

Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier 10.1109/ACCESS.2024.Doi Number

A Large Language Model-Supported Threat Modeling Framework for Transportation Cyber-Physical Systems

M Sabbir Salek¹, Member, IEEE, Mashrur Chowdhury¹, Senior Member, IEEE, Muhaimin Bin Munir², Student Member, IEEE, Yuchen Cai², Mohammad Imtiaz Hasan¹, Student Member, IEEE, Jean-Michel Tine¹, Student Member, IEEE, Latifur Khan², Fellow, IEEE, and Mizanur Rahman³, Senior Member, IEEE

¹Glenn Department of Civil Engineering, Clemson University, Clemson, SC 29634 USA

²Erik Jonsson School of Engineering and Computer Science, The University of Texas at Dallas, Richardson, TX 75080 USA

³Department of Civil, Construction, and Environmental Engineering, The University of Alabama, Tuscaloosa, AL 35487 USA

Corresponding author: M Sabbir Salek (e-mail: msalek@clemson.edu).

This work is based upon the work supported by the National Center for Transportation Cybersecurity and Resiliency (TraCR) (a U.S. Department of Transportation National University Transportation Center) headquartered at Clemson University, Clemson, South Carolina, USA. Any opinions, findings, conclusions, and recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of TraCR, and the U.S. Government assumes no liability for the contents or use thereof.

ABSTRACT Modern transportation systems are driven by cyber-physical systems (CPS), where cyber systems, such as computing infrastructure, and physical systems, such as transportation-related sensors and actuators, interact with each other seamlessly. These interactions help enhance transportation safety, mobility, energy efficiency, etc. However, increased reliance on automation and connectivity exposes transportation CPS to many cyber vulnerabilities. Existing threat modeling frameworks for transportation CPS are often narrow in scope, labor-intensive, and require substantial cybersecurity expertise. To address these challenges, we present the Transportation Cybersecurity and Resiliency Threat Modeling Framework (TraCR-TMF), a large language model (LLM)-based threat modeling framework for transportation CPS that requires limited cybersecurity expert intervention. TraCR-TMF identifies threats, potential attack techniques (i.e., methods to exploit vulnerabilities), and relevant countermeasures (e.g., attack detection and mitigation strategies) for transportation CPS. The open-source MITRE ATT&CK matrix is leveraged for these identifications using three LLM-based alternative approaches: (i) a retrieval-augmented generation (RAG)-based approach requiring no cybersecurity expert intervention, (ii) an in-context learning-based approach with low cybersecurity expert intervention, and (iii) a supervised fine-tuning approach with moderate cybersecurity expert intervention. Additionally, TraCR-TMF identifies potential attack paths, leading to the compromise of critical assets, by analyzing the vulnerabilities of different entities involved in transportation CPS using a customized LLM. Two cases were considered to evaluate the efficacy of TraCR-TMF. First, the framework was applied to identify relevant attack techniques for various transportation CPS applications. Results showed that 90% of the identified attack techniques were relevant, as validated by cybersecurity experts. Second, the framework, along with the LLM fine-tuned for the first evaluation, was used to identify potential attack paths leading to the compromise of a target asset in a real-world cyberattack incident involving industrial control systems. TraCR-TMF successfully predicted several exploitations, such as lateral movement of adversaries within information technology network, data exfiltration, and data encryption for ransomware, that occurred during a major real-world cyberattack incident. This validates the framework's potential transferability. These findings show TraCR-TMF's efficacy in transportation CPS' threat modeling, while lowering the barrier for cybersecurity expertise, and its potential for adaptation across different CPS.

INDEX TERMS Threat modeling, Cybersecurity, Large language model, Intelligent transportation systems, and Transportation cyber-physical systems

I. INTRODUCTION

With the advent of artificial intelligence (AI), automation, and the proliferation of wired and wireless communication technologies, legacy transportation systems are increasingly being supplemented by transportation cyber-physical systems (CPS). In a transportation CPS environment, different cyber systems, such as onboard, roadside and cloud-based computing infrastructure, interact with different physical systems, such as transportation-related sensors and actuators. While these systems promise enhanced efficiency, safety, and sustainability [1], they also introduce a substantially expanded attack surface that can be exploited by malicious actors [2], [3]. Cyberattacks on transportation CPS could pose serious safety risks, as compromised transportation systems can jeopardize human lives and disrupt critical daily operations.

Recent trends highlight a dramatic surge in cyberattacks targeting multimodal transportation systems. Between 2017 and 2022, cyberattacks on road transportation systems increased by 400% [4]. Similarly, the aviation sector experienced a 530% increase in cyberattacks between 2019 and 2020 [5], and maritime-related cyber incidents escalated by an alarming 900% between 2017 and 2021 [6]. These statistics underscore the growing cybersecurity risks faced by transportation CPS. Furthermore, transportation CPS could be interconnected with other critical infrastructures, including power grids, electronic banking systems, and healthcare networks. This interconnectivity increases the risk of attack propagation across multiple domains, which could lead to widespread disruptions orchestrated by individual threat actors or nation-state adversaries. As a result, securing transportation CPS against evolving cyber threats has become a critical concern for transportation stakeholders, security professionals, and policymakers.

Defending multimodal transportation CPS against a broad spectrum of potential cyber threats necessitates the deployment of both proactive and reactive cybersecurity measures. Proactive measures aim to prevent security breaches by identifying and mitigating vulnerabilities during the design and development phases, while reactive measures focus on detection, response, and recovery from cyber incidents. One of the most effective proactive strategies is threat modeling, which systematically identifies potential threats and vulnerabilities before they can be exploited [7], [8]. Threat modeling enables cybersecurity professionals to analyze security requirements, communicate risks effectively, and recommend appropriate resilience, detection, mitigation strategies. By incorporating threat modeling early in the transportation CPS development lifecycle, we could reduce security risks and enhance the overall cybersecurity posture of our transportation systems.

Despite its critical importance, the adoption of threat modeling in the transportation CPS domain remains limited. Although threat modeling has gained significant traction in the software development community, its application to

transportation CPS has been relatively underexplored. Existing studies on threat modeling for transportation CPS, which we highlighted in the related work section of this paper, are often constrained to specific modes of transportation or limited threat scenarios [9], [10], [11]. Moreover, traditional threat modeling frameworks require cybersecurity professionals to develop a comprehensive understanding of involved systems, their interactions, potential vulnerabilities, and corresponding threats—a process that is time-consuming and complex when performed manually. These challenges hinder the widespread adoption of threat modeling within transportation CPS, despite its well-documented benefits [12].

To address these challenges, this study presents a novel threat modeling framework to help enhance the cybersecurity and resilience of transportation CPS. In this study, we present the Transportation Cybersecurity and Resiliency Threat Modeling Framework (TraCR-TMF), which leverages existing cybersecurity tools, established threat models, and the logical reasoning and generative capabilities of large language models (LLMs). TraCR-TMF is a multi-stage threat modeling approach that is specifically tailored to assist cybersecurity professionals dealing with transportation CPS in conducting comprehensive threat assessments and identifying relevant countermeasures that could help detect and mitigate the identified threats.

The TraCR-TMF framework offers several advantages. First, it incorporates three alternative LLM-based approaches to identify relevant attack techniques, i.e., methods to exploit vulnerabilities. These LLM-based approaches can extract relevant details from structured knowledge bases, such as the MITRE Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK) matrix, with zero to moderate cybersecurity expert interventions. This capability enables the automatic identification of attack techniques and their corresponding detection and mitigation strategies from the MITRE ATT&CK matrix. Additionally, TraCR-TMF offers autonomous mapping of potential attack paths and associated attack techniques, retrieved from the MITRE ATT&CK matrix. These attack paths, identified by analyzing transportation CPS application vulnerabilities, could lead to the compromise of critical transportation CPS assets. By integrating these capabilities, TraCR-TMF streamlines the threat modeling process and enhances the accessibility of cybersecurity analysis for transportation CPS.

The rest of this paper is structured as follows: Section II presents the key contributions of this study. Section III presents a review of related methods and studies. Section IV provides briefs on tools and knowledge bases relevant to TraCR-TMF and its evaluation cases considered in this study. Section V details the technical aspects of TraCR-TMF, including the framework's design and implementation. Section VI presents two evaluation cases that assess the effectiveness of TraCR-TMF in modeling threats to transportation CPS. Finally, Section VII concludes the paper by summarizing the findings and the scope for future research.

