic program synthesis: (1) *Logic primitive design.* Define high-level symbolic operations based on expert knowledge as interpretable strategic intent. (2) *Decision agent initialization.* A meta-agent is instantiated to select logical primitives based on environmental conditions. (3) *LLM-guided program pool construction.* LLM generates diverse symbolic programs across multiple simulated scenarios for each primitive. (4) *Action adaptation and validation.* When performance degrades, LLM synthesizes candidate programs for action adaptation and validated them through shadow simulation before replacement. (5) *Provenance chain construction.* Execution traces and adaptation experiences are stored in a Retrieval-Augmented Generation (RAG) system for future program synthesis. **(b) Deployment-time use case for cyber defense.** The TAPA-enabled agent monitors network performance, detects degradation, and retrieves or synthesizes validated programs as adaptive defensive operation in dynamic environments.

sibilities with their powerful reasoning and generalization capabilities (Hadi et al. 2023; Guo et al. 2024b), current research largely focuses on end-to-end LLM decision agent, overlooking their potential to agents’ adaptiveness. Moreover, LLMs struggle in safety-critical, latency-sensitive environments, such as Distributed Denial of Service (DDoS) attack mitigation in cybersecurity scenarios where their inference overhead and the possibility of hallucinated output lead to inferior real-time performance compared to well-trained RL agents (Hager et al. 2024; Castro et al. 2025). This reveals a critical gap where existing autonomous systems lack adaptability, and LLMs lack efficiency and reliability. Motivated by these complementary issues, we pursue a new paradigm shown in Figure 1: *Can we harnesses LLMs’ generative and reasoning capabilities for dynamic program synthesis and action space moderation, enabling rapid adaptation while preserving efficiency and reliability of well-trained policies or expert-defined configurations?*

To address this challenge, we propose the TAPA framework, illustrated in Figure 2. Overall, the key contributions of this paper are summarized as follows:

- We introduce LLMs as intelligent action space moderators for autonomous agents, achieving a paradigm shift from costly policy to efficient action adaptation.
- We propose TAPA, a training-free framework enabling agents to adapt to evolving environments through LLM-guided symbolic program synthesis while enhancing adaptability, interpretability and traceability.
- We validate TAPA across safety-critical scenarios including cyber defense and swarm control, demonstrating superior performance in DDoS mitigation and formation control under adversarial disturbances.

## 2 Related Work

### 2.1 Autonomous Decision Agents

Autonomous intelligent agents for real-time decision-making predominantly rely on traditional planning methods, reinforcement learning, or LLMs, each exhibiting critical limitations in dynamic environments requiring rapid adaptation. Traditional planning approaches such as STRIPS (Bylander 1994; Aineto, Jiménez, and Onaindia 2018), PDDL-

based planners (Höller et al. 2020), and Hierarchical Task Networks (Hogg, Kuter, and Munoz-Avila 2009; Kelly, Botea, and Koenig 2008) require comprehensive domain specifications with predefined action sets, relying heavily on expert knowledge for domain modeling. While theoretically robust, their static nature and labor-intensive requirements significantly constrain practicality when rapid operational adaptations are required (Bhuyan et al. 2024).

Neuro-symbolic approaches integrate RL with symbolic AI to enhance capabilities through environmental interaction. These hybrid methods allow RL agents to operate over symbolic action spaces, or guide symbolic planning using learned value functions (Yang et al. 2018; Lyu et al. 2019; Shindo et al. 2025). However, they encounter critical limitations in adaptability under adversarial or highly dynamic conditions, where symbolic priors become obsolete and RL agents struggle to generalize across different operational scales due to rigid logic constraints (Hakim et al. 2025; Hu et al. 2025). End-to-end LLM agents leverage natural language understanding to directly produce action sequences, enabling flexible reasoning and context-aware interaction (Yao et al. 2023). Despite their versatility, these agents face critical limitations in safety-critical deployments: inference latency hinders real-time responsiveness (Liang and Tong 2025), and outputs are prone to uncertainty and hallucination, undermining reliability even when guardrails are equipped (Dong et al. 2024, 2025). In contrast, our TAPA framework positions LLMs as intelligent action space moderators, dynamically adapting action via programs synthesis. This approach addresses the adaptability-efficiency trade-off by enabling real-time symbolic program evolution without costly retraining, allowing agents to safely handle arbitrary scenarios including unexpected emergent situations through symbolic program adaptation.

## 2.2 Large Language Models for Code Generation

Recent advances in LLMs have revolutionized code generation capabilities, with models like GPT (Achiam et al. 2023), Claude (Anthropic 2024), and Code Llama (Roziere et al. 2024) demonstrating impressive performance on standard benchmarks such as HumanEval and MBPP (Jiang et al. 2025). These developments have led to widespread adoption of LLM-based coding assistants in software development workflows, including GitHub Copilot and various IDE integrations (El Haji, Brandt, and Zaidman 2024). However, code generation in safety-critical domains presents unique challenges that differentiate it from general-purpose programming tasks (Soroush et al. 2024). While RAG-based approaches (Parvez et al. 2021; Zhou et al. 2023) have enhanced code generation with external knowledge, the scarcity of high-quality, domain-specific code samples in safety-critical domains means RAG can only offer limited descriptive guidance, resulting in practically ineffective code generation. Recent work has explored various methods to address domain-specific code generation challenges, including specialized training datasets and domain-specific fine-tuning (Guo et al. 2024a; Zan et al. 2023; Hou et al. 2025). However, these approaches primarily focus on improving syntactic and functional correctness rather than en-

suring domain-specific operational effectiveness in dynamic environments where code must adapt to rapidly evolving scenarios and threats. To enhance the practicality of generated programs, our framework constructs symbolic program pools through domain-specific multi-scenario simulation and employs provenance chains to guide LLM moderators in making informed decisions for symbolic program selection, composition, adaptation, and even new generation during runtime based on validated execution experiences.

