

RACONTEUR: A Knowledgeable, Insightful, and Portable LLM-Powered Shell Command Explainer

Jiangyi Deng^{1♣}, Xinfeng Li^{1♣}, Yanjiao Chen^{1✉}, Yijie Bai¹, Haiqin Weng²,
Yan Liu², Tao Wei², Wenyuan Xu¹

¹Zhejiang University, ²Ant Group

{jydeng, xinfengli, chenyanjiao, baiyj, wyxu}@zju.edu.cn

{haiqin.wenghaiqin, bencao.ly, lenx.wei}@antgroup.com

Abstract—Malicious shell commands are linchpins to many cyber-attacks, but may not be easy to understand by security analysts due to complicated and often disguised code structures. Advances in large language models (LLMs) have unlocked the possibility of generating understandable explanations for shell commands. However, existing general-purpose LLMs suffer from a lack of expert knowledge and a tendency to hallucinate in the task of shell command explanation. In this paper, we present RACONTEUR, a knowledgeable, expressive and portable shell command explainer powered by LLM. RACONTEUR is infused with professional knowledge to provide comprehensive explanations on shell commands, including not only what the command does (*i.e.*, behavior) but also why the command does it (*i.e.*, purpose). To shed light on the high-level intent of the command, we also translate the natural-language-based explanation into standard *technique & tactic* defined by MITRE ATT&CK, the worldwide knowledge base of cybersecurity. To enable RACONTEUR to explain unseen private commands, we further develop a documentation retriever to obtain relevant information from complementary documentations to assist the explanation process. We have created a large-scale dataset for training and conducted extensive experiments to evaluate the capability of RACONTEUR in shell command explanation. The experiments verify that RACONTEUR is able to provide high-quality explanations and in-depth insight of the intent of the command.

I. INTRODUCTION

Ongoing cyber-attacks are significant threats to an organization's data, applications and other valuable assets [52], [21], [23]. Malicious shell commands often act as springboards for attackers to gain remote control of victim computers. To thwart such attempts, intrusion detection systems (IDS) usually capture and alert suspicious shell command for further investigation by human analysts. Manual auditing on alerted shell commands is a key part of security operation. Security analysts are expected to fully understand the commands, from their literal semantics to high-level intents. However, the complex syntax and stealthy nature of malicious commands make it difficult for junior technicians and even security experts to understand.

♣ Both authors contributed equally to the paper.

The recent breakthrough in large language models (LLMs) has unlocked the potential of code explanation for programmers and non-programmers. It has been found that commercial LLMs (*e.g.*, ChatGPT) can generate understandable and fairly accurate explanations for commonly-used commands via user prompts [49], [33], [34], [10], [25]. Nonetheless, when presented with complicated or unseen commands, general-purpose LLMs will provide erroneous and even hallucinated explanations. This situation may be worsened for companies that use special shell commands containing private utilities or sensitive parameters. For example, Samsung has banned the use of ChatGPT and other AI-powered chatbots after an accidental leakage of sensitive internal source code [48].

In this paper, we develop an LLM-powered shell command explainer, named RACONTEUR, which provides knowledgeable and insightful explanations on shell commands, especially malicious ones. RACONTEUR is also portable to private shell commands that are absent from its training dataset. These design goals are realized through addressing the following challenges.

- *How to endow a general-purpose LLM with expert knowledge in shell command explanation?*

Existing general-purpose LLMs have unsatisfying performance on shell command explanation, especially for malicious commands. Fine-tuning a general-purpose LLM is a promising and affordable solution, but there is a lack of domain data for the task of shell command explanation. To tackle this problem, we construct a fine-tuning dataset of $\langle \text{prompt}, \text{response} \rangle$ pairs. To account for diversity in user prompts, we design both rule-based and model-based methods to generate a large variety of prompts. The *responses* are composed with professional knowledge from prestigious code libraries. After fine-tuning, RACONTEUR is able to explain the action of each step in the shell command as well as summarize the overall behavior of the shell command.

- *How to provide insightful interpretations of the intent of a shell command?*

An in-depth understanding of commands not only includes step-by-step description of the semantics but also an abstraction of the intents, *i.e.*, to understand what the attacker wants to achieve (*tactic*) by what means (*technique*). Translating the natural-language-based description into the

technique & tactic defined by the MITRE ATT&CK framework helps analysts quickly seek for corresponding mitigations. For example, the semantic explanation of the command `rundll32.exe keymgr, KRShowKeyMgr` may be “The command executes the exported function `KRShowKeyMgr` located in `keymgr.dll` using `rundll32.exe`.” It is more helpful if analysts also understand the technique behind this command, *i.e.*, to export the stored Windows credentials from the credential manager to a file, which belongs to the “T1003 - OS Credential Dumping” technique. The ultimate intent of this action is “TA0006 - Credential Access” tactic in the MITRE ATT&CK framework. By pinpointing the technique and tactic, a wealth of prior knowledge can be harnessed to respond to this threat. Unfortunately, there is a gap between the natural-language-based explanation and the standard description given by MITRE ATT&CK documentations. To address this gap, we establish an embedding model that maps the description of RACONTEUR and that of MITRE ATT&CK into the same embedding space, based on which we match the command to the technique and tactic with the most similar descriptions.

- How to provide accurate information about private shell commands that are unseen in the training phase?

A large portion of commands executed in an organization are proprietary, involving utilities, parameters, and files that are developed within the organization and are for internal use only. An LLM may provide inaccurate and fictitious information when presented with unseen commands that are not publicly available. To enable RACONTEUR to be portable to private shell commands, we design a documentation retriever to glean relevant contexts in complementary documentations (*e.g.*, private documentations of the company) to assist the explanation process of RACONTEUR. In this way, RACONTEUR can analyze commands according to the provided documentations and provide faithful explanations.

We implement a fully-functional prototype of RACONTEUR and create a dataset to conduct extensive experiments to evaluate the capability of RACONTEUR. We compare RACONTEUR with three baseline models, *i.e.*, GPT-3.5-Turbo, GPT-4 and ChatGLM2-6B. Among them, GPT-4 is the most powerful commercial LLM so far. Our extensive experiments demonstrate that RACONTEUR achieves superior explanation performance than the baselines in both English and Chinese languages. RACONTEUR also performs well in technique and tactic identification, which helps shed light on the high-level intents of commands.

We summarize our main contributions as follows:

- We propose RACONTEUR, an LLM-powered shell command explainer that can provide knowledgeable and insightful descriptions on shell commands, especially malicious ones, to assist security analysts in identifying potential cyber-attacks.
- We equip RACONTEUR with a holistic toolkit including behavior explainer, intent identifier, and documentation retriever, allowing RACONTEUR to provide comprehensive

and faithful explanations on public and private shell commands.

- We have conducted extensive experiments to verify that RACONTEUR is able to provide high-quality explanations and in-depth insight of the intent of the command. The dataset we curated is open-source to boost further research in LLM-aided code explanation¹.

II. BACKGROUND AND MOTIVATION

A. Shell Command Explanation

Shell commands allow users to interact with operating systems, *e.g.*, enter commands and receive responses. Many crucial cyber-attacks, *e.g.*, masquerade attacks, are enabled by malicious shell commands [50]. Intrusion detection systems (IDS) often capture and report suspicious shell commands to security analysts for further confirmation. However, the complicated formations of shell commands usually make it difficult for junior technical personnel or even experts to truly understand its ultimate purpose. A shell command consists of *utilities* (a.k.a., *executables*) followed by *options* and *parameters*. As shown in the following benign and malicious shell commands, `bash` is a utility, `-c` is an option, and the entire subsequent string within the single quotes is a parameter.

A Benign Shell Command

```
bash -c 'exec 137<>/dev/tcp/ip_addr/port
&& echo -e "GET /_health/HTTP/1.1\r\nhost:
ip_addr" >&137 && grep ok -s -m 1 <&137'
```

A Malicious Reverse Shell Command

```
bash -c '0<&137-;exec 137<>/dev/tcp/ip_addr
/port;sh <&137 >&137 2>&137'
```

Technique & tactic. To analyze shell commands, especially malicious ones, it is essential to understand both the technique and the tactic. The *technique* is the concrete actions of the commands, *i.e.*, “how” the commands are executed step by step. The tactic is the ultimate purpose of the commands, *i.e.*, “why” the commands are performed. Take a shell command for running malicious codes as an example. The tactic is executing the malicious codes. In order to achieve this goal, the technique may involve abusing Unix shell scripts for execution, *e.g.*, the reverse shell example above. Identifying both the tactic and the technique of a shell command helps analysts better understand the purpose and the habits of attackers. The most commonly-used tactics and techniques of malicious commands are taxonomized in MITRE ATT&CK’s², a globally-accessible knowledge base according to real-world observations. MITRE ATT&CK defines 14 tactics and about 200 techniques, which can be leveraged to interpret malicious shell commands.

Automatic command explanation. To assist human experts in understanding the technique and tactic of a shell command, automated explanation systems have been developed. Current breakthrough in large language models (LLMs) also gives rise to LLM-based command explanation algorithms. However,

¹<https://raconteur-ndss.github.io/>

²<https://attack.mitre.org/>

there are two main challenges facing existing systems. First, LLM-based systems suffer from intrinsic defects, especially regarding *faithfulness* and *factuality*. The output of an LLM-based system may not faithfully follow user instructions (*i.e.*, the answer is irrelevant) or provide true facts (*i.e.*, the answer is incorrect) [35]. Second, the queried shell commands may be unseen to the explanation system trained with public datasets. For example, many companies design special shell utilities, *e.g.*, used within the enterprise. The shell parameters, *e.g.*, IP address (`ip_addr`) and port (`port`), may be private information of a company. These special shell commands are usually absent from the training dataset of a general command explanation system, leading to poor generalization problems.

B. Generative Language Models

Generative language models output texts given an input *prompt*, widely applied to translation, summarization, and question answering. With recent advances of large-scale pre-training, commercial LLMs such as ChatGPT exhibit remarkable performance across a broad spectrum of tasks [6].

A typical training pipeline of generative LLMs includes pre-training and fine-tuning [45]. Due to its prohibitive requirements of training data and computational resources [22], pre-training is usually performed by professional AI companies like OpenAI [45]. Fine-tuning is an important way of turning a general-purpose LLM into a domain-specific LLM, *e.g.*, for malicious command analysis. Specifically, a pre-trained LLM can be fine-tuned on a labelled dataset of $\langle \textit{prompt}, \textit{response} \rangle$ pairs for the target task in the given domain. In this way, the domain-specific LLM will achieve better performance on specific tasks.