II. CONTRIBUTIONS

This study aims to address the need for an easy-to-adapt, cost-effective threat modeling framework for transportation CPS that could assist in identifying vulnerabilities, potential threats, and their available countermeasures. TraCR-TMF offers LLM-supported identification of attack techniques and their corresponding detection and mitigation strategies from the widely utilized MITRE ATT&CK matrix, i.e., a continuously evolving knowledge base of known adversaries, requiring limited cybersecurity expert intervention. The key contributions of this framework are as follows:

- A multi-stage threat modeling framework for transportation CPS, incorporating widely used open-source tools and knowledge base, along with three LLM-supported adoption strategies.
- A retrieval-augmented generation (RAG)-based architecture to identify specific attack techniques that attackers could use to exploit vulnerabilities in a transportation CPS, along with the corresponding countermeasures, requiring no cybersecurity expert intervention.
- An in-context learning-based approach for attack technique and countermeasure identification, leveraging an engineered prompt and requiring minimal cybersecurity expert intervention.
- A supervised fine-tuning approach to facilitate the identification of relevant attack techniques and corresponding countermeasures, requiring moderate cybersecurity expert intervention.
- An LLM-based, asset-centric threat modeling approach that leverages an engineered prompt to identify potential attack paths and associated attack techniques. This asset-centric threat modeling offers insights into the multi-layered vulnerabilities of transportation CPS, i.e., a chain of vulnerabilities across multiple CPS entities that can be exploited together to perform an attack that aims to compromise one or more critical assets.

TraCR-TMF enables (i) systematic identification of potential threats, and their categories based on a well-established threat model, known as STRIDE, (ii) identification of known attack techniques, and their existing detection and mitigation strategies from the MITRE ATT&CK matrix based on the STRIDE-identified threats, and (iii) identification of potential attack paths and associated attack techniques for each step of the attack path that could compromise a critical transportation CPS asset.

III. RELATED WORK

Threat analysis or modeling is a well-studied topic in cybersecurity; consequently, many studies focus on different methods of threat modeling. Systematic reviews of these studies have been presented in several survey papers, for example, [7], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], which categorized the existing threat modeling methods

and frameworks from different perspectives, such as reviewing them based on the underlying methods or the type of systems or applications to which threat modeling is applied. In this section, we first present a general overview of different existing threat modeling approaches. Second, we present a review of the existing threat modeling studies that particularly relate to transportation CPS and associated applications. Third, we present a review of notable studies that leveraged LLMs to automate threat modeling.

A. THREAT MODELING IN GENERAL

Threat modeling or analysis has been defined from different perspectives in literature. Uzunov and Fernandez [23] defined threat modeling as a structured analysis of potential threats or attacks in various contexts, such as analysis of threats or attacks specific to systems and/or technologies. Another study [24] described threat modeling as a technique for identifying and documenting security threats in a software system while systematically uncovering the system's strengths and weaknesses. In [25], Xiong et al. referred to threat modeling as a process for assessing potential threats, risks, and attacks. Although threat and risk are used interchangeably in some studies, there lies a distinction between them. Al-Fedaghi and Alkandari [26] defined risk as a function of vulnerability and threat. Thus, analyzing threats is essential in risk assessment. In contrast, threat modeling is a structured process for identifying, analyzing, and mitigating security threats in a system, application, or network. It helps security professionals understand potential attack vectors, assess risks, and implement appropriate defenses.

The key aspects of threat modeling can include [20], [27]: (i) identifying assets, which implies determining what needs protection, for example, data, hardware, and software; (ii) determining threats, which refers to recognizing potential adversaries and attack vectors; (iii) analyzing potential attack paths, which implies exploring how an attacker could exploit the existing system vulnerabilities. (iv) assessing risks, which involves evaluation of the likelihood and impact of the threats; (v) identifying and implementing mitigation strategies: this refers to the identification and implementation of appropriate security controls based on different analyzed threats.

The existing threat modeling approaches can be broadly categorized as (i) formula-based threat modeling and (ii) model-based threat modeling, as presented in [22]. Formula-based methods are used for threat analysis and risk assessment of a system, primarily utilizing tables, textual descriptions, and mathematical formulas. In contrast, model-based methods are a type of threat analysis approach that employs various models to assess system threats and risks, utilizing data flow diagrams, graphs, and tree models for modeling and analysis.

The formula-based methods can be further classified into (i) asset-centric, (ii) vulnerability-centric, and (iii) attacker-centric threat modeling approaches [22]. An asset-centric approach, which can be thought of as a top-down approach, identifies target assets first and then maps potential attack

paths using expert knowledge, which enables early threat mitigations. A popular asset-centric method, OCTAVE, was introduced by Alberts et al. [28] for analyzing threats in a large organization with a multi-layered hierarchy or geographically distributed systems. This method involves three phases, as follows: (i) building asset-centric threat profiles, (ii) identifying infrastructure vulnerabilities, and (iii) developing security strategies and plans. In contrast to an asset-centric approach, a vulnerability-centric threat modeling approach is a bottom-up approach that initiates with a vulnerability of a system and analyzes other potential vulnerabilities or failures that could be caused by that vulnerability. For instance, the Common Vulnerability Scoring System (CVSS) assesses the severity of different system-level vulnerabilities so that appropriate countermeasures can be prioritized accordingly. Finally, an attacker-centric threat modeling approach focuses on analyzing potential attackers by assessing their knowledge, attack paths, motivations, and available resources. This approach enables threat modeling and risk assessment by identifying the root cause of attacks.

The model-based methods can also be further classified into (i) graph-based and (ii) attack tree-based threat model approaches [22]. Graph-based approaches utilize graphs with nodes and directional edges, representing quantitative relationships among the nodes. Among the most popular graph-based approaches is STRIDE, a Microsoft-developed threat modeling tool [29], which utilizes six categories of threats, i.e., spoofing, tampering, repudiation, information disclosure, denial-of-service, and elevation of privilege, to conduct threat modeling. LINDDUN, developed by Deng et al. [30], is another popular graph-based threat modeling framework that helps identify privacy-related threats. Another example of a graph-based threat modeling approach is PASTA, which was introduced by UcedaVelez and Morana [31]. PASTA is primarily focused on addressing high-risk threats to optimize investment and resource utilization.

In contrast to graph-based threat modeling, tree-based approaches utilize a tree to represent the hierarchical relationship among nodes. Attack tree is the most used tree-based threat modeling approach. This method represents threats as hierarchical trees, where each branch represents a potential attack path. Attack trees are especially useful in visualizing complex attack scenarios and understanding dependencies. For instance, Saini et al. [32] utilized the attack trees concept, originally developed by Bruce Schneier [33], to conduct threat modeling for a grid computing security subsystem called MyProxy. The MyProxy subsystem, within the Globus grid computing toolkit, serves as an online credential repository and a certificate authority for its users. Saini et al. [32] utilized the SecureITree tool to develop an Attack Tree for the MyProxy subsystem.

B. THREAT MODELING FOR TRANSPORTATION SYSTEMS

Although threat modeling is widely studied for software security in general, its applicability to multimodal transportation systems is still limited. In this section, we present an overview of notable existing studies related to transportation systems threat modeling.

Ramazanazadeh et al. [11] introduced an automated security assistant for transportation CPS' threat modeling based on their developed quantitative threat modeling algorithm called security object-oriented colored Petri nets (SOOCPN). SOOCPN classifies threats within a transportation CPS into different levels and associated color-coded sublevels based on the systems' security risk quantification. However, this approach relies on user-assigned arbitrary ranges of risks for defining the different security levels and their sublevels. Since such interpretations can vary from person to person, this approach poses a challenge for real-world implementations.

He et al. [9] developed a threat modeling approach for maritime transportation systems based on Markov chains to characterize domain name system (DNS) rebinding attack behaviors (i.e., malicious website tricking a web browser into thinking that a public domain name is associated with a private IP address) and extract key attack attributes. Simulations demonstrated the effectiveness of the approach in detecting and mitigating DNS rebinding attacks, enhancing the security of maritime internet-of-things (IoT) systems. However, this approach primarily focuses on DNS rebinding threats for maritime IoT networks, limiting its applicability to identifying a wide variety of threats for multimodal transportation CPS.