## 3 Problem Formulation

Consider an autonomous agent operating in an environment  $E$  characterized by state space  $\mathcal{S}$ , action space  $\mathcal{A} = \{P_1, P_2, \dots, P_m\}$  where  $m$  denotes the number of available symbolic programs, and reward function  $R(s, p)$ . The agent maintains policy  $\pi : \mathcal{S} \rightarrow \mathcal{A}$  and aims to maximize cumulative rewards:

$$\max_{\pi} \mathbb{E}_{\tau \sim (E, \pi)} \left[ \sum_{t=0}^T \gamma^t R(s_t, a_t) \right] \quad (1)$$

where  $\tau$  represents a trajectory,  $T$  is the time horizon, and  $\gamma \in [0, 1]$  is the discount factor. However, real-world deployment environments for autonomous decision agents are typically dynamic and evolve rapidly, especially in safety-critical scenarios. When the environment evolves across multiple versions  $\mathcal{E} = \{E_1, E_2, \dots, E_V\}$ , where  $V$  denotes the number of environment variants, a fixed action space and its corresponding trained policies may become insufficient, thereby limiting the agent’s generalizability. Environmental evolution involves multiple dimensions such as scale, complexity, adversarial patterns, and operational constraints, causing significant performance degradation  $R_{v+1} \ll R_v$ , where  $v$  is the index of environmental version. To maintain agent stability and address the aforementioned challenges posed by environmental evolution, we reframe autonomous agent training or configuration by abstracting the concrete action space  $\mathcal{A}$  into a set of logical primitives  $\mathcal{L} = \{L_1, L_2, \dots, L_N\}$ , where  $N$  denotes the number of available logical actions.

**Definition 1** (Logical Primitive). *A logical primitive  $L_i \in \mathcal{L}, i \leq N$  represents a high-level strategic intent that specifies what the agent should accomplish in a given context, without prescribing how to implement. These primitives abstract over concrete actions and symbolic program implementations, enabling decision-making to remain generalizable and interpretable across evolving environments.*

Abstract policies over  $\mathcal{L}$  enhance generalizability by decoupling strategic intent from low-level execution, allowing the symbolic layer to adapt independently to environmental changes while preserving policy stability. Given this abstraction, the challenge becomes finding optimal symbolic program mappings for each logical primitive that enable performance stability in evolved environments. For an effective abstract policy  $\pi^* : \mathcal{S} \rightarrow \mathcal{L}$  and evolved environment  $E_{v+1}$ , we seek the optimal mapping:

$$\mathcal{M}^* = \arg \max_{\mathcal{M}} \mathbb{E}[R_{v+1} | \pi^*, E_{v+1}, \mathcal{M}(\mathcal{L})] \quad (2)$$

where  $\mathcal{M} : \mathcal{L} \rightarrow \mathcal{A}_{mod}$  maps each logical primitive to moderated symbolic programs, where  $\mathcal{A}_{mod}$  represents the space of all possible program compositions. Formally, for each logical primitive  $L_i$ :

$$\mathcal{M}(L_i) = \bigcup_j (P_j \odot \text{op}_j) \quad (3)$$

where  $P_j \in \mathcal{A}$  are selected symbolic programs,  $\odot$  represents composition operations that apply  $\text{op}_j$  to programs, and  $\text{op}_j \in \{\wedge, \vee, +, -, \delta\}$  denotes logical operations including conjunction, disjunction, addition, deletion, and modification applied to program  $P_j$ . This formulation shifts adaptation from expensive policy retraining to efficient symbolic action space moderation, ensuring adaptive decision-making in dynamic scenarios.

## 4 TAPA Framework

In the following, we describe each step of the TAPA framework shown in Figure 2 in details.

**Logic Primitive Design by Expert Knowledge** We begin by defining a set of domain-specific logical primitives  $\mathcal{L}$  according to Definition 1 based on scenario and task requirements. Each primitive represents a strategic intent derived from domain expertise and serves as interpretable abstractions over concrete actions, enabling more stable and generalizable strategy across varying environments.

**Example 1** (Logical Primitives in Cyber Defense). *In cybersecurity scenarios, experts typically define logical primitives as  $\mathcal{L} = \{\text{Observe}, \text{Defend}, \text{Validate}, \text{Alert}\}$ , where *Observe* monitors network traffic patterns, *Defend* applies countermeasures like filtering and rate limiting, *Validate* verifies threat authenticity, and *Alert* triggers human verification.*

**Decision Agent Initialization** Based on this abstraction, we formally define a meta-agent that operates over  $\mathcal{L}$  rather than directly over the raw action space as follows:

**Definition 2** (Meta-Agent and Policy). *A meta-agent is an autonomous agent that operates over abstract logical primitives rather than concrete actions, learning or being configured with a meta policy  $\pi_{meta} : S \rightarrow \mathcal{L}$  based on Eq. (1).*

Under this definition, the meta-policy can be optimized in a simplified environment  $E_{simple}$  by replacing the original action space  $\mathcal{A}$  with logical primitives  $\mathcal{L}$ , thereby enhancing generalizability across evolving environments.