Faithfulness and *factuality* are two serious concerns about LLMs, which means the output may diverge from the prompt (violating *faithfulness*) or misalign with established world knowledge (violating *factuality*) [35]. Pre-trained models, *e.g.*, ChatGPT, are shown to have a high degree of faithfulness but not factuality [26]. To enhance factuality, retrieval-augmented generation (RAG) [58] has been proposed, which complements the original user prompt with related facts and knowledge via web browsing [40], [31] or document retrieval [43]. The original user prompt is encoded into an embedding vector, which is used to retrieve relevant texts from available sources based on semantic similarity. The semantic similarity between two texts can be computed by measuring the distance between two embedding vectors. The user prompt and the retrieved texts are combined as a new prompt to be fed into the LLM to elicit more factual output.

C. A Real-World Motivating Example

In this part, we give an example of explaining a real-world malicious shell command by RACONTEUR. The command is presented in the above malicious reverse shell command example in §II-A. This malicious command can enable an attacker to gain remote control, which was captured by the IDS in real-world security operation.

The aim of the malicious shell command is to establish an illegal connection to a remote computer at `ip_addr:port` and redirect the input and output of the victim system's shell so that the attacker can access it remotely. As we can see in the example, the malicious shell command has a similar form to the benign shell command, *e.g.*, they both have IP address and port and redirect the standard input and output. To make matters worse, the malicious shell command contains obfuscating file descriptors (*e.g.*, `&137`) and complicated redirections (*e.g.*, `<&137 >&137 2>&137`). Therefore, it is hard for analysts to identify the ulterior motive of the command.

The explanation generated by RACONTEUR for this command is shown in the following example. The explanation consists of two sections, *i.e.*, behavior explanation and intent prediction. The behavior explanation section first describes the behavior of each action in the command (*i.e.*, **Step by step explanation**), and then summarizes the overall behavior of the command (*i.e.*, **Overall**). The intent prediction section (*i.e.*, **Tactic & Technique**) derives the technique and tactic of the command in a standard framework based on its behaviors. The holistic explanation provided by RACONTEUR enables analysts to quickly comprehend both semantics and intents of the command, assisting them in making informed decisions for further countermeasures.

An Explanation on the Malicious Shell Command

Step by step explanation:

- `bash -c`: This command is invoking the Bash shell and running a command within it.
- `0<&137-`: This is redirecting file descriptor 0 (stdin) to file descriptor 137 and closing file descriptor 137. This is a common technique used in reverse shells to redirect input/output to a network socket.
- `exec 137<>/dev/tcp/ip_addr/port`: This is creating a new file descriptor 137 and opening a TCP connection to the IP address `ip_addr` on port `port`. This is the reverse shell connection.
- `sh <&137 >&137 2>&137`: This is executing the shell command and redirecting input, output, and errors to file descriptor 137. This is the reverse shell payload.

Overall: The command is attempting to establish a reverse shell connection to a remote host and execute a shell command on that host. This could be used for malicious purposes such as remote access or data exfiltration.

Tactic: Execution **Technique:** Unix Shell

III. SYSTEM MODEL

We define the system model in terms of the actors, the assumptions of their capabilities, and the security objectives.

A. Actor

In the context of security operation within an organization, we refer to the party attempting to execute malicious shell commands for assets such as sensitive information or control

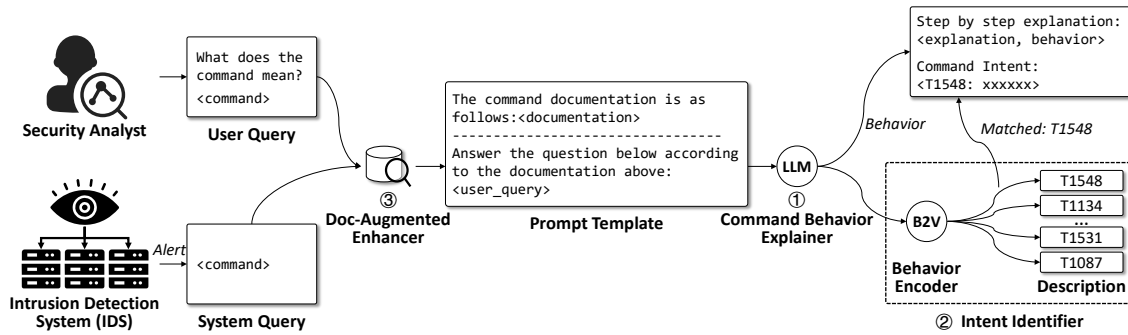


Fig. 1: Inference pipeline of RACONTEUR. The system receives queries from analysts and IDS. First, related documentations of the requested command are retrieved. Then a documentation-augmented prompt is created and fed into the command behavior explainer for analysis. The summarized behavior of the command is utilized for intent identification.

access as the *attacker*, and the party analyzing such suspicious commands as the *analyst*. In a typical security operation scenario, the malicious commands executed by the attacker are captured and flagged by the IDS for violating specific rules. Subsequently, this alert is sent to the analyst to confirm whether it constitutes a genuine attack.

B. Capability

Attacker. We assume that the attacker can operate the shell within the system, *e.g.*, they may be insiders with malicious intent or intruders who have gained initial access. They may design complex and potentially obfuscated commands to mislead the IDS and the analyst.

Analyst. We assume that the analyst possesses basic knowledge and skills in security operation, including the ability to comprehend the semantics of shell commands with the help of natural language explanations. Furthermore, the analyst can assess whether these commands exhibit malicious or benign behavior based on their descriptions in natural language. Additionally, we assume that the analyst can decipher obfuscated commands when their obfuscation is recognized, thereby revealing the underlying undisguised commands.

C. Security Objective

The primary goal of RACONTEUR is to aid analysts in understanding the semantics and intents of shell commands, encompassing Unix Shell and PowerShell. This goal can be delineated into four key security objectives:

- **Comprehensiveness.** This denotes that the explanations provided by RACONTEUR should facilitate the analyst’s comprehension of the command’s semantics.
- **Insightfulness.** This signifies that RACONTEUR should help the analyst to determine whether the command is malicious or benign, elucidating the tactics and techniques employed if the command is indeed malicious.
- **Correctness.** This indicates that the content generated by RACONTEUR should be factually correct even when handling proprietary and private commands.

- **Portability.** This implies that RACONTEUR should function within a local environment without necessitating reliance on online services.

It is important to note that our aim is not to directly classify whether a command is malicious or not, but rather to aid the analyst in making that determination.

IV. RACONTEUR: DETAILED CONSTRUCTION

As shown in Figure 1, RACONTEUR mainly consists of three core modules, *i.e.*, behavior explainer, intent identifier, and doc-augmented enhancer.

Behavior Explainer (①) takes the user prompt as input and outputs the behavior explanation section (*i.e.*, **Step by step explanation & Overall**). The behavior explainer is obtained by fine-tuning a general LLM with a carefully-constructed dataset consisting of $\langle \text{prompt}, \text{response} \rangle$ pairs. To transform the general LLM into an expert LLM in shell command explanation, we extract knowledge from code libraries to compose high-quality responses.

Intent Identifier (②) takes the behavior summary from the behavior explainer and outputs the intent prediction (*i.e.*, **Tactic & Technique**). By mapping behavior descriptions in natural language into the standardized MITRE ATT&CK framework, RACONTEUR assists analysts in further analyzing potential attacks.

Doc-Augmented Enhancer (③) improves the performance of the behavior explainer and the intent identifier by augmenting user prompt with relevant information retrieved from command documentations. To retrieve appropriate information, we design both rule-based and model-based methods to match the prompted shell command with useful texts in command documentations.

A. Behavior Explainer

The behavior explainer gives a detailed step-by-step explanation of the command, including descriptions of the utilities, options, and parameters. The behavior explainer can point out potential malicious attempt of the command in the overall behavior summary. When a malicious attempt is specified,

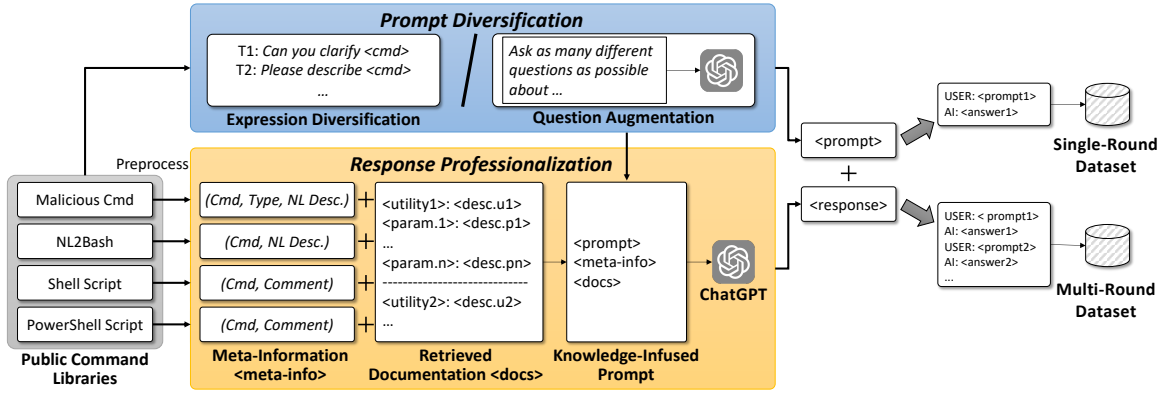


Fig. 2: Dataset generation pipeline of RACONTEUR. The prompt diversification module is utilized to generate prompts with diversified expressions and questions. In the response professionalization module, commands from different data sources are prompted with corresponding meta-information and command documentations for accurate responses.

the behavior explainer is designed to describe the type of the malicious attempt and the countermeasures to be taken.

Existing general-purpose LLMs are not professional enough for the task of command explanation. We leverage fine-tuning to establish an expert behavior explainer with a carefully-constructed dataset of $\langle \text{prompt}, \text{response} \rangle$ pairs. Given a prompt q_i , the fine-tuning process aims to maximize the conditional probability of producing the response r_i ,

$$\max_{\theta} \mathbb{E}_{\langle q_i, r_i \rangle \sim \mathbb{D}} p_{\theta}(r_i | q_i), \quad (1)$$

where θ denotes the parameters of the behavior explainer model, initialized by a pre-trained general-purpose LLM. \mathbb{D} is the fine-tuning dataset.

A major challenge is how to construct a high-quality fine-tuning dataset. First, user prompts have diversified forms of descriptions. Therefore, we need to enrich q_i to account for prompt diversity. Second, to ensure factuality of the explanation, we should infuse distilled expert knowledge into response r_i . We tackle these two problems by prompt diversification and response professionalization, as shown in Figure 2. Specifically, we diversify the expression and augment the question of prompts in prompt diversification module. For each prompt, we infuse both meta-information (extracted from the command libraries) and retrieved related documentation to construct a knowledge-infused prompt. Subsequently, we use an online LLM service (GPT-3.5-Turbo) to organize the expert knowledge we provide and generate accurate answers/responses. Single $\langle \text{prompt}, \text{response} \rangle$ pairs constitute the single-round dialogue dataset. Multiple $\langle \text{prompt}, \text{response} \rangle$ pairs are amalgamated within one dialogue to constitute the multi-round dialogue dataset. The details of the two modules are as follows.