Kumar et al. [34] developed a deep learning-supported threat modeling framework for maritime IoT to automate manual threat analysis and aimed to address the low detection and high false alarm rates of existing threat modeling solutions. Their approach integrates a long short-term memory-based variational autoencoder (LSTM-VAE)-supported feature extraction, a bi-directional gated recurrent unit (Bi-GRU)-based threat detection, and another Bi-GRU-based attack type identification. However, this approach requires extensive training datasets to train the AI models and once trained, the models can only identify the type of threats they are trained on, limiting the framework's applicability to transportation CPS with a continuously evolving cyber threat landscape.

In another study, Subran et al. [35] applied the STRIDE model to analyze threats in electric vehicle charging infrastructure by decomposing the charging infrastructure architecture and generating threat scenarios using the Microsoft threat modeling tool. The authors of [35] developed an automated attack simulation platform, Mininet, to replicate electric vehicle communication and simulate denial-of-service attacks. However, the study primarily focuses on denial-of-service type threats, leaving other types of threats unaddressed.

Hamad et al. [36] presented a threat modeling framework for modern vehicles by integrating different subsystem-specific threat models into a generalized model. The authors relied on manual threat modeling by reviewing existing literature and used the threats as roots for generating attack trees. This study focuses on reusability, where attack trees can be applied across different vehicle systems. However, the authors acknowledged the lack of automation and recommended that future work should explore automated tools to generate and evaluate attack trees dynamically.

In summary, the existing threat modeling frameworks for transportation CPS are often limited to specific threat types and lack automation or wide adoption potential, as discussed above. This study aims to address this gap by developing an LLM-supported threat modeling framework for transportation CPS that offers alternative adoption strategies, reducing the responsible organizations' cybersecurity expertise-related constraints.

C. THREAT MODELING USING LLMs

With the advent of LLMs, several recent studies have utilized the advanced reasoning and text generation abilities of LLMs to automate threat modeling approaches. This section presents a review of notable related studies.

Munir et al. [37] presented a RAG-enhanced, LLM-based cyberattack detection approach for transit security. The authors utilized RAG to retrieve texts from the MITRE ATT&CK matrix based on relevant information associated with a transit security transportation CPS, specifically focusing on the descriptions of physical objects involved in the transportation CPS and the information flows among them. However, their evaluation was limited to a standard transit security application and did not include comparisons with other LLM-based strategies, such as in-context learning and supervised fine-tuning.

Gabrys et al. [38] presented a method to utilize a LLM to convert intrusion detection systems (IDS) rules into a format that can be interpreted by humans and trained several AI models to map the observed attacker behaviors to corresponding MITRE ATT&CK techniques. This study [38] utilized ChatGPT and BERT models for converting the Wazuh IDS predefined rules set. However, the requirement of a manually labeled training dataset is also indispensable for this study.

Branescu et al. [39] experimented with several BERT-based models trained on large cyber security corpus and compared their performance against ChatGPT while mapping from the MITRE Common Vulnerabilities and Exposures (CVE) to the MITRE ATT&CK techniques. The authors claimed their strategy can systematically analyze CVE descriptions and map them to corresponding ATT&CK tactics. This mapping helps understand the potential impact of vulnerabilities and assists in prioritizing mitigation efforts accordingly. The authors of [39] acknowledged that further exploration including more

data or metadata is needed to improve the performance of their AI models.

Another study, Nir et al. [40], explored how LLMs can be utilized to enhance network IDS by automating threat labeling. Nir et al. [40] developed a threat modeling strategy leveraging machine learning, deep learning, and behavioral analysis to identify anomalies and classify attacks by analyzing network traffic data. However, challenges like scalability, high false positives, and domain-specific adaptations remain. The authors recommended further improvements of their AI models incorporating transfer learning, hybrid AI models, and adversarial defense mechanisms.

Based on these reviews, this study is motivated to develop a threat modeling framework for transportation CPS that can be adopted without being constrained by extensive cybersecurity expertise and substantial manual labor requirements. To this end, our TraCR-TMF offers different alternative adoption strategies leveraging different levels of cybersecurity expert intervention and logical reasoning capabilities of LLMs, which we present in Section V.

IV. PRELIMINARIES

In this section, we briefly discuss a few preliminaries that are essential for presenting the TraCR-TMF and its evaluation cases considered in this study. The following subsections provides the readers with quick introductions to (i) the Microsoft Security Development Lifecycle (MS SDL) threat modeling tool, which identifies different types of threats in a transportation CPS at the first stage of the TraCR-TMF (details are presented in Section V-A), (ii) the MITRE ATT&CK matrix, which serves as a comprehensive cyberattack knowledge base for the second stage of the TraCR-TMF (details are presented in Section V-B), and (iii) the national intelligent transportation systems reference architecture, which provides reference architecture, including physical and functional objects, data flow, etc. for various transportation CPS applications.

A. MS SDL THREAT MODELING TOOL

Developed by Microsoft in 2018, the MS SDL threat modeling tool [29] serves as a core component of MS SDL, an engineered approach primarily to help web-based application and/or software developers identify cyber vulnerabilities, threats, potential cyberattacks, and their countermeasures. This open-source threat modeling tool allows developers to identify and mitigate potential cybersecurity issues early in their software or application development lifecycle.

MS SDL threat modeling tool utilizes the STRIDE model that categorizes threats into six categories, as follows: (i) spoofing, (ii) tampering, (iii) repudiation, (iv) information disclosure, (v) denial-of-service, and (vi) elevation of privilege [41]. Table I provides a brief description of each of these threat categories. This graph-based threat modeling tool provides users with a graphical user interface (GUI) for

developing data flow diagrams (DFDs) using different stencils, such as process, external interactor, data store, data flow, and trust boundary [42]. Once the DFD is provided for an application, the MS SDL threat modeling tool generates a threat report, showcasing the threats associated with each interaction observed in the given DFD based on the STRIDE model. The tool utilizes a “STRIDE per element” approach, which refers to the identification of STRIDE threats for each DFD element. A DFD element can be a process, an external interactor, a data store, or a data flow. For each threat identified by the MS SDL threat modeling tool, the report presents the associated STRIDE threat category (as presented in Table I), a brief description of the threat, along with its potential mitigations.

B. MITRE ATT&CK MATRIX

MITRE ATT&CK is a globally recognized cybersecurity knowledge base that provides a structured approach to understanding and analyzing cyber threats [43]. Developed by the MITRE Corp., ATT&CK provides comprehensive documentation of real-world adversary behaviors to help organizations improve their detection, response, and

mitigation strategies. This framework is widely used by cybersecurity professionals for threat intelligence, security operations, incident response, and adversary emulation.

MITRE ATT&CK is organized into matrices that cover different environments, including enterprise, mobile, and industrial control systems. Each matrix outlines various adversary behaviors that attackers use to infiltrate, persist, and execute malicious activities within a network. Each matrix includes some common key components, such as tactics, techniques, sub-techniques, and procedures. Table II presents a brief description of each component, along with some examples. For each technique, the MITRE ATT&CK provides a list of associated sub-techniques and procedure examples as well as the existing detection and mitigation strategies for those attacks.

C. NATIONAL REFERENCE ARCHITECTURE FOR INTELLIGENT TRANSPORTATION SYSTEMS

Transportation CPS plays a pivotal role in modern urban infrastructure by facilitating efficient traffic management, enhancing safety, and improving mobility. These systems rely on a sophisticated network of interconnected physical and

TABLE I
STRIDE THREAT CATEGORIES (ADOPTED FROM [41])

Threat Category	Threat Description
Spoofing	The unauthorized access and use of another user's authentication credentials, such as a username and password.
Tampering	The malicious alteration of data, including unauthorized modifications to persistent data (e.g., in a database) or changes to data in transit over an open network like the Internet.
Repudiation	A situation where a user denies performing an action without any means to verify the claim. For example, in a system without proper logging, an unauthorized operation can be executed without accountability.
Information disclosure	The unauthorized exposure of information to individuals who should not have access. Examples include users reading files they lack permission for or attackers intercepting data transmitted between systems.
Denial of service	The disruption of access to a system or service, rendering it temporarily unavailable or unusable to legitimate users; for example, making a web server inaccessible.
Elevation of privilege	An unprivileged user gaining higher-level access, potentially compromising or taking control of the entire system. This includes scenarios where an attacker bypasses all security defenses and integrates into the trusted system, posing a significant threat.