**Example 2** (Meta-Agent in Cyber Defense). *Continuing from Example 1, a programmatic meta-agent learns to select appropriate logical primitives based on network states. For instance, when detecting abnormal traffic patterns, the meta-policy  $\pi_{meta}(s_t) = \text{Defend}$ , triggering the defensive intent rather than directly specifying which symbolic programs or rules to apply.*

### LLM-Guided Symbolic Program Pool Construction

We firstly construct diverse scenarios across multiple simulated environments to capture various operational conditions and challenges. For each logical primitive  $L_i$  and simulated environment  $E_v$  from the set of environmental versions  $\mathcal{E}$ ,

we employ an LLM augmented with domain-specific expert knowledge to generate candidate symbolic programs. The generation process is formalized as:

$$\mathcal{P}_{i,v} = \text{LLM}_{gen}(L_i, \xi(E_v), \mathcal{V}_{RAG}) \quad (4)$$

where  $\mathcal{P}_{i,v}$  represents the candidate program set for primitive  $L_i$  in environment  $E_v$ ,  $\xi(E_v)$  denotes environmental context features, and  $\mathcal{V}_{RAG}$  contains expert knowledge and experience from past simulations. The LLM agent leverages this knowledge to ensure generated programs are domain-compliant and contextually appropriate.

**Example 3** (Program Pool in Cyber Defense). *Continuing from Example 2, for the logical primitive  $L_1$  across multiple DDoS scenarios, the LLM agent generates a series of concrete programmatic programs such as  $\{P_{1,1}, P_{1,2}, P_{1,3}, \dots\}$  to construct the program pool. For instance,  $P_{1,1}$  implements rate limiting,  $P_{1,2}$  implements traffic filtering, and  $P_{1,3}$  implements blacklisting and rate limiting.*

**Action Adaptation and Validation** We conduct multi-scenario simulations to validate the symbolic programs in the constructed program pool. During this process, when performance degradation is detected typically through reward  $R(s_t, L_t)$  below threshold  $\xi$ , the LLM agent leverages its reasoning capabilities to invoke the RAG system and determine optimal symbolic program compositions following Eq. (2) to perform action adaptation by replacing the programs equipped for corresponding logical primitives. This process involves either composing existing programs using logical operations when similar cases exist in the RAG system, or synthesizing new programs guided by RAG-stored experiences when no comparable scenarios are found.

All candidate programs must undergo validation through shadow simulation (Kuchta, Palikareva, and Cadar 2018; Microsoft 2023), a standard practice in industrial deployments where new algorithms run in parallel with production systems without affecting live operations. Within this shadow simulation environment, when significant environmental shifts are detected, the framework can also generate several alternative meta-policies  $\{\pi_{meta}^{(1)}, \pi_{meta}^{(2)}, \dots, \pi_{meta}^{(k)}\}$  where  $k$  represents the number of viable strategic alternatives, as backup options to maintain policy stability and ensure uninterrupted decision-making in safety and time-critical scenarios. This holistic design guarantees safety and performance before deployment while enabling quick strategic switches based on real-time feedback.

**Example 4** (Action Adaptation in Cyber Defense). *Continuing from Example 3, when performance drops below threshold during DDoS scenarios, action adaptation is triggered. The LLM agent analyzes environmental and performance context (traffic volume: 2.3K packets/sec, network uptime: 22%) and retrieves similar cases from the RAG system to compose multiple potential programs  $\{P_{adapt}^{(1)}, P_{adapt}^{(2)}, \dots\}$ . After shadow simulation validation shows  $P_{adapt}^{(1)} = P_{1,1} \wedge P_{1,3}$  for  $L_1$  can recover performance to 72% (e.g., acceptable threshold), this adapted program replaces the original to serve as the new execution action for  $L_1$  while continuously searching for better-performing actions or backup strategy combinations.*

**Provenance Chain Construction** Each simulation generates a detailed provenance chain (PC) stored in the RAG system as experience for future program synthesis. We record comprehensive execution traces based on simulation logs, including invoked logical primitives, environmental context, original and modified programs, performance profiles, and adaptation rationales as shown in Figure 3 and Appendix B. The PC and constructed program pool together enable full traceability, facilitate system debugging, and enhance LLM-driven program generation by leveraging prior adaptation experiences, *addressing the scarcity of domain-specific program samples and experience*. This iterative optimization process across diverse scenarios progressively enriches the RAG system, establishing a comprehensive knowledge base for runtime adaptation.

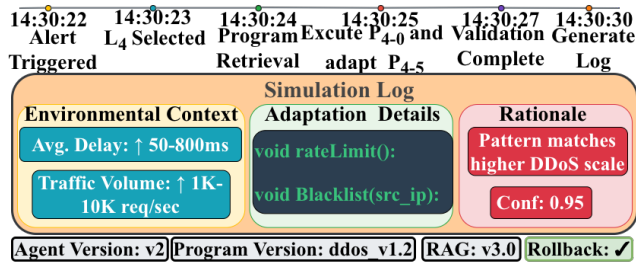


Figure 3: Provenance chain example for DDoS attack.