1) *Prompt Diversification:* We consider two aspects for prompt diversification. First, users may use different expressions to try to obtain command explanations, e.g., “Please provide a detailed explanation for $\langle \text{command} \rangle$.” or “Could you please shed some light on $\langle \text{command} \rangle$?”. Second, apart from a general command explanation, users may be interested

in a specific aspect of the command, e.g., “What can I do with $\langle \text{command} \rangle$?” or “What does -c mean in $\langle \text{command} \rangle$?”

To account for different expressions, we craft a set of templates (Appendix A) to construct varied prompts as

$$\langle q_i, r_i \rangle \longrightarrow \langle \mathcal{Q}(q_i, t), r_i \rangle, t \sim \mathbb{T}, \quad (2)$$

where $\mathcal{Q}(\cdot, \cdot)$ converts the original prompt q_i based on the template t sampled from the template set \mathbb{T} . Different expressions of the same prompt correspond to the same response r_i . Note that the set of templates can be easily extended using sentence rephrasing techniques in existing works [30].

To consider other aspects of a shell command that users may be interested in, we construct an extended set of potential prompts by querying general-purpose commercial LLMs as follows.

Question Augmentation Prompt

Ask as many different questions as possible about the following command from all perspectives, and respond in the format of one question per line.

Command: $\langle \text{command} \rangle$

The extended prompt set also enables us to construct $\langle \text{multi-round prompt}, \text{response} \rangle$ pairs that allows RACONTEUR to engage in a multi-round dialogue with users as follows.

An Example of Multi-Round Dialogue

```

USER: What does command bash -c '0<&137
    -;exec ... <&137 >&137 2>&137' do?
RACONTEUR: The command bash can ...
USER: What is the meaning of -c?
RACONTEUR: In bash, -c is used to ...
... ..

```

In this way, the behavior explainer can answer questions on various aspects of the command in a multi-round interactive manner. This allows analysts to continue asking questions if there is anything unclear about the explanation.

2) *Response Professionalization*: To guarantee factuality of the explanation, we construct responses by referring to professional explanations in code libraries.

Malicious shell command libraries. We resort to three libraries that contain different kinds of malicious Unix Shell or PowerShell commands.

- `atomic-red-team`. Atomic Red Team³ is a library of tests mapped to the MITRE ATT&CK framework. Each test includes the test command, the test name, the test description, and the attack technique it belongs to.
- `metta`. Metta⁴ is another library of tests in the MITRE ATT&CK framework. Each test includes the command, the name, the description, the attack technique and the tactic it belongs to.
- `reverse-shell`. We generate this dataset via the Metasploit framework⁵ as reverse shells require special attention in cybersecurity attacks.

Benign shell command libraries. We adopt two libraries of benign shell commands.

- `NL2Bash`. We utilize NL2Bash [29], which is originally designed for translating natural language descriptions into Bash commands.
- `The Stack`. We utilize the Unix Shell and PowerShell subsets of The Stack [24], which is a large code library. Note that, the shell scripts might contain comments, providing extra information about the commands.

The meta-information provided by the command libraries and the documentation related to each command are combined to generate accurate responses. The prompt template we use is as follows.

Knowledge-Infused Prompt

Please refer to the command documentations and command descriptions, answer the following questions:
<prompt>

Command documentation:
<docs>

Command description:
 MITRE ATT&CK Technique: <meta:type>
 cmd name: <meta:name>
 cmd description: <meta:desc>

B. Intent Identifier

In this section, we aim to identify the technique and tactic of the shell command in a standard way. First, we leverage the behavior explanation and search for the technique with the most similar description. Note that a tactic consists of a series of techniques and a technique may be used to achieve multiple tactics. Then, we predict the tactic of the command based on the identified technique.

³<https://github.com/redcanaryco/atomic-red-team>

⁴<https://github.com/uber-common/metta>

⁵<https://www.metasploit.com/>

As shown in Figure 3, to match the behavior descriptions given by the behavior explainer and the standard technique descriptions in the MITRE ATT&CK framework, we build a translator based on the Text2Vec model. Text2Vec is a natural language processing technique used to acquire vector representations of texts. Text2Vec models can map two texts with different syntax but similar semantics to vectors that are in close proximity. Conversely, texts with different semantics will be mapped to vectors that are distantly positioned from each other. Our intuition is to map the behavior description and the technique description into the same embedding space and then compute the distance between the embedding vectors. However, there may be a gap between the generated behavior description and the standard technique description that cannot be well bridged by a vanilla Text2Vec model. To tackle this problem, we fine-tune a Text2Vec model on a dataset of $\langle \textit{behavior}, \textit{technique description}, \textit{label} \rangle$ triples to obtain an enhanced behavior-and-description-to-vector (BD2Vec) model. $\textit{label} \in \{0, 1\}$ denotes whether or not the *behavior* matches the *technique description*. More specifically, for each command sourced from `atomic-red-team` and `metta`, we acquire its behavior also using response professionalization as mentioned in §IV-A2, wherein we substitute the `<prompt>` in the knowledge-infused prompt with a behavior description prompt. The behavior description prompt is tailored to facilitate the generation of descriptions aligned with the style of the MITRE ATT&CK framework through in-context learning (ICL) [36], as depicted below.

Behavior Description Prompt

Very briefly describe what adversaries attempt to do by the following shell command:
 Command: <command>

Example (follow the style of the following DESCRIPTION):
 Command: <example-cmd>
 DESCRIPTION: <example-std-desc>

The BD2Vec model encodes the behavior description and the standard description into vectors. In this way, the technique t_e of the shell command is identified by

$$t_e = \arg \max_{i \in \mathbb{I}} \mathcal{S}(\mathcal{V}(d), \mathcal{V}(s_i)), \quad (3)$$

where $\mathcal{V}(\cdot)$ denotes the BD2Vec model, $\mathcal{S}(\cdot, \cdot)$ is the function to compute the similarity score, and \mathbb{I} is the set of all techniques, e.g., 196 techniques in the MITRE ATT&CK framework. d is the description of the shell command and s_i is the standard description of technique i .

To identify the tactic t_a of a shell command, we construct a vector set \mathbb{J}_j for each tactic $j \in \mathbb{J}$, consisting of all technique description vectors belonging to that tactic. Then the tactic is identified by

$$t_a = \arg \max_{j \in \mathbb{J}} \text{avg} \left(\text{Top}_k [\mathcal{S}(\mathcal{V}(d), \mathcal{V}(s_i))] \right), \quad (4)$$

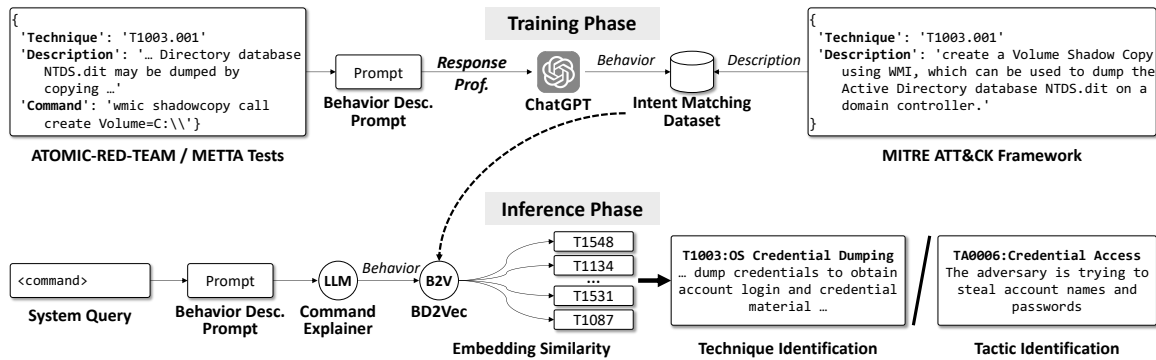


Fig. 3: Intent identification design of RACONTEUR. In the training phase, a Text2Vec model is fine-tuned to acquire an enhanced BD2Vec model, which maps the behavior description from LLMs and the standard description into the same embedding space. In the inference phase, the behavior description is encoded into a vector for comparison with a set of standard technique descriptions.

where $\text{Top}_k(\cdot)$ is a function that picks the k largest elements. $\text{avg}(\cdot)$ is the average function.

C. Doc-Augmented Enhancer

The user prompt alone provides limited hint for RACONTEUR. To enhance the quality of explanation, we augment the original user prompt with relevant information retrieved from command documentations as shown in Figure 1. The user prompt and the retrieved texts are combined as a new prompt to be fed into RACONTEUR.

The key of command documentation retrieval is to search for most relevant information to the user prompt. To achieve this goal, we match the command in the user prompt with contents in available documentations. Similar to intent identification, we fine-tune a Text2Vec model into a command-and-documentation-to-vector (CD2Vec) model with a carefully-constructed dataset of $\langle \text{command}, \text{doc}, \text{label} \rangle$ triples, where *label* denotes whether the command and the documentation are related.

We utilize the Linux manual pages of Bash commands to curate a dataset as follows.

1) *Extraction of utility and option descriptions:* Given the consistent structure of Linux manual pages across utilities, we utilized regular expression matching to extract the descriptions of utilities and their respective options. This method enabled us to establish the ground-truth documentation for each utility and its associated options.

2) *Generation of shell commands:* We randomly create shell commands by combining utilities and their options, which are not necessary to be real-world commands, e.g., `ls --ignore-backups` and `ls --color -t`. As we possess the knowledge regarding the utilities and options associated with each shell command, we are able to ascertain the ground-truth documentation for each command. (from Step 1).

3) *Documentation chunking:* We partition the raw documentations into chunks by certain rules, i.e., word count or line feeds, simulating documentation snippets.

4) *Creation of triples:* We obtain $\langle \text{command}, \text{doc}, \text{label} \rangle$ triples by verifying whether the documentation snippet is in the ground-truth set (from Step 2).

The doc-augmented enhancer enables RACONTEUR to explain private shell commands that are not in public datasets by retrieving relevant information from private documentations. Note that a company may define new private shell commands with the same name as but different meanings from public shell commands in the training set of LLMs. In this case, we expect RACONTEUR to explain the private shell commands mainly according to the private documentation but not on the memorized knowledge extracted from public shell commands. To realize this purpose, in the training phase, after documentation retrieval, we substitute the objects (e.g., utilities) in both the command and the retrieved contexts with random strings as if they are unseen by RACONTEUR. In this way, RACONTEUR can be trained to create explanations relying on the documentations faithfully.

V. EXPERIMENT SETUP

A. Prototype

We have implemented a fully-functional prototype of RACONTEUR.