TABLE II
Key Components of MITRE ATT&CK (adopted from [43])

Component	Description	Examples
Tactic	Tactics represent the high-level objectives or goals that an adversary seeks to achieve during an attack. A tactic provides a context on why an attacker is performing a particular action.	<ul style="list-style-type: none"> Initial Access – Gaining entry into a system or network (e.g., phishing, exploiting vulnerabilities). Persistence – Establishing long-term access within a compromised environment.
Technique	Techniques describe how an adversary achieves a specific tactic. Each technique represents a method or approach used by attackers to execute their objectives.	<ul style="list-style-type: none"> Phishing (T1566) – Sending deceptive emails to trick users into providing credentials. Process Injection (T1055) – Injecting malicious code into legitimate processes to evade detection.
Sub-technique	Sub-techniques break down techniques into more specific methods, offering deeper granularity into how adversaries operate. They help organizations refine their detection and response strategies by addressing precise attack behaviors.	Sub-techniques of phishing (T1566): <ul style="list-style-type: none"> Spearphishing Attachment (T1566.001) – Delivering malware via an email attachment. Spearphishing Link (T1566.002) – Tricking users into clicking a malicious link.
Procedure	Procedures define the specific ways in which threat actors or malware implement techniques and sub-techniques. They provide real-world examples of how known adversary groups execute attacks.	Procedures of phishing (T1566): <ul style="list-style-type: none"> Axiom (G0001) – Axiom has used spear phishing to initially compromise victims. Hikit (S0009) – Hikit has been spread through spear phishing

functional entities that exchange data to help provide efficient and safe transportation services. To analyze the cybersecurity challenges within transportation CPS, this study utilizes the U.S. Department of Transportation's (USDOT) Architecture Reference for Cooperative and Intelligent Transportation (ARC-IT) [44]. ARC-IT offers a comprehensive framework for transportation CPS design and implementation, encompassing detailed service packages and associated components. Here, we present a brief introduction to the different components of these reference architectures:

1) SERVICE PACKAGES

ARC-IT categorizes transportation CPS functionalities into 156 service packages (to date), distributed across 12 distinct domains such as Public Safety, Data Management, and Vehicle Safety. Each service package on ARC-IT is equivalent to a transportation CPS application that integrates physical systems and functional components in a transportation CPS environment to address specific transportation services. This list of service packages is updated over time to include newly introduced transportation CPS.

2) PHYSICAL AND FUNCTIONAL OBJECTS

Physical objects represent the tangible entities involved in transportation CPS, including centers (e.g., traffic management center, emergency management center, payment administration center), fields (e.g., roadway equipment, intermodal terminal, electric charging station), personal devices (e.g., pedestrian, payment device, remote access device), support systems (e.g., data distribution system, archived data system), and vehicles (e.g., basic transit vehicle, basic commercial vehicle). These objects serve as the core infrastructure of transportation CPS, enabling critical information exchange for operational and security purposes. Each physical object is associated with one or more functional objects, which are deployment-specific units designed to fulfill the functional requirements of the system.

3) INFORMATION AND DATA FLOWS

Information and data flows serve as the backbone of transportation CPS communications, facilitating data

exchange among physical and functional objects. Each information flow originates from a physical object, referred to as the initiator or source, and is received by another physical object, referred to as the acceptor or destination. In contrast, data flow refers to the data exchange between two functional objects in a transportation CPS application. According to these definitions of information and data flows used in ARC-IT, each information flow (between two physical objects) can be broken down into several data flows (between two functional objects).

V. OVERVIEW OF TRACR-TMF

The TraCR-TMF, presented in Fig. 1, is a threat modeling framework that requires limited cybersecurity expert intervention and leverages different LLM-enabled approaches to (i) identify potential attack techniques that could exploit the vulnerabilities within transportation CPS, and (ii) analyze these attack techniques to identify potential attack paths leading to critical transportation CPS assets. As depicted in Fig. 1, this multi-stage threat modeling is carried out by augmenting a primary threat report generated from the MS SDL threat modeling tool with relevant attack techniques and their corresponding detection and mitigation strategies retrieved from the MITRE ATT&CK matrix. Once relevant MITRE ATT&CK techniques have been identified, a customized LLM is utilized to perform an asset-centric threat modeling in which the LLM identifies potential attack paths, along with associated attack techniques, leading to the compromise of specified assets. The different stages of the TraCR-TMF, presented in Fig. 1, are detailed in the following subsections.

A. STRIDE-BASED THREAT IDENTIFICATION (STAGE 1)

The TraCR-TMF utilizes the open-source MS SDL threat modeling tool [29] to model threats based on the STRIDE model. MS SDL threat modeling tool requires the systems or reference architecture as a DFD, which is modeled through the tool's GUI. As mentioned in Section IV-A, a DFD can include four types of elements, as follows: (i) Processes: functional objects in a transportation CPS application are modeled as

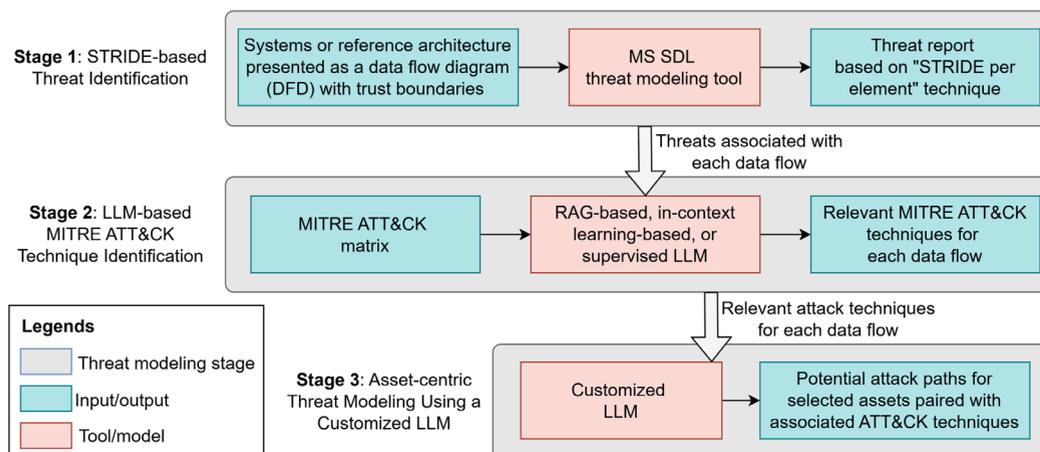


FIGURE 1. Overview of TraCR-TMF for transportation CPS.

processes; (ii) External interactors: terminators, i.e., entities that are involved (e.g., a human user) but fall outside the design scope of a transportation CPS application, are modeled as external interactors; (iii) data stores: in-house, external, or cloud-based data stores are modeled as data stores; and (iv) data flows: interactions between functional objects, or between a functional object and an external interactor, are modeled as data flows. In addition, a trust boundary is used to indicate a physical object involved in a transportation CPS application that may include one or more functional objects.

Once the DFD, along with the trust boundaries, of a transportation CPS application is provided, the MS SDL threat modeling tool generates a STRIDE-based threat report. This report outlines spoofing, tampering, repudiation, information disclosure, denial-of-service, and elevation of privilege-related threats for each interaction in a DFD. These threats are identified for any processes, external interactors, data stores, or data flows associated with each interaction in a DFD. However, this report does not provide information on specific attack techniques associated with each threat that attackers could utilize to exploit the vulnerabilities. Knowledge of these attack techniques, along with their existing detection and mitigation methods, would enable application developers to implement targeted countermeasures for effective prevention, detection, and mitigation of potential threats. In this study, we address this gap through the TraCR-TMF’s LLM-based specific MITRE ATT&CK technique identification approaches (presented in Section V-B) and further extend it to identify exploitable multi-layered vulnerabilities that could compromise critical transportation CPS assets (presented in Section V-C).