## 5 Experiments

To evaluate the effectiveness of TAPA across safety-critical domains, we conduct experiments on two representative cases: (1) *Autonomous cyber defense*, where agents must rapidly adapt to evolving DDoS attack patterns; and (2) *Swarm formation control*, where UAVs must maintain consensus under environmental disturbances and adversarial interference. These domains exemplify settings in which static symbolic systems struggle to adapt effectively (Dehghantanha, Yazdinejad, and Parizi 2023; Stodola, Nohel, and Horák 2025). Our results demonstrate that TAPA enables training-free action adaptation while preserving interpretability and operational efficiency. Implementation details for both domains are provided in Appendix A. In addition, we perform ablation studies to validate the design rationale behind key components of the framework.

### 5.1 Autonomous Cyber Defense

This cybersecurity defense use case is designed to evaluate the effectiveness of action adaptation under evolving DDoS attack patterns and changing network conditions, and to assess the agent’s ability to maintain high detection accuracy, low false positives, and sustained network uptime (without policy retraining, which typically needs hours to adapt to unknown environments).

**Experimental Setup** We conduct our experiments using the NS3AI simulation platform (Yin et al. 2020), which combines high-fidelity network simulation (Henderson et al.

2008) with Python-based AI frameworks via shared memory linking, enabling real-time simulation and protocol-level traffic modeling. Our network topology mirrors standard enterprise architectures used in cyber defense testbeds (e.g., CIC-IDS2017, CIC-DDoS2019) (Sharafaldin et al. 2018), comprising a victim LAN with router/firewall, legitimate TCP clients, and a botnet launching high-rate attacks. To evaluate defense performance, we design five progressively complex scenarios (E2–E5) varying in attack sophistication, network size, and operational constraints, as summarized in Table 1. Following the NIST and MITRE cybersecurity configuration (Möller 2023), we model DDoS mitigation through four core defensive operations abstracted as logical primitives as shown in Example 1. We compare TAPA with several baseline methods, all of which are trained or configured in environment **E1** and then evaluated in environments **E2–E5** without prior exposure. These baselines operate with fixed action spaces consisting of one program per abstract intent defined in the example, denoted as  $\mathcal{A} = \{P_{0,0}, P_{1,0}, P_{2,0}, P_{3,0}\}$ . In contrast, while TAPA also starts with the same initial program set  $\mathcal{A}$ , it activates action adaptation for the meta-agent when abnormal network performance is detected across consecutive time windows to optimize or replace the equipped programs.

Env	Servers	Bots	Clients	Attack Type
E1	10	10	30	TCP Flood
E2	10	10	30	UDP Flood
E3	10	20	30	Mixed Attack
E4	10	10	80	UDP Flood
E5	20	40	80	UDP Flood

Table 1: Environmental Evolution Scenarios Configuration

**Baselines & Evaluations** The baseline methods include:

- *Static Symbolic*. A symbolic rule-based system with no learning or adaptation.
- *Symbolic-Classic*. A Q-learning agent operating over predefined symbolic actions.
- *Symbolic-Neural*. A Transformer-in-Transformer (TiT) based neural policy network (Mao et al. 2022).
- *End-to-End LLM*. A GPT-4o-based agent that makes real-time decisions.

Evaluation is conducted using three key metrics: (1) *Network uptime*: the percentage of time the victim server remains responsive; (2) *Detection accuracy (Acc)*: the proportion of correctly identified bot nodes; (3) *False positives (FP)*: the proportion of legitimate nodes mistakenly flagged as malicious.

**Effectiveness analysis of Action Adaptation** Table 2 validates our action adaptation strategy, demonstrating that TAPA achieves superior overall uptime without costly retraining compared to policy adaptation approaches. While pure symbolic AI and neural-symbolic methods show promising results in baseline environments (E1), they exhibit gradual performance degradation in evolved environments (E2-E5) without rule redefinition or policy retraining,

Method	E1		E2		E3		E4		E5		Overall Network Uptime
	Acc	FP	Acc	FP	Acc	FP	Acc	FP	Acc	FP	
Static Symbolic	100.0	16.7	100.0	36.7	100.0	33.3	70.0	15.0	15.0	15.0	48.1
Symbolic-Classic	90.0	10.0	100.0	30.0	85.0	50.0	80.0	45.0	92.5	87.5	55.7
Symbolic-Neural	100.0	6.7	100.0	16.7	100.0	43.3	60.0	63.8	97.5	56.3	70.5
End-to-End LLM	60.0	43.3	70.0	26.7	65.0	33.3	60.0	27.5	45.0	46.3	27.6
TAPA (Ours)	100.0	0.0	100.0	0.0	100.0	0.0	100.0	0.0	100.0	12.5	77.7

Table 2: Detection Accuracy, False Positives, and Overall Network Uptime (%) across Evolutionary Environments.

highlighting the brittleness of fixed action spaces. Figure 4 further illustrates this contrast: the fixed neural-symbolic approach exhibits sparse defense activations concentrated around attack peaks, reflecting static symbolic modules’ limitations under increased environmental complexity. Although policy retraining can relearn alternative strategies, such as sustained defend or continuous validate behaviors for improved network management, it lacks the timeliness and responsiveness of our action adaptation mechanism. In contrast, TAPA achieves adaptive and frequent defensive activations through dynamic symbolic program synthesis, with notably superior suspicious list management that maintains dynamic regulation rather than monotonic growth.