- *Behavior explainer* (①). The behavior explainer is obtained by fine-tuning ChatGLM2-6B, an open-source bilingual (Chinese-English) chat model of 6 billion parameters [63], [12]. We conduct the fine-tuning process using four NVIDIA A100 (80GB) GPUs for four days with a batch size of 16 and a learning rate of $1e-4$ for 42,000 steps. We set the maximum length of queries and responses as 1,024 tokens. The total number of tokens for fine-tuning is 232 millions.
- *Intent identifier* (②). The key component of the intent identifier, i.e., the BD2Vec that aligns the embedding space of the generated behavior explanation and the standard description of the MITRE ATT&CK framework, is materialized with five Text2Vec models, i.e., Sentence-T5_{large} [41], GTR-T5_{XL} [42], SGPT [38], E5_{large} [54], and E5_{large}(FT).

- *Doc-augmented enhancer* (③). The key component of the doc-augmented enhancer, *i.e.*, the CD2Vec model that relates the command and relevant information in documentations, is obtained by fine-tuning the second version of E5_{large} model. E5_{large} is the state-of-the-art text embedding model of 330 million parameters [54]. The embedding size is 1,024. We fine-tune the model using low-rank adaptation (LoRA) [18], a popular parameter-efficient fine-tuning method.

B. Datasets

We have curated three datasets for constructing the behavior explainer, intent identifier, and doc-augmented enhancer, respectively.

1) *Dataset for command explainer*: As mentioned in §IV-A2, we construct the dataset for the command explainer based on five data sources. The dataset contains a total of 254,000 samples. We split the dataset into 9:0.5:0.5 for training, validation, and testing respectively. The test set is built to fulfil the following properties.

- *Diversity*. The test set includes both malicious and benign shell commands. More than 300 types of malicious Unix Shell and PowerShell commands are included.
- *Bilingual*. The test set includes both Chinese and English queries and responses.
- *Single-round & multi-round*. The test set consists of both single-round and multi-round samples for command explanation. The multi-round samples enhance the user-system interaction.
- *Inspected by human experts*. Four computer science researchers have been recruited to inspect and refine 200 of the test samples to be served as evaluation references.

More details of the test set is presented in Table IX. Note that all commands in our test set are real-world commands. Furthermore, there is no overlap between commands in our training set and test set. Therefore, the commands in the test set can be regarded as new and unseen.

2) *Dataset for intent identifier*: We utilize the `atomic-red-team` dataset for training and validating the BC2Vec model of the intent identifier, while the `metta` dataset is used for testing. In total, `atomic-red-team` and `metta` contain malicious shell commands of 129 techniques (293 sub-techniques) and 14 tactics, all of which have been labelled with ground-truth techniques and tactics by security experts. We have manually verified the labels of the test set and discovered that some of the labels were out-of-date due to the continuous maintenance and updates of the MITRE ATT&CK framework. Consequently, we have replaced these labels with the correct ones from the most recent version of the MITRE ATT&CK framework. Note that the original test set exhibits non-uniformity. A naive model that predicts the densest techniques can achieve a Top-1 ACC of 10.7% and a Top-5 ACC of 41.4%. To mitigate the impact of data distribution imbalance, we create a balanced version of the test set through resampling.

3) *Dataset for doc-augmented enhancer*: As mentioned in §IV-C, we construct a documentation retrieval dataset with the Linux manual pages of 1,662 Bash utilities, curating around 952,000 triples. We split the dataset into 9:0.5:0.5 for training, validating, and testing the CD2Vec model of the doc-augmented enhancer.

C. Evaluation Methods

To comprehensively assess the performance of RACONTEUR, we conduct both quantitative and qualitative evaluations.

1) *Quantitative evaluations*: We use the following quantifiable metrics to evaluate the performance of RACONTEUR.

- *Quantifiable metric for behavior explainer*. We adopt four well-established quantifiable metrics, *i.e.*, ROUGE [28], BLEU [46], METEOR [2] and CIDEr [51] to evaluate the performance of the behavior explainer, which provide token-level interpretable evaluations of natural language outputs. METEOR and CIDEr are two advanced metrics that have demonstrated a high correlation with human judgments [2], [51].
 - ROUGE (Recall-Oriented Understudy for Gisting Evaluation) [28] is an NLP metric that compares a machine-produced text against the ground-truth text. A higher ROUGE- k means a larger k -grams overlapping between the generated text and the ground-truth text. k is usually set as 1 or 2. ROUGE- ℓ measures the longest common sub-sequence (LCS) of the generated text and the ground-truth text.
 - BLEU (Bilingual Evaluation Understudy) [46] is an NLP metric that is precision-oriented. A higher BLEU- k indicates a larger k -grams overlapping between the generated text and the ground-truth text. k is usually set to 4.
 - METEOR (Metric for Evaluation of Translation with Explicit ORdering) [2] is an advanced NLP metric that takes into account word order, stemming matching and synonymy matching, aspects which are not considered by BLEU.
 - CIDEr (Consensus-Based Image Description Evaluation) [51] is an advanced NLP metric originally designed to compute the similarity between a machine-generated image description and the ground-truth description.

We also conduct end-to-end evaluation on the performance of the behavior explainer by using the generated command explanations for command classification, simulating the scenario when security analysts are presented with the command explanations. In this way, accuracy, precision, and recall are utilized as metrics.

- *Quantifiable metric for intent identifier*. We use Top- k ACC to assess the accuracy of the intent identifier in detecting the techniques and tactics of malicious commands. Top- k ACC represents the proportion of cases where the ground-truth label is among the top k labels

TABLE I: The Overall Performance of the Command Explainer.

Model	Malicious Command [‡]						Benign Command [†]					
	ROUGE-1	ROUGE-2	ROUGE- ℓ	BLEU-4	METEOR	CIDEr	ROUGE-1	ROUGE-2	ROUGE- ℓ	BLEU-4	METEOR	CIDEr
GPT-3.5-Turbo	48.7	24.3	34.5	36.3	39.1	30.2	54.3	27.0	35.2	29.5	40.9	16.8
GPT-4	45.5	20.3	30.5	40.5	32.5	14.6	51.8	25.8	36.2	34.2	34.5	12.0
ChatGLM2-6B	42.5	18.0	26.9	35.2	31.5	10.4	50.5	24.3	32.1	30.9	32.5	9.3
RACONTEUR	68.9	51.5	58.8	59.5	51.1	128.5	69.3	46.1	53.1	48.5	50.5	43.0
Increased [§]	62.1%	186.1%	118.5%	69.0%	62.2%	1137.1%	37.2%	89.7%	65.4%	57.0%	55.7%	362.9%
cf. GPT-4*	151.4%	253.7%	192.8%	146.9%	157.1%	880.5%	133.8%	178.7%	146.7%	141.8%	146.4%	358.7%

[‡] Malicious Command consists of data from `atomic-red-team`, `metta`, and `reverse-shell`.

[†] Benign Command consists of data from `NL2Bash`.

[§] Increased: the percentage improvement of RACONTEUR over the original ChatGLM2-6B model.

* cf. GPT-4: the achieved percentage of GPT-4 performance.

TABLE II: The Overall Performance of the Command Explainer on HumanCheck Test Set.

Model	HumanCheck					
	ROUGE-1	ROUGE-2	ROUGE- ℓ	BLEU-4	METEOR	CIDEr
GPT-3.5-Turbo	62.1	37.5	45.9	48.8	43.3	30.3
GPT-4	54.4	29.1	36.9	38.8	31.1	6.8
ChatGLM2-6B	55.7	29.4	34.9	38.3	32.6	6.8
RACONTEUR	69.6	47.7	53.3	53.6	48.9	47.0
Increased [§]	25.0%	62.2%	52.7%	39.9%	49.8%	594.4%
cf. GPT-4*	127.9%	163.9%	144.4%	138.1%	157.0%	693.5%

[§] Increased: the percentage improvement of RACONTEUR over the original ChatGLM2-6B model.

* cf. GPT-4: the achieved percentage of GPT-4 performance.

TABLE III: End-to-End Evaluation of The Performance of the Command Explainer.

Model	Classification Metrics		
	Precision (%)	Recall (%)	Accuracy (%)
GPT-3.5-Turbo	78.7	62.6	72.8
GPT-4	76.7	59.8	70.8
ChatGLM2-6B	77.6	47.1	66.7
RACONTEUR	83.7	79.2	81.8

predicted (ranked by similarity scores). A higher Top- k ACC means that the matching method can better infer the attacker’s intent.

- *Quantifiable metric for doc-augmented enhancer.* AUC-ROC is used to evaluate the performance of the doc-augmented enhancer. A higher AUC-ROC (Area Under the ROC curve) score indicates that the CD2Vec model can effectively encode related command-documentation pairs to be in close proximity, while distinctly separating unrelated pairs.

2) *Qualitative evaluations:* We have recruited 52 undergraduate and graduate students majoring in computer science to assess the responses of RACONTEUR. The detailed evaluation process and results are provided in §VI-D.

D. Baselines

1) *Baselines for behavior explainer:* For behavior explanation, we compare RACONTEUR with GPT-3.5-Turbo [6],

[45], GPT-4 [44] and the original ChatGLM2-6B [12], [63]. GPT-3.5-Turbo and GPT-4 are two commercial state-of-the-art LLMs and ChatGLM2-6B is an open-source bilingual (Chinese-English) model.

2) *Baselines for intent identifier:* For intent identifier, we compare RACONTEUR with GPT-3.5-Turbo and GPT-4. In particular, we demonstrate a chain-of-thought [61] process in the prompt together with an example of a command and its corresponding technique. Then, we query GPT-3.5-Turbo and GPT-4 to identify a given command’s technique and tactic.

3) *Baselines for doc-augmented enhancer:* For doc-augmented enhancer, we compare RACONTEUR with four state-of-the-art embedding models, *i.e.*, Sentence-T5_{large} [41], GTR-T5_{XL} [42], SGPT [38], and E5_{large} [54], which are state-of-the-art Text2Vec models on the massive text embedding benchmark [39]. Detailed information of the baseline models is summarized in Table VII and VIII.

VI. EVALUATION RESULTS

A. Command Explanation

1) *Overall Performance:* As shown in Table I, we compare the explanation capability of RACONTEUR with baselines regarding both malicious and benign commands. We can see that RACONTEUR achieves the best explanation performance on both malicious and benign commands in terms of all six metrics, improving the vanilla ChatGLM2-6B by 37.2%~1137.1%. Note that, the three baselines perform worse on malicious commands than on benign commands in terms of ROUGE- k , which means some expected information related to the malicious commands is not mentioned in their responses. In contrast, RACONTEUR shows equal and even better explanation performance on malicious commands, *e.g.*, RACONTEUR outperforms the second-best by 41% and 32% in terms of ROUGE- ℓ and BLEU-4, respectively. This may result from the response professionalization that we design in §IV-A2 to incorporate knowledge of malicious commands into RACONTEUR. We also evaluate RACONTEUR and the three baselines on the HumanCheck test set, which is inspected and corrected by four computer science researchers. We can see in Table II that RACONTEUR also achieves the best performance on HumanCheck test set.

TABLE IV: The Performance on Three Tasks.