B. LLM-BASED MITRE ATT&CK TECHNIQUE IDENTIFICATION (STAGE 2)

The TraCR-TMF utilizes an LLM to identify attack techniques from the MITRE ATT&CK matrix that could be relevant in exploiting vulnerabilities within transportation CPS based on the threats identified by the STRIDE model. In this study, we consider three different approaches with LLMs to map the STRIDE-based threats to the potentially associated MITRE ATT&CK techniques that require different levels of cybersecurity expert intervention. These approaches are as follows: (i) a RAG-based approach with an LLM, (ii) an in-context learning-based approach with an LLM, and (iii) a supervised fine-tuning approach using an LLM. The input that is included in the prompts of all these approaches, which we refer to as the “basic input” in this study, includes the following components (as presented in Fig. 2):

- **Data Flow:** A name is specified for the data flow being analyzed.
- **Data Flow Definition:** An explanation of the data flow is provided
- **Initiator and Acceptor:** Descriptions of the physical objects responsible for initiating and receiving the

```
Data Flow: <data_flow_name>
Initiator: <initiator_name>
Acceptor: <acceptor_name>
Requires Authentication?: ...
Requires Encryption?: ...

Description of the Initiator: <initiator_description>
Function: <function_name_and_desc>
    <process1>: <description>
    <process2>: <description>
    ...

Description of the Acceptor: <acceptor_description>
Function: <function_name_and_desc>
    <process1>: <description>
    <process2>: <description>
    ...

Definition of <data_flow_name>: <data_flow_definition>

STRIDE-based threats associated with the data flow:
...
```

FIGURE 2. Basic input for LLM.

associated information flow, which includes the data flow under analysis, are included.

- **Functional Objects and Processes:** Descriptions of the functional objects that initiate or receive the data flow and the processes associated with these functional objects are provided, if known.
- **Security Attributes:** Required or recommended security attributes, such as confidentiality, integrity, availability, authentication, and encryption, are provided, if known.
- **STRIDE-based Threats:** The initially identified threats based on the STRIDE model are included from the MS SDL threat report.

In addition to the basic input, the different LLM-based approaches used in the TraCR-TMF utilize different sets of instructions, queries, examples, etc., in the prompts. The following sections explain these approaches in detail.

1) RAG-BASED APPROACH WITH AN LLM

RAG, first introduced by a group of META AI researchers in 2020 [45], became a widely popular technique in natural language processing for its ability to enhance text generation models. RAG provides a way of incorporating relevant external knowledge from a database or document repository in addition to relying on pre-trained knowledge. To this end, we consider an RAG-based approach to complement the generative power of LLMs with retrieval mechanisms to identify pertinent attack techniques from the MITRE ATT&CK matrix. Fig. 3 presents the architecture of our RAG-enabled LLM-based mapping to the MITRE ATT&CK techniques, comprising the following steps:

- **Initial Identification of Relevant Cyberattacks:** The process initiates with a user input, known as the vanilla prompt, which primarily includes a query to identify the cyberattacks that are potentially relevant to a given data flow. The query provided in the vanilla prompt is as follows: “*What are the possible cyberattacks that can be used to attack this information flow? Return them in*

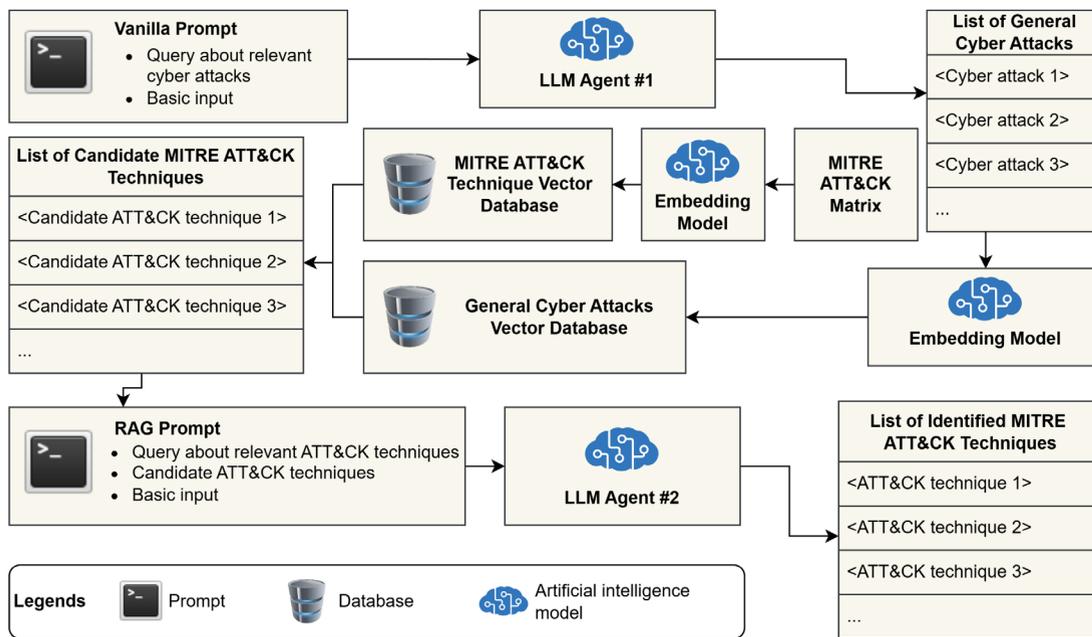


FIGURE 3. RAG-based LLM architecture for ATT&CK technique mapping.

a Python dictionary format with the key being the cyberattack technique and the value being the technique description.” In addition to this query, the prompt contains the basic input, which we presented in Fig. 2 before. The vanilla prompt is processed by an LLM, i.e., GPT-4o, which we denote as the LLM agent #1 in Fig. 3. GPT-4o is optimized for inference-based reasoning [46], making it well-suited for identifying implicit cyberattack techniques in our framework. In addition, the model’s faster response time, enhanced coherence, and context retention ability compared to its legacy models [47] motivated us to consider it for the TraCR-TMF. As shown in Fig. 3, the LLM agent #1 serves as an intelligent agent that generates a list of general cyberattack techniques based on the query and the basic input.

- Construction of a Vector Database from the MITRE ATT&CK Matrix:** MITRE ATT&CK descriptions are retrieved from files using a data connector, i.e., SimpleDirectoryReader [48]. The files are then processed in batches to optimize memory usage. Next, the text data is converted into embeddings utilizing one of OpenAI’s latest embedding models, i.e., text-embedding-3-small model [49]. This embedding model is chosen due to its cost-effectiveness and balance between speed and accuracy [50]. If an index already exists, new documents are inserted into it incrementally. The indices persist, enabling future use without requiring reprocessing. A VectorIndexRetriever [51] is set up to retrieve relevant embeddings based on similarity. Finally, a SimilarityPostprocessor [52] is applied with a cutoff of 0.6 to filter out low-relevance results [52]. A cutoff of 0.6 was selected as it offers a balanced trade-off between precision and recall. A

higher threshold excludes potentially useful but moderately similar entries, reducing recall, especially in cases where relevant descriptions use varied phrasing. Conversely, a lower threshold introduces semantically weak matches, reducing precision and increasing noise. The 0.6 cutoff showed to effectively filter out irrelevant content while retaining diverse yet meaningful results.

- Retrieval Candidate ATT&CK Techniques from Vector Databases:** The list of general cyberattacks produced by the LLM agent #1 is transformed into vector embeddings using OpenAI’s text-embedding-3-small model [49]. These embeddings are then matched against pre-vectorized descriptions of MITRE ATT&CK techniques using cosine similarity to identify the most relevant attack techniques. The top three similar MITRE ATT&CK techniques identified by the LLM agent #1 for each general cyberattack are then selected and compiled into a candidate list, as depicted in Fig. 3.
- Identification of ATT&CK Techniques using a RAG Prompt:** The candidate MITRE techniques are utilized to create a new prompt, which we refer to as the RAG prompt, as shown in Fig. 3. This prompt incorporates the candidate technique IDs, names, and descriptions retrieved from the MITRE ATT&CK matrix. It is followed by the query: “Which MITRE ATT&CK techniques from the table above can be used to attack the data flow? Provide the technique IDs in Python list format.” Additionally, the RAG prompt includes descriptions of the associated data flow, the initiator and the acceptor of the data flow to enrich the context. The RAG prompt is then processed by another LLM agent, which we refer to as the LLM agent #2 (also based on GPT-4o). LLM agent #2 refines the information and

generates a final list of specific MITRE ATT&CK techniques relevant to the data flow. Again, the GPT-4o model is chosen for this task due to its superior logical reasoning compared to OpenAI’s legacy models [46].