Figure 4: Temporal evolution of defensive patterns under static vs. adaptive action spaces. Blue regions indicate normal state; red regions indicate attack periods.

Overall, our TAPA framework introduces LLM generalization capability to facilitate action adaptation under evolving environments without over-reliance on perfectly predefined symbolic modules. The adaptive mechanism consistently maintains higher performance across dynamic cybersecurity environments through context-aware program synthesis and refinement, while the meta-policy enables rapid execution. Its balanced, neuro-symbolic paradigm addresses the inherent limitations of static action spaces, constrained generalization capacity, and inefficient decision-making processes, providing a principled and scalable solution for training-free autonomous adaptation in safety-critical dynamic environments.

## 5.2 Formation Control of Swarm Intelligence

Formation control of swarm intelligence enables multiple autonomous agents (e.g., aircraft or robots) to main-

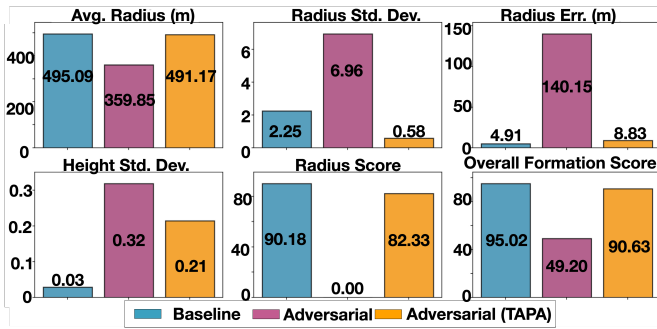
tain targeted formations through well-defined control algorithms. However, even carefully designed algorithms face significant limitations in dynamic environments with changing conditions, environmental disturbances, and adversarial threats. TAPA addresses these challenges by enabling real-time program synthesis for swarm coordination, dynamically adapting control parameters and logic. We evaluate TAPA against baseline methods across diverse scenarios, demonstrating its effectiveness in maintaining formation coherence under varying operational conditions.

**Experimental Setup** We evaluate TAPA in multi-aircraft formation control, where 10 aircraft maintain circular formation while adapting to environmental changes and adversarial threats. The baseline strategies follow Reynolds’ flocking rules (Reynolds 1987; Braga et al. 2017; Olfati-Saber and Murray 2004) and defensive strategies as implemented in (HU et al. 2025). Performance is evaluated across scenarios in Table 3 using formation radius and height metrics. We model formation control through two logical primitives:  $L_1$  (Formation Control) for adjusting coordination parameters like cohesion distance and behavioral weights, and  $L_2$  (Defend) for implementing defensive programs including outlier detection and noise injection for adversarial resilience. Each aircraft uses classical flocking behaviors (separation, cohesion, alignment, goal-seeking) through distributed decision-making within communication ranges, while TAPA enables real-time adjustment of formation variables including behavioral weights, distance thresholds, and defensive programs. When overall formation score based on radius and height continuously below threshold, TAPA’s LLM moderator analyzes environmental context and synthesizes appropriate parameter or program adaptations guided by expert knowledge and provenance chain experiences.

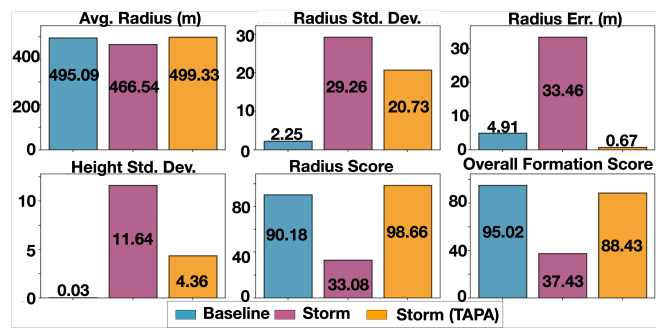
Env	Aircraft	Scenario description
E1	10	Stable weather, No any disturbance
E2	10	Severe weather with wind and rain affecting flight dynamics
E3	10	Malicious aircraft propagating infection to disrupt neighbor decisions

Table 3: Environmental Scenarios Configuration

**Formation Adaptation Analysis** Figure 6 and Appendix C visualizes consensus achievement across scenarios, comparing baseline formation control, scenarios without TAPA,



(a) Adversarial scenario with malicious UAV attacks



(b) Storm weather scenario with environmental disturbances

Figure 5: Formation control performance across changing scenarios. TAPA consistently recovers performance in both (a) adversarial environments with malicious interference and (b) severe weather conditions, maintaining near-baseline formation quality.

and TAPA-enabled adaptation through 3D trajectory visualization. Figure 5 provides quantitative analysis, illustrating TAPA’s consistent performance recovery across adversarial attacks and storm weather conditions. The results demonstrate that while baselines fail to maintain formation consensus under environmental changes, TAPA successfully triggers multiple action adaptations in dynamic environments, continuously adjusting programs and parameters to preserve aircraft consensus. TAPA’s ability to maintain near-baseline performance while adapting to diverse operational challenges validates symbolic program synthesis as an effective approach for safety-critical swarm intelligence applications requiring effectiveness, interpretability, and adaptiveness.