Model	Explanation & Explanation w/ Doc. [†]			
	ROUGE-1	ROUGE-2	ROUGE- ℓ	BLEU-4
GPT-3.5-Turbo	49.3	24.0	31.4	28.8
GPT-4	48.5	23.7	32.7	36.9
ChatGLM2-6B	46.3	21.9	28.9	31.6
RACONTEUR	68.4	48.1	54.6	52.9
Increased [§] cf. GPT-4*	47.7%	119.6%	88.9%	67.4%
	141.0%	203.0%	167.0%	143.3%

Model	Behavior Summarization			
	ROUGE-1	ROUGE-2	ROUGE- ℓ	BLEU-4
GPT-3.5-Turbo	51.1	26.0	39.5	43.8
GPT-4	44.8	18.2	30.5	42.7
ChatGLM2-6B	41.6	15.7	26.8	38.5
RACONTEUR	70.3	53.9	62.1	63.8
Increased [§] cf. GPT-4*	69.0%	243.3%	131.7%	65.7%
	156.9%	296.2%	203.6%	149.4%

[†] Explanation & Explanation w/ Doc. include both the explanation task and the explanation with documentation task.


[§] Increased: the percentage improvement of RACONTEUR over the original ChatGLM2-6B model.

* cf. GPT-4: the achieved percentage of GPT-4 performance.


2) *End-to-End Evaluation of Overall Performance:* In order to evaluate the insightfulfulness of the command explanations for aiding security analysts in making judgments, we provide the command explanations to GPT-3.5-Turbo for command classification. As depicted in Table III, RACONTEUR achieves the highest precision (83.7%), recall (79.2%) and accuracy (81.8%). It demonstrates that the explanations generated by RACONTEUR are more effective in facilitating command comprehension for identifying malicious commands.

3) *Different Tasks:* We compare RACONTEUR with the baselines on two tasks, *i.e.*, step-by-step explanation and behavior summarization. We present the results in Table IV. As shown in Table IV, RACONTEUR outperforms the three baselines on the step-by-step explanation and behavior summarization tasks by a large margin (*e.g.*, over 40% and 30% in terms of ROUGE- ℓ and BLEU-4, respectively).

4) *Different Expressions of Query:* Since users might use different ways to ask RACONTEUR to explain a command, in this part, we evaluate RACONTEUR’s robustness against diversified expressions of explanation queries. We present the evaluation results of three types of query expressions in Table X in Appendix D, including the original query, the diversified query and the diversified query with documentations. We create diversified queries by converting the original queries into different expressions using Equation (2). We can see that RACONTEUR performs consistently on all three query expressions and achieves the best performance comparing with the baselines. This suggests that RACONTEUR can generate comprehensive explanations without explicit or intentional prompting. Moreover, RACONTEUR demonstrates the best performance on diversified queries with documentation. This underscores the robustness of RACONTEUR to query expressions



(a) Different top- p for inference (temperature=0.8)



(b) Different temperature for inference (top- p =0.8)

Fig. 4: The impact of temperature and top- p on the performance of RACONTEUR.

and its proficiency in leveraging command documentations. We notice that the performance of the baselines degrades when queries are in different expressions. It is because existing general-purpose LLMs are not trained to account for the variance in queries. In contrast, RACONTEUR provides complete analysis when requested by different expressions of explanation queries.

5) *Different Languages:* We also evaluate the bilingual capability of RACONTEUR comparing with the three baselines. As shown in Table XI in Appendix E, we test four models on the same tasks both in Chinese and English. RACONTEUR achieves the best performance in both languages. Note that it seems that these four models have better performance in English than in Chinese, however, the metrics for Chinese and English are computed in different ways in terms of word segmentation (*a.k.a.*, *tokenization*) and thus are not comparable.

6) *Different Inference Hyper-Parameters:* In this part, we evaluate the impact of the inference hyper-parameters, including the temperature and the top- p parameters. We set temperature and top- p from 0.1 to 0.9, evaluate RACONTEUR on the HumanCheck test set, and present the evaluation results in Figure 4. We find that these two parameters have little impact on the performance of RACONTEUR. When we vary the temperature, the performance of RACONTEUR only changes by 1%~3%, and reaches the best result when temperature=0.3. When we vary p , the performance of RACONTEUR only changes




Fig. 5: ROC curves of five Text2Vec models. Our CD2Vec model obtains the highest AUC, indicating the best discriminating ability.

by 2%~4%, and achieves the best result when $p=0.5$.

B. Intent Identification

We compare the intent identification capability of RACONTEUR with baselines GPT-3.5-Turbo and GPT-4. We present the evaluation results for the original test set in Table V and those for the balanced test set in Table VI. Overall, the fine-tuned version of E5_{large} achieves the best performance on both technique and tactic identification. All five of our materializations of the intent identifier show superior performance compared to GPT-3.5-Turbo and GPT-4. For example, SGPT, though performing the worst among our five materializations of the intent identifier, still surpasses GPT-3.5-Turbo and GPT-4 by 16.4% and 7.1% in terms of Top-1 ACC on technique and tactic identification, respectively. We attribute this to our behavior-oriented intent identification method. Note that, the Top-1 and Top-5 ACC of RACONTEUR (52.4% and 83.0%) significantly surpass those of the naive model that predicts the densest techniques mentioned in §V-B, *i.e.*, 10.7% and 41.4%. Additionally, we observe only minor differences between the results of Table V and Table VI. This demonstrates that the superiority of RACONTEUR is not contingent upon data distribution imbalance.

C. Doc-Augmented Enhancer

A Text2Vec model attributing high similarity scores to related command-documentation pairs and low scores to unrelated pairs is more proficient in retrieving the correct documentation snippets within a documentation database. From the dataset detailed in §IV-C, we randomly select 9,782 $\langle \text{command}, \text{doc}, \text{label} \rangle$ triples, both related and unrelated, as our test set. We compare our CD2Vec with four baselines. The ROC curves are presented in Figure 5, where our CD2Vec achieves the highest AUC of 0.981. This performance significantly surpasses that of the baselines, of which the average AUC is 0.858. Note that the baseline models each have a larger number of model parameters than ours. This result highlights that our lightweight doc-augmented enhancer is able to determine the relevance between a given shell command and a documentation snippet.

TABLE V: Technique and Tactic Identification Performance on the Original Test Set.

Model	Technique (ACC)			Tactic (ACC)
	Top-1	Top-5	Top-10	Top-1
GPT-3.5-Turbo	25.6	33.8	35.0	45.7
GPT-4	26.0	32.5	36.2	55.5
Sentence-T5 _{large}	45.4	80.1	87.3	69.4
GTR-T5 _{XL}	50.8	79.6	87.5	70.5
SGPT	42.4	75.0	84.3	62.6
E5 _{large}	48.5	78.2	87.1	69.3
E5 _{large} (FT)	52.4	83.0	90.4	75.0

TABLE VI: Technique and Tactic Identification Performance on the Balanced Test Set.

Model	Technique (ACC)			Tactic (ACC)
	Top-1	Top-5	Top-10	Top-1
GPT-3.5-Turbo	27.0	32.9	34.1	43.3
GPT-4	28.1	32.0	34.9	52.3
Sentence-T5 _{large}	51.2	78.8	86.6	66.7
GTR-T5 _{XL}	51.3	80.2	88.1	69.4
SGPT	48.5	78.4	86.8	67.3
E5 _{large}	50.5	81.3	87.8	69.0
E5 _{large} (FT)	56.0	82.1	89.3	74.3

D. User Study

We conduct a user study to qualitatively evaluate the performance of RACONTEUR. We have recruited 52 undergraduate and graduate students majoring in computer science to answer three sets of questions. The participants are asked to evaluate their proficiency in shell command, *i.e.*, at elementary, intermediate, and advanced levels⁶. The proportions of elementary, intermediate and advanced participants are 35%, 38% and 27%, respectively. In the first part, each question consists of one command and the corresponding explanation (from three baselines or RACONTEUR). Participants are requested to provide three scores, *i.e.*, to rate the explanation based on whether it helps them understand the details and intent of the command, as well as whether it helps them determine whether the command is malicious or not, using a scale of 1~5 points. In the second part, participants are asked to compare the reference answer and the response from LLMs, and rate the correctness of the response on a scale of 1~5 points. The first two parts consist of 40 questions, with 20 malicious commands and 20 benign commands, along with responses from RACONTEUR, GPT-4, GPT-3.5-Turbo, and ChatGLM2. In the third part, each question consists of one command and two responses from different methods. Participants are asked to choose the one they prefer and after that describe the reason why a specific response

⁶In general, advanced level indicates more than 7 years of Unix Shell or PowerShell experience, intermediate is 3~7 years, and elementary is 0~3 years.

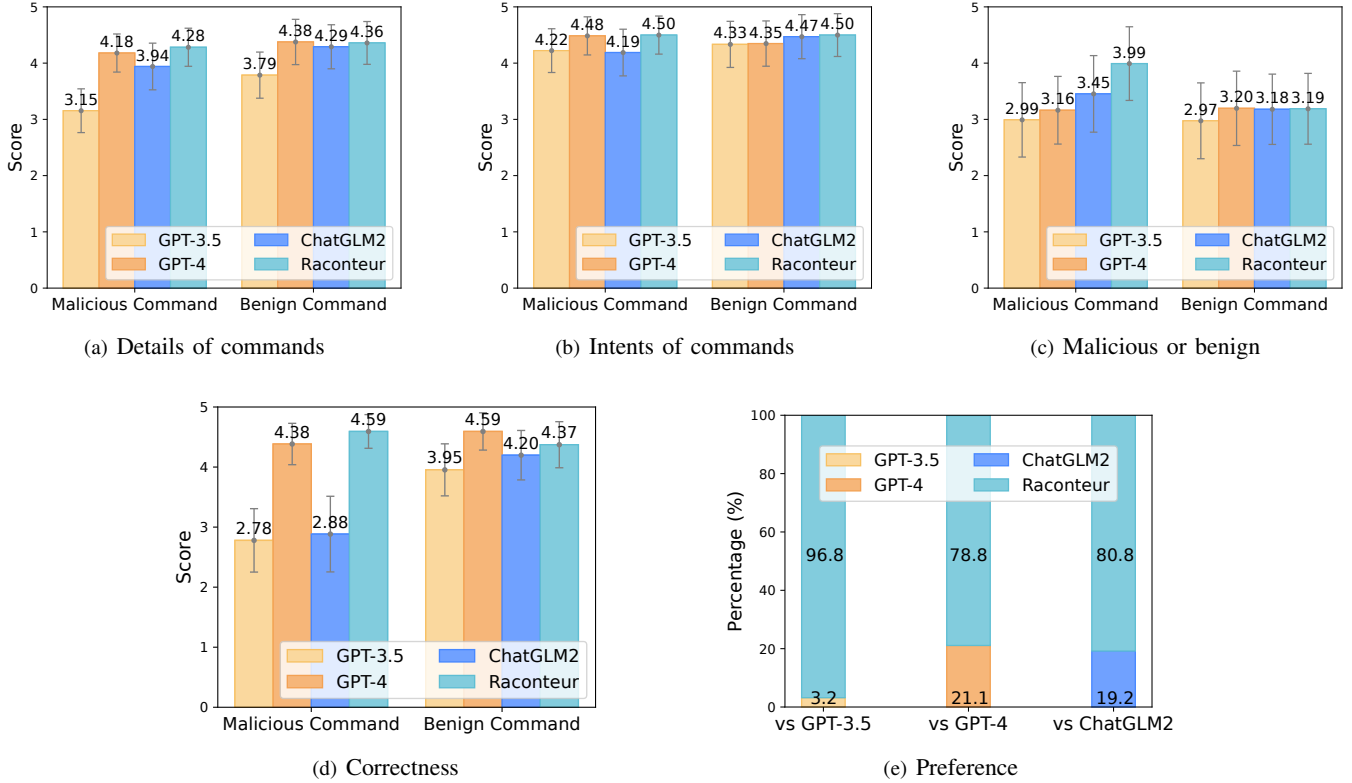


Fig. 6: The results of the user study. RACONTEUR achieves the best comprehensiveness and correctness for malicious commands and the second-best for benign commands (slightly lower than GPT-4) (a & d). RACONTEUR demonstrates superior insightfulness on both malicious and benign commands (b & c). Participants favor the explanations provided by RACONTEUR over those of baselines (e).

is considered good. This part includes 12 questions comparing RACONTEUR to GPT-4, GPT-3.5-Turbo, and ChatGLM2.