The RAG-based approach leverages the complementary strengths of generative and retrieval-based approaches, enabling contextually informed predictions of relevant MITRE ATT&CK techniques. Unlike the other two approaches we present in the following sections, this RAG-based approach identifies relevant MITRE ATT&CK techniques without any specific examples or retraining that would require a human cybersecurity expert’s intervention. Therefore, this approach is the most automated, requiring no cybersecurity expert intervention for relevant attack technique identification among the three approaches offered by the TraCR-TMF. In addition, this RAG-based approach enables the retrieval of relevant information from the most up-to-date list of attack techniques available in the MITRE ATT&CK matrix at the time of retrieval. Consequently, the information retrieved by this approach is not limited to the training cutoff period, as is the case with supervised fine-tuning approaches.

2) IN-CONTEXT LEARNING APPROACH WITH AN LLM

In-context learning (ICL) is a technique that leverages the capabilities of LLMs to perform tasks without requiring retraining. Instead, by providing examples and instructions directly within the prompt, ICL enables an LLM to generalize and produce more accurate results for new inputs. We consider the gpt-4o-mini-2024-07-18 model with ICL to identify the MITRE ATT&CK techniques that are potentially relevant to a given transportation CPS data flow. Three variations of ICL are considered in this context: (i) zero-shot, (ii) one-shot, and (iii) few-shot learning.

In zero-shot learning, the LLM is provided with only task instructions along with the basic input shown in Fig. 2, without any examples. This approach utilizes the model’s ability to generalize solely based on its pre-trained knowledge and the given context. Thus, zero-shot learning is equivalent to using the basic input without any examples.

In one-shot learning, the LLM is given a single example of data flow along with its corresponding MITRE ATT&CK techniques identified by human cybersecurity experts. This example provides the model with a reference for understanding the task and the expected output format.

In a few-shot learning, the LLM is supplied with multiple examples of data flows and their corresponding MITRE ATT&CK technique identified by human cybersecurity experts. By exposing the model to various instances, this approach improved the LLM’s comprehension of the task and enhanced its ability to generalize.

Fig. 4 illustrates the ICL prompt given to the LLM. The prompt is designed to ensure that the LLM receives adequate context to perform the requested task effectively. Since the ICL-based approach includes a few worked-out examples verified by cybersecurity experts, this approach is considered to require low cybersecurity expert intervention.

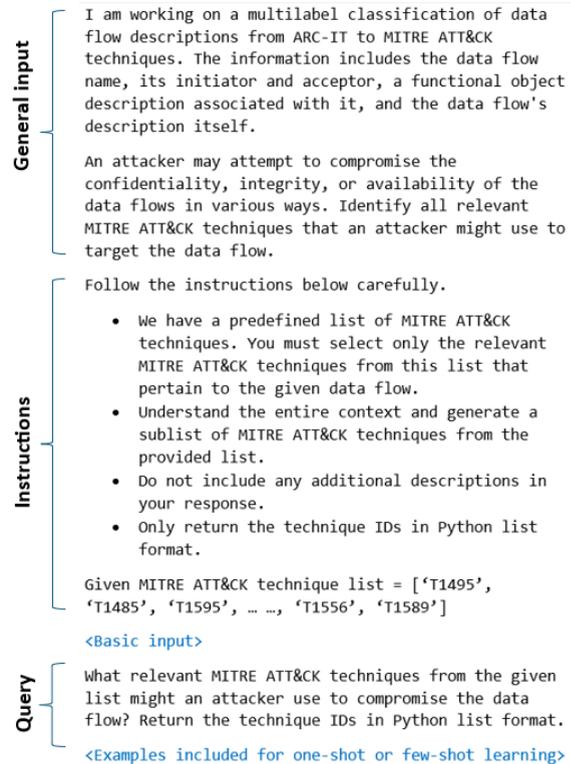


FIGURE 4. Prompt for LLM while using ICL.

3) SUPERVISED FINE-TUNING APPROACH WITH AN LLM

The supervised learning approach utilizes the power of pre-trained transformer-based models. Due to the complexity and lengthiness of our input texts, we selected ModernBERT over BERT. ModernBERT is a cutting-edge, encoder-only model that was trained on a wide range of English texts. Unlike the original BERT, it can handle longer input sequences [53], which is important in our case as the longest input sequence in our dataset requires nearly seven times BERT’s 512-token limit. The base version of ModernBERT, with its 22 layers and 149 million parameters, is chosen for this task since ModernBERT has been shown to strike a balance between strong performance and computational demands [53]. Unlike other state-of-the-art LLMs with superior reasoning capability, such as GPT-4o, the ModernBERT is open source and can be downloaded and run on a local machine. All these advantages of using ModernBERT make it an appropriate candidate for our attack technique identification task.

To perform supervised fine-tuning of the ModernBERT model, we follow a transfer learning approach, i.e., we start with a pre-trained model and fine-tune it for the task of our interest. The primary advantage of transfer learning in this case over training a model from scratch is that by leveraging the knowledge already embedded in the pre-trained model from its initial training on large online datasets, the model does not need to learn basic features from the ground up, which yields an improved performance especially when the target dataset is comparatively smaller.

In machine learning, supervised methods can be used when the training dataset contains ground-truth labels. In a multi-label classification problem, each instance can be associated with more than one label. Consider, we denote the dataset as follows: $D_{train} = \{(\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2), (\mathbf{x}_3, \mathbf{y}_3), \dots, (\mathbf{x}_N, \mathbf{y}_N)\}$, where $\mathbf{x}_i \in \mathbb{R}^N$ is the representation of the i^{th} instance, and $\mathbf{y}_i = [y_{i,1}, y_{i,2}, y_{i,3}, \dots, y_{i,K}]$ is a binary vector with $y_{i,j} = 1$ indicating the j^{th} label is associated with the i^{th} instance based on the ground-truth list, and $y_{i,j} = 0$ indicating otherwise. Then, our goal is to train a multi-label classifier based on ModernBERT that can perform the mapping from \mathbf{x}_i to \mathbf{y}_i . In the context of mapping a basic input associated with a transportation CPS data flow to a list of relevant MITRE ATT&CK techniques, \mathbf{x}_i represents the i^{th} basic input and \mathbf{y}_i represents the relevance of the MITRE ATT&CK techniques present in the ground-truth list. The binary cross-entropy loss function considered for retraining the ModernBERT model is given by,

$$L = \frac{1}{NK} \sum_{i=1}^N \sum_{j=1}^K [y_{i,j} \log(\hat{y}_{i,j}) + (1 - y_{i,j}) \log(1 - \hat{y}_{i,j})] \quad (1)$$

where $\hat{y}_{i,j}$ is the probability assigned by the classifier for the j^{th} MITRE ATT&CK technique to be associated with the i^{th} data flow. To enhance the model's robustness and generalizability, a five-fold cross-validation is performed. The model is trained for 20 epochs with a batch size of 8, a learning rate of 0.00002, and a weight decay of 0.01. The number of epochs was set to 20 as further training yielded diminishing F1-score in early tests, and a batch size of 8 was selected to ensure all input representations could be processed within the memory capacity of an A100 GPU. The learning rate and weight decay were empirically tuned to optimize model performance. The model with the highest F1 score is selected to perform the multi-label classification of the information flows to the relevant MITRE ATT&CK techniques.