Without EK, the system exhibits trial-and-error learning patterns, starting poorly but gradually improving through iterative discovery. Conversely, without PC, the system achieves reasonable initial performance due to expert-guided program generation but lacks systematic experience accumulation, causing performance to degrade over time as it cannot adapt upon previous experiences. When relying purely on LLM reasoning capabilities without either component, the system fundamentally fails to achieve meaningful improvement, as evidenced by near-zero performance throughout all rounds. These results demonstrate our framework achieved sustained improvement by combining domain expertise for reliable initialization with systematic experience learning mechanisms for continuous adaptation.

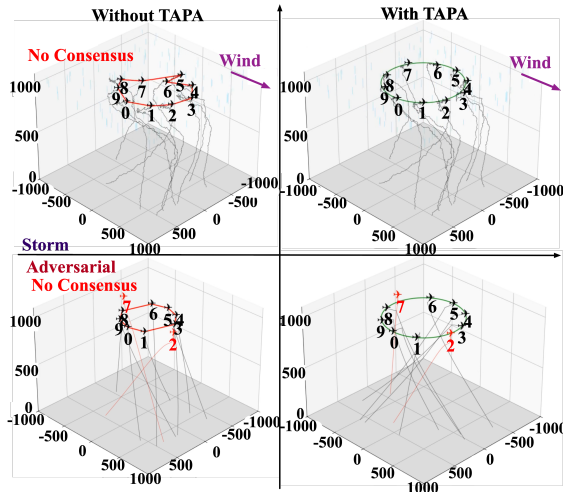


Figure 6: Consensus visualization of aircraft formation.

### 5.3 Ablation Study

To validate the design rationale of our framework components, we conduct ablation studies on formation control under storm conditions, where the system dynamically adjusts aircraft’s programs and parameters across multiple adaptation rounds. Table 4 reveals the complementary roles of Expert Knowledge (EK) and PC in adaptation effectiveness.

Configuration	Round 1	Round 2	Round 3
<b>TAPA-Full</b>	72.3 (+72.3)	81.7 (+9.4)	86.2 (+4.5)
<b>w/o PC</b>	65.1 (+65.1)	68.9 (+3.8)	62.4 (-6.5)
<b>w/o EK</b>	38.6 (+38.6)	54.2 (+15.6)	62.7 (+8.5)
<b>w/o Both</b>	4.2 (+4.2)	6.8 (+2.6)	0.0 (-6.8)

Table 4: Multi-Round Ablation on Formation Adaptation

## 6 Conclusion

We presented TAPA, a novel framework for training-free adaptation of programmatic agents through LLM-guided action space moderation. By positioning LLMs as intelligent moderators rather than direct decision-makers, TAPA achieves superior adaptability without costly retraining while preserving interpretability and efficiency. Our framework decouples strategic intent from execution through logical primitive abstraction, defining execution as interpretable programmatic actions to enable dynamic symbolic program synthesis in response to environmental changes. Experimental validation across cyber defense and swarm intelligence demonstrates TAPA’s effectiveness in real-world scenarios. This paradigm opens new avenues for developing adaptive autonomous systems that evolve with their environments while maintaining reliability for symbolic modules.

## Acknowledgments

This work is partially funded by the European Union (under grant agreement ID 101212818). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Health and Digital Executive Agency (HADEA). Neither the European Union nor the granting authority can be held responsible for them. This work is partially supported by Innovate UK through AI-PASSPORT under Grant 10126404. Yi's contribution is partially supported through the Royal Society international exchanges programme and in part by the Engineering and Physical Sciences Research Council, through funding from RAi UK [EP/Y009800/1].