1) *Comprehensiveness & Correctness:* Comprehensiveness indicates that the explanation helps individuals understand the details of the command. As depicted in Figure 6(a), RACONTEUR achieves the highest score for malicious commands and the second-highest score for benign commands (slightly lower than GPT-4), which aligns with the results in §VI-A1. This demonstrates that RACONTEUR can generate comprehensive explanations, aiding the understanding of both malicious and benign commands. As shown in Figure 6(d), RACONTEUR achieves the best correctness on malicious commands and the second-best on benign commands. We notice that GPT-3.5-Turbo gets relatively unsatisfying results. We inspect the responses of GPT-3.5-Turbo and find that it tends to return brief and short explanations when not explicitly prompted.

2) *Insightfulness:* Insightfulness indicates that the explanation helps individuals understand the intent of the command, enabling them to determine whether the command is malicious or benign. As depicted in Figure 6(b) and Figure 6(c), RACONTEUR demonstrates superior performance on both malicious and benign commands. It is noteworthy that RACONTEUR outperforms baselines in assisting individuals to identify malicious commands.

3) *Preference:* Preference directly represents the likability of RACONTEUR compared with those of baselines. As shown

in Figure 6(e), participants favor the explanations provided by RACONTEUR over those of baselines. This validates the high quality of explanations generated by RACONTEUR.

Based on the free descriptions regarding the reasons for preference, we identify two factors that are deemed beneficial and valuable for command analysis, *i.e.*, detailed formatted explanations and explicit warnings. It shows that more detailed and formatted responses with a clear hierarchical structure are preferred. Furthermore, explicit warnings about the malicious nature of commands can enhance analysts' alertness.

VII. RELATED WORK

In this section, we briefly review the most related works on human-oriented code explanation and machine-oriented log analysis.

A. Human-Oriented Code Explanation

Code Explanation. The category of works within code explanation endeavors to generate comprehensive explanations that aid in the better understanding of code by humans. Over the recent years, the emergence of large-scale pre-trained models has catalyzed several initiatives in the field of code explanation. Notably, Sarse *et al.* [49] and MacNeil *et al.* [33], [34] explored and evaluated generating code explanations by large language models (LLMs) for education purposes. Denny *et al.* [10] introduced and evaluated robo-sourcing for educational resources,

i.e., querying LLMs to replace some of the work traditionally performed by the crowd. Leinonen *et al.* [25] compared the explanation generated by LLMs with those created by students, and found that LLMs outperformed students in terms of accuracy and understandability.

Existing works above [49], [33], [34], [10], [25] utilized cloud services of closed-source LLMs (*e.g.*, OpenAI’s Codex [7] and GPT-3 [6]) for normal code explanation. However, the use of cloud services is restricted when it comes to explaining proprietary commands, due to their sensitive nature and close association with company operations and internal network states. In contrast, RACONTEUR is founded upon a customized local LLM and incorporates support for handling malicious commands, thereby enhancing its capability to assist in shell log auditing. This unique approach not only ensures a more secure and tailored solution for code explanation, especially in scenarios involving proprietary commands, but also aligns with the imperative need for heightened security measures in contemporary software development practices.

Code Summarization. The objective of code summarization is to generate natural language comments for code, which aligns with the goal of code explanation to a certain extent. Early endeavors in code summarization were information retrieval-based. Haiduc *et al.* [17], Eddy *et al.* [13] and Wong *et al.* [62] proposed to search comments from similar code snippets to create summaries. Moreno *et al.* [37] extracted keywords in code snippets for Java class summaries. However, these information retrieval-based methods falter when no similar code snippet exists or when identifiers in the code are poorly named, a common occurrence in the context of malicious shell commands.

In recent years, learning-based methods have gained traction. Allamanis *et al.* [1] proposed a convolutional attention network for extreme (very short) code summarization. Iyer *et al.* [20] used long short-term memory (LSTM) networks with attention to produce summaries of C# and SQL code snippets. Hu *et al.* [19] and Wan *et al.* [53] incorporated an abstract syntax tree structure as well as sequential content of code snippets into an encoder-decoder model for better comment generation. Wei *et al.* [60] proposed a joint model to take advantage of the duality of code generation and code summarization. Given the success of transformer in NLP, Wang *et al.* [55] and Clement *et al.* [8] applied transformer in code-to-text (code summarization) and text-to-code (code generation) tasks.

Considering the promising effectiveness of graph neural networks (GNNs) on structured data, Fernandes *et al.* [15], Wang *et al.* [56] and Wang *et al.* [57] utilized GNN to reason about the long-distance relationships in code. Stepping into the era of pre-training language models (PLMs), a series of multi-tasking PLMs [32], [14], [59], [16], [58], [27] for code have been proposed for various code understanding and generation tasks, including code summarization.

While code summarization can be viewed as a form of brief explanation, detailed explanations are preferred in practice for better comprehension. Moreover, existing methods are developed for general code, whereas RACONTEUR supports

explaining proprietary commands based on documentation and analyzing malicious commands at a high-level through technique and tactic identification.

B. Machine-Oriented Log Analysis

In the realm of cybersecurity, there has been a significant focus on the automated analysis of logs for various downstream tasks, with an emphasis on providing explanations tailored for machine understanding rather than human consumption. Notably, Crespi *et al.* [9] applied unsupervised NLP methods on honeypot command logs to cluster IP addresses aiming at botnet detection. Boffa *et al.* [4] leveraged bag of words and Word2Vec to learn representations from honeypot logs and identify groups of similar SSH/Telnet sessions and attacks. These endeavors predominantly represent data mining efforts aimed at offering initial insights derived from log data.

Expanding on this, Boffa *et al.* [5] took a step further by delving into the identification of specific tactics employed by attackers within individual segments of shell sessions. While these efforts mark significant progress, they primarily yield trace-based or label-like outcomes. In contrast, our proposed system goes beyond mere traces or labels, offering comprehensive and easily comprehensible natural language explanations for commands. Moreover, RACONTEUR not only identifies the tactics employed by attackers but also discerns more granular techniques from individual commands, rather than analyzing entire sessions.

In a somewhat different vein, other research efforts [11], [3] have focused on learning representations from network data for predictive tasks such as delay prediction. These pursuits, while valuable, are orthogonal to our primary objective of explicating and understanding malicious activities within command logs. Our emphasis lies in providing clear, human-readable insights into the nature of potentially harmful commands, thereby enhancing the interpretability of cybersecurity log analysis.

VIII. CONCLUSION & FUTURE WORK

In this work, we present the design, implementation, and evaluation of RACONTEUR, a knowledgeable, insightful and portable shell command explainer powered by LLM, which provides bottom-up explanations of shell commands and assists shell log auditing in security operation. RACONTEUR is infused with professional knowledge to provide comprehensive explanations on shell commands, including not only what the command does but also why the command does it. Our extensive experiments demonstrate that RACONTEUR achieves much better explanation performance than the original base model, which is even comparable to the GPT series on benign and malicious shell commands in both English and Chinese. RACONTEUR also achieves superior technique and tactic identification performance. However, RACONTEUR can be further improved in four aspects:

Obfuscation. We have observed that when presented with obfuscated commands, RACONTEUR can notify the analyst within the explanation that the commands have indeed been obfuscated. Moreover, RACONTEUR can effectively explain

commands with payloads obfuscated using techniques such as *Base64* encoding. This success may be attributed to the emerging ability of existing LLMs to comprehend syntactic transformations [47], such as string splicing, common encoding, and simple encryption. However, due to the absence of a benchmark for obfuscated commands, we consider a more comprehensive evaluation of the analysis of obfuscated commands as a focus for our future work.

Shell session. A shell session consists of continuous commands executed by potential attackers, which contains more information about the attack. Although RACONTEUR has already shown satisfying performance on analyzing a standalone shell command and a compound command that consists of multiple commands, extending to shell sessions that consists of sequential shell commands might be more conducive, *e.g.*, for honey-pot analysis and APT analysis.

Other logs. RACONTEUR is now designed for assisting shell log auditing. A future direction is to extend the system to analyze other kinds of logs, *e.g.*, network logs, database logs, web server logs, *etc.* Extending RACONTEUR to fine-tuning the base model on other log modalities requires extensive labelled data which is expensive. A possible solution is through incremental pre-training on unsupervised data of network logs, database logs, web server logs, *etc.* This is our future direction.

Base model selection. In this paper, we choose ChatGLM2-6B as our base model and further fine-tune it into RACONTEUR for professional command explanation. We choose ChatGLM2-6B because it is the state-of-the-art within open-source bilingual (Chinese-English) LLMs. But a limitation is that ChatGLM2-6B only has 6 billion parameters, which may not be able to beat the performance of larger models, *e.g.*, GPT-3.5-Turbo and GPT-4. Although we are happy to see that RACONTEUR has already achieved satisfying performance compared with our baselines in §VI, exploration of larger open-source models as the base model may improve the performance further.

ACKNOWLEDGEMENT

We sincerely thank all the anonymous reviewers for their valuable comments. This work was supported by China NSFC Grant 61925109 and by Ant Group. The authors from Ant Group were supported by the Leading Innovative and Entrepreneur Team Introduction Program of HangZhou (Grant No. TD2020001). Yanjiao Chen is the corresponding author.