C. ASSET-CENTRIC THREAT MODELING USING CUSTOMIZED LLM

Asset-centric threat modeling focuses on identifying threats associated with specified assets that could lead the assets to be compromised. These assets can be physical or functional objects, data stores, or data flows in a transportation CPS application. This type of modeling helps cybersecurity professionals realize potential attack paths and associated attack methods or tools that could be utilized by attackers to compromise an organization's valuable assets. Predicting these attack paths helps identify proactive and reactive security measures that need to be incorporated to minimize associated risks and impacts. To this end, the TraCR-TMF offers an asset-centric threat modeling approach to predict potential attack paths leading to compromising critical assets. In this context, an attack path can consist of multiple steps where an attack propagates from one entity to another within a transportation CPS network until it reaches the target asset

that the attacker wants to compromise. This is obtained by analyzing the potential MITRE ATT&CK techniques identified for different data flows in a transportation CPS application or infrastructure using a customized LLM. Each predicted attack path is broken down into individual attack steps that complete the attack path. Here, an attack step within a predicted attack path implies the attack propagating from one entity (e.g., an initiator or an acceptor of a data flow) to another entity. In addition, relevant attack techniques that could be used by attackers to conduct each step of the predicted attack paths are explored. The following sections detail how this asset-centric threat modeling is carried out by an LLM in the TraCR-TMF.

1) BUILDING A CUSTOMIZED GPT WITH CHATGPT

Powered by GPT-4o, OpenAI now offers users the opportunity to build their customized versions of ChatGPT with a single line of instruction [54]. This enables users to build their customized GPT with expertise in specialized tasks. In addition, enabling OpenAI's memory feature helps the customized GPT retain insights over sessions to refine interactions over time. To build our customized GPT for asset-centric threat modeling, we provide the following one-line instruction to ChatGPT: "Build a customized GPT for me that will serve as an expert-level cybersecurity analyst." Given this prompt, ChatGPT builds a customized GPT, specialized for cybersecurity analysis.

2) PROMPT CONSTRUCTION FOR ASSET-CENTRIC THREAT MODELING

Once the customized GPT is ready for user interactions, the GPT is provided with a prompt to perform the asset-centric threat modeling. This prompt includes the names of the initiator and the acceptor of each data flow, along with their associated MITRE ATT&CK techniques, in a table format, as shown in Fig. 5. This information is followed by a set of instructions and a query, asking the customized LLM to identify potential attack paths leading to a specified asset to be compromised along with the associated MITRE ATT&CK techniques that can be utilized in each step of the attack path.

As shown in Fig. 5, the customized GPT is also instructed to present its output in a table format with one column showing the attack paths it predicted that could compromise the specified asset, and another column explaining the attack execution steps with associated MITRE ATT&CK techniques. Thus, the output from the customized GPT provides users with asset-centric threats that would help cybersecurity experts think critically about how such attacks could be detected, mitigated, or prevented in the first place.

VI. EVALUATION OF TRACR-TMF

To evaluate the TraCR-TMF, we consider two evaluation cases as follows: (i) the evaluation of ARC-IT transportation CPS applications, and (ii) the evaluation of a major real-world cyberattack incident. The first evaluation focuses on identifying the potential MITRE ATT&CK techniques that could exploit the vulnerabilities of different transportation

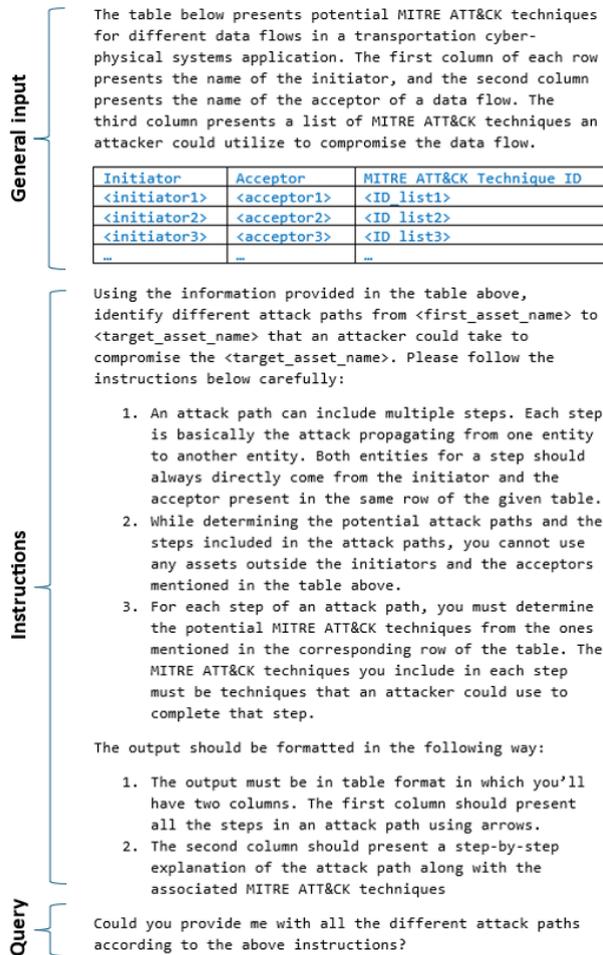


FIGURE 5. Prompt for asset-centric threat modeling.

CPS applications using the applications' planning-level systems architecture. In contrast, the second evaluation, i.e., evaluation of a major real-world cyberattack incident, focuses on asset-centric threat modeling to identify potential attack paths leading to compromising a critical transportation CPS asset, in addition to the identification of potential MITRE ATT&CK techniques that could exploit the vulnerabilities. As the real-world cyberattack incident has already taken place and the compromised assets are known for the incident, the attack paths leading to the compromised assets, as predicted by the asset-centric threat modeling, are verifiable. However, this verification is not possible with the first evaluation that focuses only on the planning-level transportation CPS applications present in ARC-IT. In addition, the LLMs, fine-tuned or customized for the ARC-IT transportation CPS application evaluation, are directly applied without any further modifications to evaluate the real-world incident considered in this study. This provides an assessment of the TraCR-TMF's transferability beyond what the framework's constituent LLMs are fine-tuned or customized for.

A. EVALUATION METRICS

To evaluate the performance of the TraCR-TMF when we utilize LLMs to identify attack techniques that are relevant to a given transportation CPS data flow, we consider three standard metrics of classification tasks: (i) precision, (ii) recall, and (iii) accuracy. For a binary classification task, these metrics are calculated as follows:

$$Precision = \frac{TP}{TP + FP} \quad (2)$$

$$Recall = \frac{TP}{TP + FN} \quad (3)$$

$$F1 = \frac{2 \cdot Precision \cdot Recall}{Precision + Recall} \quad (4)$$

where, TP represents the number of true positives, FP represents the number of false positives, and FN represents the number of false negatives. Precision and recall measure the proportion of correctly predicted positive instances among all predicted positives, and among all actual positives, respectively. F1-score is the harmonic mean of precision and recall.

To further clarify how these metrics are calculated for a multi-label classification task, let us consider that P_i denotes the predicted set, G_i denotes the ground-truth set, and O_i denotes the overlap between P_i and G_i . Then, precision and recall are given as follows:

$$Precision = \frac{1}{N} \sum_{i=1}^N \frac{|O_i|}{|P_i|} \quad (5)$$

$$Recall = \frac{1}{N} \sum_{i=1}^N \frac{|O_i|}{|G_i|} \quad (6)$$

where, N is the total number of information flows considered.

B. EVALUATION OF ARC-IT TRANSPORTATION CPS APPLICATIONS

The objective of this evaluation is to assess the TraCR-TMF's efficacy in identifying the relevant attack techniques by applying the framework to a wide range transportation CPS applications available on ARC-IT. To this end, we consider 26 different randomly selected transportation CPS applications from ARC-IT [44]. In total, we extracted 433 different data flows that are related to these transportation CPS applications along with their associated initiator, acceptor, functional objects and processes, and recommended security attribute information, such as confidentiality, integrity, availability, encryption, and authentication requirements. This information helps us construct the basic input, shown in Fig. 2, except for the part dedicated to STRIDE-based associated threats that come from the MS SDL threat modeling tool.

Since ARC-IT serves as a database for the U.S. national reference architectures for transportation CPS applications that are widely adopted by transportation CPS professionals for

implementation, it provides a detailed physical view of the architectures, outlining different involved entities and their interactions. These reference architectures provide adequate information to develop their corresponding DFDs in the MS SDL threat modeling tool. Once a transportation CPS application architecture is transferred to the MS SDL threat modeling tool, the tool automatically identifies potential threats for each data flow within the application according to the STRIDE model. Fig. 6 presents a snapshot of the threat report generated by the MS SDL tool for an example transportation CPS application that is referred to as the “CVO03: Electronic Clearance” service package in ARC-IT [55]. The threat information obtained from the MS SDL threat modeling report completes the basic input (shown in Fig. 2) requirements for each data flow.