## References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Aineto, D.; Jiménez, S.; and Onaindia, E. 2018. Learning STRIPS action models with classical planning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 28, 399–407.
- Anthropic. 2024. Claude 3 Technical Report. <https://www.anthropic.com/news/claude-3-family>. Accessed: 2025-07-20.
- Bhuyan, B. P.; Ramdane-Cherif, A.; Tomar, R.; and Singh, T. 2024. Neuro-symbolic artificial intelligence: a survey. *Neural Computing and Applications*, 36(21): 12809–12844.
- Braga, R. G.; Da Silva, R. C.; Ramos, A. C.; and Mora-Camino, F. 2017. Collision avoidance based on Reynolds rules: A case study using quadrotors. In *Information Technology-New Generations: 14th International Conference on Information Technology*, 773–780. Springer.
- Bylander, T. 1994. The computational complexity of propositional STRIPS planning. *Artificial Intelligence*, 69(1-2): 165–204.
- Carnevali, L.; and Lippi, M. 2024. Neuro-Symbolic Artificial Intelligence for Safety Engineering. In *International Conference on Computer Safety, Reliability, and Security*, 438–445. Springer.
- Castro, S. R.; Campbell, R.; Lau, N.; Villalobos, O.; Duan, J.; and Cardenas, A. A. 2025. Large Language Models are Autonomous Cyber Defenders. In *CAI*, 1125–1132.
- Dehghantanha, A.; Yazdinejad, A.; and Parizi, R. M. 2023. Autonomous cybersecurity: Evolving challenges, emerging opportunities, and future research trajectories. In *Proceedings of the Workshop on Autonomous Cybersecurity*, 1–10.
- Dong, Y.; Mu, R.; Jin, G.; Qi, Y.; Hu, J.; Zhao, X.; Meng, J.; Ruan, W.; and Huang, X. 2024. Position: Building Guardrails for Large Language Models Requires Systematic Design. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, 11375–11394. PMLR.
- Dong, Y.; Mu, R.; Zhang, Y.; Sun, S.; Zhang, T.; Wu, C.; Jin, G.; Qi, Y.; Hu, J.; Meng, J.; et al. 2025. Safeguarding large language models: A survey. *Artificial Intelligence Review*, 58(12): 382.
- Duan, Z.; and Wang, J. 2025. Enhancing multi-agent consensus through third-party llm integration: Analyzing uncertainty and mitigating hallucinations in large language models. In *International Conference on Advanced Algorithms and Control Engineering*, 2222–2227. IEEE.
- El Haji, K.; Brandt, C.; and Zaidman, A. 2024. Using github copilot for test generation in python: An empirical study. In *AST 2024*, 45–55.
- Guo, C.; Liu, X.; Xie, C.; Zhou, A.; Zeng, Y.; Lin, Z.; Song, D.; and Li, B. 2024a. Redcode: Risky code execution and generation benchmark for code agents. *Advances in Neural Information Processing Systems*, 37: 106190–106236.
- Guo, T.; Chen, X.; Wang, Y.; Chang, R.; Pei, S.; Chawla, N. V.; Wiest, O.; and Zhang, X. 2024b. Large language model based multi-agents: a survey of progress and challenges. In *IJCAI*, 8048–8057.
- Hadi, M. U.; Qureshi, R.; Shah, A.; Irfan, M.; Zafar, A.; Shaikh, M. B.; Akhtar, N.; Wu, J.; Mirjalili, S.; et al. 2023. A survey on large language models: Applications, challenges, limitations, and practical usage. *Authorea Preprints*.
- Hager, P.; Jungmann, F.; Holland, R.; Bhagat, K.; Hubrecht, I.; Knauer, M.; Vielhauer, J.; Makowski, M.; Braren, R.; Kaissis, G.; et al. 2024. Evaluation and mitigation of the limitations of large language models in clinical decision-making. *Nature medicine*, 30(9): 2613–2622.
- Hakim, S. B.; Adil, M.; Velasquez, A.; and Song, H. H. 2025. ANSR-DT: An Adaptive Neuro-Symbolic Learning and Reasoning Framework for Digital Twins. *arXiv preprint arXiv:2501.08561*.
- Henderson, T. R.; Lacage, M.; Riley, G. F.; Dowell, C.; and Koppena, J. 2008. Network simulations with the ns-3 simulator. *SIGCOMM demonstration*, 14(14): 527.
- Hogg, C.; Kuter, U.; and Munoz-Avila, H. 2009. Learning Hierarchical Task Networks for Nondeterministic Planning Domains. In *IJCAI*, 1708–1714.
- Höller, D.; Behnke, G.; Bercher, P.; Biundo, S.; Fiorino, H.; Pellier, D.; and Alford, R. 2020. HDDL: An extension to PDDL for expressing hierarchical planning problems. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 9883–9891.
- Hou, Z.; Liu, H.; Bian, J.; He, X.; and Zhuang, Y. 2025. Enhancing medical coding efficiency through domain-specific fine-tuned large language models. *npj Health Systems*, 2(1): 14.
- Hu, J.; Dong, Y.; Ao, S.; Li, Z.; Wang, B.; Singh, L.; Cheng, G.; Ramchurn, S. D.; and Huang, X. 2025. Stop Reducing Responsibility in LLM-Powered Multi-Agent Systems to Local Alignment. *arXiv:2510.14008*.
- HU, J.; DONG, Y.; DING, Z.; and HUANG, X. 2025. Enhancing robustness of LLM-driven multi-agent systems through randomized smoothing. *Chinese Journal of Aeronautics*, 103779.
- Hu, J.; Dong, Y.; and Huang, X. 2025. Trust-Oriented Adaptive Guardrails for Large Language Models. *arXiv:2408.08959*.