REFERENCES

- [1] M. Allamanis, H. Peng, and C. Sutton, "A Convolutional Attention Network for Extreme Summarization of Source Code," in *International Conference on Machine Learning*. JMLR.org, 2016.
- [2] S. Banerjee and A. Lavie, "METEOR: an automatic metric for MT evaluation with improved correlation with human judgments," in *Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*. Association for Computational Linguistics, 2005.
- [3] Z. Ben-Houidi, R. Azorin, M. Gallo, A. Finamore, and D. Rossi, "Towards a Systematic Multi-modal Representation Learning for Network Data," in *ACM Workshop on Hot Topics in Networks*, 2022.
- [4] M. Boffa, G. Milan, L. Vassio, I. Drago, M. Mellia, and Z. Ben-Houidi, "Towards NLP-based Processing of Honeypot Logs," in *IEEE European Symposium on Security and Privacy*, 2022.
- [5] M. Boffa, R. V. Valentim, L. Vassio, D. Giordano, I. Drago, M. Mellia, and Z. B. Houidi, "LogPr\`ecis: Unleashing Language Models for Automated Shell Log Analysis," *arXiv preprint arXiv:2307.08309*, 2023.
- [6] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language Models are Few-shot Learners," in *Conference on Neural Information Processing Systems*. PMLR, 2020.
- [7] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. de Oliveira Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman *et al.*, "Evaluating Large Language Models Trained on Code," *arXiv preprint arXiv:2107.03374*, 2021.
- [8] C. B. Clement, D. Drain, J. Timcheck, A. Svyatkovskiy, and N. Sundaresan, "PyMT5: Multi-mode Translation of Natural Language and Python Code with Transformers," in *Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2020.
- [9] V. Crespi, W. Hardaker, S. Abu-El-Haija, and A. Galstyan, "Identifying Botnet IP Address Clusters Using Natural Language Processing Techniques on Honeypot Command Logs," *arXiv preprint arXiv:2104.10232*, 2021.
- [10] P. Denny, S. Sarsa, A. Hellas, and J. Leinonen, "Robosourcing Educational Resources - Leveraging Large Language Models for Learner-sourcing," in *Workshop on Learnersourcing: Student-Generated Content*. CEUR-WS.org, 2022.
- [11] A. Dietmüller, S. Ray, R. Jacob, and L. Vanbever, "A New Hope for Network Model Generalization," in *ACM Workshop on Hot Topics in Networks*, 2022.
- [12] Z. Du, Y. Qian, X. Liu, M. Ding, J. Qiu, Z. Yang, and J. Tang, "GLM: General Language Model Pretraining with Autoregressive Blank Infilling," in *Annual Meeting of the Association for Computational Linguistics*, 2022.
- [13] B. P. Eddy, J. A. Robinson, N. A. Kraft, and J. C. Carver, "Evaluating Source Code Summarization Techniques: Replication and Expansion," in *International Conference on Program Comprehension*. IEEE Computer Society, 2013.
- [14] Z. Feng, D. Guo, D. Tang, N. Duan, X. Feng, M. Gong, L. Shou, B. Qin, T. Liu, D. Jiang, and M. Zhou, "CodeBERT: A Pre-Trained Model for Programming and Natural Languages," in *Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2020.
- [15] P. Fernandes, M. Allamanis, and M. Brockschmidt, "Structured Neural Summarization," in *International Conference on Learning Representations*. OpenReview.net, 2019.
- [16] D. Guo, S. Ren, S. Lu, Z. Feng, D. Tang, S. Liu, L. Zhou, N. Duan, A. Svyatkovskiy, S. Fu, M. Tufano, S. K. Deng, C. B. Clement, D. Drain, N. Sundaresan, J. Yin, D. Jiang, and M. Zhou, "GraphCodeBERT: Pre-training Code Representations with Data Flow," in *International Conference on Learning Representations*. OpenReview.net, 2021.
- [17] S. Haiduc, J. Aponte, L. Moreno, and A. Marcus, "On the Use of Automated Text Summarization Techniques for Summarizing Source Code," in *Conference on Reverse Engineering*. IEEE Computer Society, 2010.
- [18] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "Lora: Low-rank adaptation of large language models," in *International Conference on Learning Representations*. OpenReview.net, 2022.
- [19] X. Hu, G. Li, X. Xia, D. Lo, and Z. Jin, "Deep Code Comment Generation," in *Conference on Program Comprehension*. ACM, 2018.
- [20] S. Iyer, I. Konstas, A. Cheung, and L. Zettlemoyer, "Summarizing Source Code using a Neural Attention Model," in *Annual Meeting of the Association for Computational Linguistics*, 2016.
- [21] L. A. Johnson Kinyua, "AI/ML in Security Orchestration, Automation and Response: Future Research Directions," *Intelligent Automation & Soft Computing*, vol. 28, no. 2, pp. 527–545, 2021.
- [22] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, "Scaling Laws for Neural Language Models," *arXiv preprint arXiv:2001.08361*, 2020.
- [23] M. Khalili, M. Zhang, D. Borbor, L. Wang, N. Scarabeo, and M.-A. Zamor, "Monitoring and Improving Managed Security Services inside a Security Operation Center," *EAI Endorsed Transactions on Security and Safety*, vol. 5, no. 18, 2019.
- [24] D. Kocetkov, R. Li, L. B. Allal, J. Li, C. Mou, C. M. Ferrandis, Y. Jernite, M. Mitchell, S. Hughes, T. Wolf, D. Bahdanau, L. von Werra, and H. de Vries, "The Stack: 3 TB of Permissively Licensed Source Code," *arXiv preprint arXiv:2211.15533*, 2022.

- [25] J. Leinonen, P. Denny, S. MacNeil, S. Sarsa, S. Bernstein, J. Kim, A. Tran, and A. Hellas, "Comparing Code Explanations Created by Students and Large Language Models," in *Conference on Innovation and Technology in Computer Science Education V*. ACM, 2023.
- [26] B. Li, G. Fang, Y. Yang, Q. Wang, W. Ye, W. Zhao, and S. Zhang, "Evaluating ChatGPT's Information Extraction Capabilities: An Assessment of Performance, Explainability, Calibration, and Faithfulness," *arXiv preprint arXiv:2304.11633*, 2023.
- [27] R. Li, L. B. Allal, Y. Zi, N. Muennighoff, D. Kocetkov, C. Mou, M. Marone, C. Akiki, J. Li, J. Chim *et al.*, "StarCoder: May the Source be with You!" *arXiv preprint arXiv:2305.06161*, 2023.
- [28] C.-Y. Lin, "Rouge: A Package for Automatic Evaluation of Summaries," in *Text summarization branches out*. Association for Computational Linguistics, 2004.
- [29] X. V. Lin, C. Wang, L. Zettlemoyer, and M. D. Ernst, "NL2Bash: A Corpus and Semantic Parser for Natural Language Interface to the Linux Operating System," in *International Conference on Language Resources and Evaluation*. European Language Resources Association, 2018.
- [30] C. Liu, F. Zhao, L. Qing, Y. Kang, C. Sun, K. Kuang, and F. Wu, "Goal-oriented prompt attack and safety evaluation for llms," *arXiv preprint arXiv:2309.11830*, 2023.
- [31] X. Liu, H. Lai, H. Yu, Y. Xu, A. Zeng, Z. Du, P. Zhang, Y. Dong, and J. Tang, "WebGLM: Towards An Efficient Web-Enhanced Question Answering System with Human Preferences," in *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2023.
- [32] S. Lu, D. Guo, S. Ren, J. Huang, A. Svyatkovskiy, A. Blanco, C. B. Clement, D. Drain, D. Jiang, D. Tang, G. Li, L. Zhou, L. Shou, L. Zhou, M. Tufano, M. Gong, M. Zhou, N. Duan, N. Sundaresan, S. K. Deng, S. Fu, and S. Liu, "CodeXGLUE: A Machine Learning Benchmark Dataset for Code Understanding and Generation," in *Neural Information Processing Systems Track on Datasets and Benchmarks*. PMLR, 2021.
- [33] S. MacNeil, A. Tran, A. Hellas, J. Kim, S. Sarsa, P. Denny, S. Bernstein, and J. Leinonen, "Experiences from Using Code Explanations Generated by Large Language Models in a Web Software Development E-Book," in *Technical Symposium on Computer Science Education*, 2023.
- [34] S. MacNeil, A. Tran, D. Mogil, S. Bernstein, E. Ross, and Z. Huang, "Generating Diverse Code Explanations using the GPT-3 Large Language Model," in *Conference on International Computing Education Research*. ACM, 2022.
- [35] J. Maynez, S. Narayan, B. Bohnet, and R. T. McDonald, "On Faithfulness and Factuality in Abstractive Summarization," in *Annual Meeting of the Association for Computational Linguistics*, 2020.
- [36] S. Min, X. Lyu, A. Holtzman, M. Artetxe, M. Lewis, H. Hajishirzi, and L. Zettlemoyer, "Rethinking the Role of Demonstrations: What Makes In-Context Learning Work?" in *Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2022.
- [37] L. Moreno, J. Aponte, G. Sridhara, A. Marcus, L. L. Pollock, and K. Vijayshanker, "Automatic Generation of Natural Language Summaries for Java Classes," in *International Conference on Program Comprehension*. IEEE Computer Society, 2013.
- [38] N. Muennighoff, "SGPT: GPT Sentence Embeddings for Semantic Search," *arXiv preprint arXiv:2202.08904*, 2022.
- [39] N. Muennighoff, N. Tazi, L. Magne, and N. Reimers, "MTEB: Massive Text Embedding Benchmark," in *Conference of the European Chapter of the Association for Computational Linguistics*, 2023.
- [40] R. Nakano, J. Hilton, S. Balaji, J. Wu, L. Ouyang, C. Kim, C. Hesse, S. Jain, V. Kosaraju, W. Saunders, X. Jiang, K. Cobbe, T. Eloundou, G. Krueger, K. Button, M. Knight, B. Chess, and J. Schulman, "WebGPT: Browser-assisted Question-answering with Human Feedback," *arXiv preprint arXiv:2112.09332*, 2021.
- [41] J. Ni, G. H. Ábrego, N. Constant, J. Ma, K. B. Hall, D. Cer, and Y. Yang, "Sentence-T5: Scalable Sentence Encoders from Pre-trained Text-to-text Models," *arXiv preprint arXiv:2108.08877*, 2021.
- [42] J. Ni, C. Qu, J. Lu, Z. Dai, G. H. Ábrego, J. Ma, V. Y. Zhao, Y. Luan, K. B. Hall, and M.-W. Chang, "Large Dual Encoders are Generalizable Retrievers," *arXiv preprint arXiv:2112.07899*, 2021.
- [43] OpenAI, "Chatgpt plugins," <https://openai.com/blog/chatgpt-plugins>.
- [44] —, "GPT-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.
- [45] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askell, P. Welinder, P. F. Christiano, J. Leike, and R. Lowe, "Training Language Models to Follow Instructions with Human Feedback," in *Conference on Neural Information Processing Systems*. PMLR, 2022.
- [46] K. Papineni, S. Roukos, T. Ward, and W. Zhu, "BLEU: a Method for Automatic Evaluation of Machine Translation," in *Annual Meeting of the Association for Computational Linguistics*, 2002.
- [47] A. Rao, S. Vashista, A. Naik, S. Aditya, and M. Choudhury, "Tricking llms into disobedience: Formalizing, analyzing, and detecting jailbreaks," *arXiv preprint arXiv: 2305.14965*, 2024.
- [48] S. Ray, "Samsung Bans ChatGPT Among Employees After Sensitive Code Leak," <https://www.forbes.com/sites/siladityaray/2023/05/02/samsung-bans-chatgpt-and-other-chatbots-for-employees-after-sensitive-code-leak>.
- [49] S. Sarsa, P. Denny, A. Hellas, and J. Leinonen, "Automatic Generation of Programming Exercises and Code Explanations Using Large Language Models," in *Conference on International Computing Education Research*. ACM, 2022.
- [50] M. Schonlau and M. Theus, "Detecting Masquerades in Intrusion Detection based on Unpopular Commands," *Information Processing Letters*, vol. 76, no. 1-2, pp. 33-38, 2000.
- [51] R. Vedantam, C. L. Zitnick, and D. Parikh, "Cider: Consensus-based image description evaluation," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [52] M. Vielberth, F. Böhm, I. Fichtinger, and G. Pernul, "Security Operations Center: A Systematic Study and Open Challenges," *IEEE Access*, vol. 8, pp. 227 756-227 779, 2020.
- [53] Y. Wan, Z. Zhao, M. Yang, G. Xu, H. Ying, J. Wu, and P. S. Yu, "Improving Automatic Source Code Summarization via Deep Reinforcement Learning," in *ACM/IEEE International Conference on Automated Software Engineering*, 2018.
- [54] L. Wang, N. Yang, X. Huang, B. Jiao, L. Yang, D. Jiang, R. Majumder, and F. Wei, "Text Embeddings by Weakly-Supervised Contrastive Pre-training," *arXiv preprint arXiv:2212.03533*, 2022.
- [55] W. Wang, Y. Zhang, Z. Zeng, and G. Xu, "Trans³: A Transformer-based Framework for Unifying Code Summarization and Code Search," *arXiv preprint arXiv:2003.03238*, 2020.
- [56] Y. Wang, E. Shi, L. Du, X. Yang, Y. Hu, S. Han, H. Zhang, and D. Zhang, "Cocosum: Contextual Code Summarization with Multi-relational Graph Neural Network," *arXiv preprint arXiv:2107.01933*, 2021.
- [57] Y. Wang, Y. Dong, X. Lu, and A. Zhou, "GypSum: Learning Hybrid Representations for Code Summarization," in *IEEE/ACM International Conference on Program Comprehension*, 2022.
- [58] Y. Wang, H. Le, A. D. Gotmare, N. D. Bui, J. Li, and S. C. Hoi, "CodeT5+: Open Code Large Language Models for Code Understanding and Generation," *arXiv preprint arXiv:2305.07922*, 2023.
- [59] Y. Wang, W. Wang, S. R. Joty, and S. C. H. Hoi, "CodeT5: Identifier-aware Unified Pre-trained Encoder-Decoder Models for Code Understanding and Generation," in *Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2021.
- [60] B. Wei, G. Li, X. Xia, Z. Fu, and Z. Jin, "Code Generation as a Dual Task of Code Summarization," in *Conference on Neural Information Processing Systems*. PMLR, 2019.
- [61] J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. H. Chi, Q. V. Le, and D. Zhou, "Chain-of-thought prompting elicits reasoning in large language models," in *Conference on Neural Information Processing Systems*. PMLR, 2022.
- [62] E. Wong, T. Liu, and L. Tan, "CloCom: Mining Existing Source Code for Automatic Comment Generation," in *International Conference on Software Analysis, Evolution, and Reengineering*. IEEE Computer Society, 2015.
- [63] A. Zeng, X. Liu, Z. Du, Z. Wang, H. Lai, M. Ding, Z. Yang, Y. Xu, W. Zheng, X. Xia, W. L. Tam, Z. Ma, Y. Xue, J. Zhai, W. Chen, Z. Liu, P. Zhang, Y. Dong, and J. Tang, "GLM-130B: An Open Bilingual Pre-trained Model," in *International Conference on Learning Representations*. OpenReview.net, 2023.