- Hu, J.; Tang, Z.; Jin, X.; Zhang, B.; Dong, Y.; and Huang, X. 2025. Hierarchical Testing With Rabbit Optimization for Industrial Cyber-Physical Systems. *IEEE Transactions on Industrial Cyber-Physical Systems*, 3: 472–484.
- Jiang, J.; Wang, F.; Shen, J.; Kim, S.; and Kim, S. 2025. A Survey on Large Language Models for Code Generation. *TOSEM*. Just Accepted.
- Kelly, J.-P.; Botea, A.; and Koenig, S. 2008. Offline planning with hierarchical task networks in video games. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 4, 60–65.
- Kiran, B. R.; Sobh, I.; Talpaert, V.; Mannion, P.; Al Sallab, A. A.; Yogamani, S.; and Pérez, P. 2021. Deep reinforcement learning for autonomous driving: A survey. *IEEE transactions on intelligent transportation systems*, 23(6): 4909–4926.
- Kuchta, T.; Palikareva, H.; and Cadar, C. 2018. Shadow symbolic execution for testing software patches. *TOSEM*, 27(3): 1–32.
- Lelis, L.; Chen, X.; and Sun, S. 2024. Generating Programmatic Solutions: Algorithms and Applications of Programmatic Reinforcement Learning and Code Generation. In *NeurIPS 2024 Tutorial*. East Exhibition Hall A, NeurIPS.
- Li, B.; Li, Z.; Du, Q.; Luo, J.; Wang, W.; Xie, Y.; Stepputtis, S.; Wang, C.; Sycara, K.; Ravikumar, P.; et al. 2024. LogiCity: Advancing neuro-symbolic ai with abstract urban simulation. *Advances in Neural Information Processing Systems*, 37: 69840–69864.
- Liang, G.; and Tong, Q. 2025. LLM-Powered AI Agent Systems and Their Applications in Industry. *arXiv preprint arXiv:2505.16120*.
- Lohn, A.; Knack, A.; Burke, A.; and Jackson, K. 2023. Autonomous Cyber Defense. *A roadmap from lab to ops. Online. Centre for Emerging Technology and Security (CETaS) at The Alan Turing Institute*.
- Lyu, D.; Yang, F.; Liu, B.; and Gustafson, S. 2019. SDRL: interpretable and data-efficient deep reinforcement learning leveraging symbolic planning. In *AAAI*, volume 33, 2970–2977.
- Mao, H.; Zhao, R.; Chen, H.; Hao, J.; Chen, Y.; Li, D.; Zhang, J.; and Xiao, Z. 2022. Transformer in Transformer as Backbone for Deep Reinforcement Learning. *arXiv preprint arXiv:2212.14538*.
- Marton, S.; Grams, T.; Vogt, F.; Lüdtke, S.; Bartelt, C.; and Stuckenschmidt, H. 2024. SYMPOL: Symbolic Tree-Based On-Policy Reinforcement Learning. *arXiv preprint arXiv:2408.08761*.
- Microsoft. 2023. Shadow Testing - Engineering Fundamentals Playbook. <https://microsoft.github.io/code-with-engineering-playbook/automated-testing/shadow-testing/>. Accessed: 2025-07-18.
- Möller, D. P. 2023. NIST cybersecurity framework and MITRE cybersecurity criteria. In *Guide to Cybersecurity in Digital Transformation: Trends, Methods, Technologies, Applications and Best Practices*, 231–271. Springer.
- Olfati-Saber, R.; and Murray, R. M. 2004. Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on automatic control*, 49(9): 1520–1533.
- Parvez, M. R.; Ahmad, W.; Chakraborty, S.; Ray, B.; and Chang, K.-W. 2021. Retrieval Augmented Code Generation and Summarization. In *EMNLP Findings*, 2719–2734.
- Reynolds, C. W. 1987. Flocks, herds and schools: A distributed behavioral model. In *Annual conference on Computer graphics and interactive techniques*, 25–34.
- Roziere, B.; Gehring, J.; Gloeckle, F.; Sootla, S.; Gat, I.; Tan, X. E.; Adi, Y.; Liu, J.; Sauvestre, R.; Remez, T.; et al. 2024. Code llama: Open foundation models for code. *LLM4Code*.
- Sharafaldin, I.; Lashkari, A. H.; Ghorbani, A. A.; et al. 2018. Toward generating a new intrusion detection dataset and intrusion traffic characterization. *International Conference on Information Systems Security and Privacy*.
- Shindo, H.; Delfosse, Q.; Dhami, D. S.; and Kersting, K. 2025. BlendRL: A Framework for Merging Symbolic and Neural Policy Learning. In *Proceedings of the 13th International Conference on Learning Representations*.
- Soroush, A.; Glicksberg, B. S.; Zimlichman, E.; Barash, Y.; Freeman, R.; Charney, A. W.; Nadkarni, G. N.; and Klang, E. 2024. Large language models are poor medical coders—benchmarking of medical code querying. *NEJM AI*, 1(5): AIdbp2300040.
- Stodola, P.; Nohel, J.; and Horák, L. 2025. Dynamic reconnaissance operations with UAV swarms: adapting to environmental changes. *Scientific Reports*, 15(1): 15092.
- Verma, A.; Le, H.; Yue, Y.; and Chaudhuri, S. 2019. Imitation-projected programmatic reinforcement learning. *Advances in Neural Information Processing Systems*, 32.
- Yang, F.; Lyu, D.; Liu, B.; and Gustafson, S. 2018. PE-ORL: integrating symbolic planning and hierarchical reinforcement learning for robust decision-making. In *IJCAI*, 4860–4866.
- Yao, S.; Zhao, J.; Yu, D.; Du, N.; Shafran, I.; Narasimhan, K. R.; and Cao, Y. 2023. ReAct: Synergizing Reasoning and Acting in Language Models. In *The Eleventh International Conference on Learning Representations*.
- Yin, H.; Liu, P.; Liu, K.; Cao, L.; Zhang, L.; Gao, Y.; and Hei, X. 2020. ns3-ai: Fostering artificial intelligence algorithms for networking research. In *Proceedings of the 2020 Workshop on ns-3*, 57–64.
- Zan, D.; Chen, B.; Zhang, F.; Lu, D.; Wu, B.; Guan, B.; Yongji, W.; and Lou, J.-G. 2023. Large Language Models Meet NL2Code: A Survey. In *ACL*, 7443–7464.
- Zhou, S.; Alon, U.; Xu, F. F.; Jiang, Z.; and Neubig, G. 2023. DocPrompting: Generating Code by Retrieving the Docs. In *ICLR*.