APPENDIX A
EXPRESSION DIVERSIFICATION

Template Set of Different Expressions

- Can you clarify <command>
- Please describe <command>
- Elaborate on <command>
- Can you give me more details about <command>
- Could you shed some light on <command>
- I would like to understand <command>
- Can you break down <command>
- Can you make it clear <command>
- Can you give a rundown <command>
- Please provide a detailed explanation <command>
- I would like a detailed explanation <command>
- Kindly provide a thorough explanation <command>
- Can you give a detailed explanation <command>
- Please explain in detail <command>
- Can you explain it in detail <command>
- Could you provide a comprehensive explanation <command>
- Could you go into a detail about the command <command>

APPENDIX B
BASELINES

TABLE VII: The Baselines for Behavior Explainer and Intent Identifier.

Models	Type	Accessibility	#Parameters
GPT-3.5-Turbo	Generative	Closed-Source	175B
GPT-4	Generative	Closed-Source	100T
ChatGLM2-6B	Generative	Open-Source	6B

TABLE VIII: The Baselines for Doc-Augmented Enhancer.

Models	Type	Accessibility	#Param.	Dimension [†]
Sentence-T5 _{large}	Text2Vec	Open-Source	335M	768
GTR-T5 _{XL}	Text2Vec	Open-Source	1,240M	768
SGPT	Text2Vec	Open-Source	1,300M	2,048
E5 _{large}	Text2Vec	Open-Source	330M	1,024

[†] Embedding Dimension: the models encode arbitrary-length text into a fixed-dimensional vector, *e.g.*, 1,024.

APPENDIX C
TEST SET

TABLE IX: Test Sets of the Command Explainer.

Type	Data Source	Language	Task [‡]	Prompt [†]	#Samples	Alias
Malicious Command-Line	reverse-shell	Chinese	Explanation	Original Q.	785	rs-zh-ex-or
		Chinese	Explanation	Augmented Q.	743	rs-zh-ex-au
		Chinese	Explanation	Augmented Q. w/ Doc	773	rs-zh-ex-ad
		English	Explanation	Original Q.	954	rs-en-ex-or
		English	Explanation	Augmented Q.	933	rs-en-ex-au
		English	Explanation	Augmented Q. w/ Doc	919	rs-en-ex-ad
	metta	Chinese	Explanation	Original Q.	5	me-zh-ex-or
		Chinese	Explanation	Augmented Q.	5	me-zh-ex-au
		Chinese	Explanation	Augmented Q. w/ Doc	8	me-zh-ex-ad
		English	Explanation	Original Q.	9	me-en-ex-or
		English	Explanation	Augmented Q.	7	me-en-ex-au
		English	Explanation	Augmented Q. w/ Doc	10	me-en-ex-ad
	atomic-red-team	Chinese	Explanation	Original Q.	68	rt-zh-ex-or
			Explanation	Augmented Q.	58	rt-zh-ex-au
			Explanation	Augmented Q. w/ Doc	71	rt-zh-ex-ad
English		Explanation	Original Q.	66	rt-en-ex-or	
		Explanation	Augmented Q.	57	rt-en-ex-au	
		Explanation	Augmented Q. w/ Doc	69	rt-en-ex-ad	
English		Behavior	Original Q.	1464	rt-en-be-or	
		Behavior	Augmented Q.	1418	rt-en-be-au	
		Behavior	Augmented Q. w/ Doc	1918	rt-en-be-ad	
Benign Command-Line	NL2Bash	Chinese	Explanation	Original Q.	397	nb-zh-ex-or
		Chinese	Explanation	Augmented Q.	396	nb-zh-ex-au
		Chinese	Explanation	Augmented Q. w/ Doc	369	nb-zh-ex-ad
		Chinese	Explanation w/ Doc	Original Q.	410	nb-zh-ed-or
		Chinese	Explanation w/ Doc	Augmented Q.	409	nb-zh-ed-au
		Chinese	Explanation w/ Doc	Augmented Q. w/ Doc	411	nb-zh-ed-ad
		Chinese	Multi-round	Original Q.	5162	nb-zh-mu-or
		English	Explanation w/ Doc	Original Q.	386	nb-zh-ed-or
		English	Explanation w/ Doc	Augmented Q.	396	nb-zh-ed-au
	English	Explanation w/ Doc	Augmented Q. w/ Doc	407	nb-zh-ed-ad	
Benign & Malicious	atomic-red-team & NL2Bash	Chinese & English	Explanation & Explanation w/ Doc	Original	200	HumanCheck

[‡] Explanation: detailed explanation; Explanation w/ Doc: detailed explanation given documentations; Multi-round: multi-round interation; Behavior: behavior summarization.

[†] Original Q.: use the same prompts as when requesting ChatGPT; Augmented Q.: use augmented prompts; Augmented Q. w/ Doc: use augmented prompts with documentations.

APPENDIX D
DIFFERENT EXPRESSIONS OF QUERY

TABLE X: The Performance on Diversified Queries.

Model	Original Query				Diversified Query [†]				Diversified Query w/ Doc. [†]			
	ROUGE-1	ROUGE-2	ROUGE- ℓ	BLEU-4	ROUGE-1	ROUGE-2	ROUGE- ℓ	BLEU-4	ROUGE-1	ROUGE-2	ROUGE- ℓ	BLEU-4
GPT-3.5-Turbo	61.2	37.2	46.6	50.3	45.9	19.7	29.8	30.0	43.6	18.2	27.6	24.2
GPT-4	50.3	24.6	34.6	41.7	45.3	20.2	30.3	38.2	45.7	20.1	30.7	37.6
ChatGLM2-6B	49.3	23.6	31.1	38.0	41.4	17.3	26.5	31.8	42.9	17.8	26.9	32.9
RACONTEUR	70.4	52.0	58.7	57.9	68.5	49.3	56.6	56.2	68.5	49.6	57.0	56.8

[†] Diversified Query: the form of the query is converted using Equation (2).

APPENDIX E
DIFFERENT LANGUAGES

TABLE XI: The Performance on Chinese and English.

Model	Chinese			
	ROUGE-1	ROUGE-2	ROUGE- ℓ	BLEU-4
GPT-3.5-Turbo	46.6	20.6	29.2	19.6
GPT-4	46.5	20.9	31.9	26.4
ChatGLM2-6B	43.5	18.2	26.8	21.7
RACONTEUR	65.7	44.0	51.6	42.9

Model	English			
	ROUGE-1	ROUGE-2	ROUGE- ℓ	BLEU-4
GPT-3.5-Turbo	50.8	26.9	32.5	38.6
GPT-4	49.4	26.1	32.5	49.1
ChatGLM2-6B	48.1	25.5	30.4	42.8
RACONTEUR	70.9	53.2	58.4	65.1