The basic input is then coupled with instructions and queries to prepare the prompts to be processed by the LLMs to identify relevant attack techniques from the MITRE ATT&CK matrix. Unlike the RAG-based approach discussed in Section III-B, which does not require a predetermined list of relevant attack techniques, the ICL-based approach requires a few predetermined relevant attack techniques to be provided as examples. On the other hand,

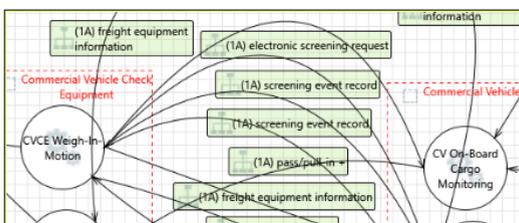
the supervised fine-tuning approach requires a ground-truth list of relevant attack techniques to train on. Preparing this predetermined or ground-truth list of relevant attack techniques for each data flow is a task that can only be done by cybersecurity experts. In this study, we established this list by consulting with two cybersecurity experts with decades of academic and industrial experience in the cybersecurity domain. Nevertheless, this list remains incomplete since the MITRE ATT&CK matrix includes over 200 different attack techniques, all of which could not be reviewed by the cybersecurity experts for each of the 433 data flows considered in this evaluation. However, unlike traditional supervised machine learning or deep learning models, LLMs are not only able to perform multi-label classification tasks based on their learning from provided examples or supervised training, but LLMs can go beyond the given examples or its training domain due to its original training on a vast amount of data and ability to understand the contexts [56].

As mentioned in Section III-B, TraCR-TMF utilizes three different LLM-based approaches to map the STRIDE-based threats to the potentially relevant attack techniques from the MITRE ATT&CK matrix, as follows: (i) a RAG-based

Diagram 1 Diagram Summary:

Not Started	341
Not Applicable	0
Needs Investigation	0
Mitigation Implemented	0
Total	341
Total Migrated	0

Interaction: (1A) electronic screening request



1. Data Flow Sniffing [State: Not Started] [Priority: High]

Category: Information Disclosure

Description: Data flowing across (1A) electronic screening request may be sniffed by an attacker. Depending on what type of data an attacker can read, it may be used to attack other parts of the system or simply be a disclosure of information leading to compliance violations. Consider encrypting the data flow.

Justification: <no mitigation provided>
2. Potential Data Repudiation by CV On-Board Cargo Monitoring [State: Not Started] [Priority: High]

Category: Repudiation

Description: CV On-Board Cargo Monitoring claims that it did not receive data from a source outside the trust boundary. Consider using logging or auditing to record the source, time, and summary of the received data.

Justification: <no mitigation provided>
3. Potential Lack of Input Validation for CV On-Board Cargo Monitoring [State: Not Started] [Priority: High]

Category: Tampering

Description: Data flowing across (1A) electronic screening request may be tampered with by an attacker. This may lead to a denial of service attack against CV On-Board Cargo Monitoring or an elevation of privilege attack against CV On-Board Cargo Monitoring or an information disclosure by CV On-Board Cargo Monitoring. Failure to verify that input is as expected is a root cause of a very large number of exploitable issues. Consider all paths and the way they handle data. Verify that all input is verified for correctness using an approved list input validation approach.

Justification: <no mitigation provided>

FIGURE 6. Snapshot of the STRIDE-based threat report of “CVO03: Electronic Clearance” application.

approach, (ii) an ICL-based approach, and (iii) a supervised fine-tuning approach. The LLM-predicted relevant attack techniques are compared with the corresponding ground-truths to determine the precision, recall, and F1 score for each LLM-based approach. Table III presents these results for the three LLM-based attack technique identification approaches.

As observed from Table III, the supervised fine-tuning approach outperformed the RAG-based and the ICL-based attack technique identification approaches. While the RAG-based approach performed the worst, it should be noted that this approach solely relied on the basic input and utilized its multistage attack identification architecture, as shown in Fig. 3, to identify relevant attack techniques from the MITRE ATT&CK matrix. In contrast, the ICL-based approach leveraged in-prompt examples to learn relevant attack techniques and the supervised fine-tuning approach leveraged retraining the pre-trained LLM to enhance the model’s attack technique identification performance. Among zero-, one-, and few-shot ICL approaches, the few-shot learning performed the best. While we tested few-shot learning by providing two to nine examples within a single prompt, few-shot learning with eight examples outperformed the other ICL approaches, which is reported in Table III as few-shot learning. This better performance of eight example-based few-shot learning can be attributed to striking the right balance between better learning from additional examples and potential hallucinations due to a context overload issue induced by lengthy prompts with excess examples [57]. Although the supervised fine-tuning approach outperformed the best-performing ICL-based approach in this evaluation, fine-tuning requires establishing ground truths, which could be challenging to establish with limited cybersecurity knowledge and expertise. In contrast, providing a handful of examples to leverage ICL might seem easier for adoption in some cases.

Another insight into the performance of the LLM-based relevant attack identification approaches is that the LLMs identified several attack techniques outside their training domain, several of which were found to be relevant. As mentioned earlier, the ground-truth list is not complete. The cybersecurity experts, whom we consulted to develop the ground-truth list, reviewed several, but not all, of the existing MITRE ATT&CK techniques. Even if we had established a ground-truth list by considering every attack technique from the MITRE ATT&CK matrix, new attack techniques would be introduced and added to the matrix in the future. Thus, LLMs’ ability to identify relevant attack techniques beyond their training domain is useful in this regard.

To better assess LLM’s performance in identifying relevant attack techniques from the MITRE ATT&CK matrix, we randomly selected 50 data flows out of the 433 data flows and again consulted with the cybersecurity experts to evaluate each LLM-predicted attack technique that is not present in our ground-truth list of relevant attack techniques. Based on this second stage validation from the cybersecurity experts, we reevaluated the performance of the supervised fine-tuning

approach on the randomly selected 50 data flows, which is presented in Table IV. From Table IV, it is observed that once we updated the ground-truth list with the help of the cybersecurity experts considering all the attack techniques identified by the supervised LLM, the precision improved from 0.61 to 0.90 (i.e., a 47.5% improvement) for the randomly selected 50 data flows. This high precision indicates that about 90% of the attack techniques identified by the supervised LLM were correct. This demonstrates that the supervised LLM can identify more types of attack techniques from the MITRE ATT&CK matrix than it was originally trained on during fine-tuning, which can be attributed to its prior extensive knowledge base and contextual reasoning ability. This helps us consider the supervised LLM-based approach for our second evaluation case related to a major real-world cyber incident.

C. EVALUATION OF A MAJOR REAL-WORLD CYBERATTACK INCIDENT

The objectives of this real-world cyberattack incident evaluation are as follows: (i) to assess the transferability of the TraCR-TMF by applying its underlying supervised LLM to transportation CPS outside the model’s training domain, (ii) to perform an asset-centric threat modeling of transportation CPS using the TraCR-TMF to identify potential attack paths and associated attack techniques leading to a target asset, (iii) to assess the attack paths and the associated attack techniques identified by the framework. To this end, we selected the Colonial Pipeline double extortion ransomware attack, a major real-world cyberattack incident that took place in 2021.

This Colonial Pipeline incident serves as an evaluation case for which we can assess the attack paths and the associated techniques identified by the TraCR-TMF. To perform an

TABLE III
PERFORMANCE OF TRACR-TMF IN IDENTIFYING RELEVANT ATT&CK TECHNIQUES

LLM-based Approach	Precision	Recall	F1 Score
RAG	0.33	0.20	0.25
ICL	Zero-shot	0.18	0.18
	One-shot	0.38	0.31
	Few-shot	0.47	0.48
Supervised fine-tuning	0.77	0.62	0.69

TABLE IV
Performance of Supervised LLM on 50 Randomly Selected Transportation CPS Data Flows

Supervised LLM-based Approach	Precision	Recall	F1 Score
Based on initial ground truths set by the experts	0.61	0.48	0.54
Based on the LLM-identified attack techniques with subsequent expert validation	0.90	0.55	0.69

asset-centric threat modeling using our framework, a target asset must be specified for which we aim to identify (i) the potential attack paths leading to that asset, and (ii) the associated attack techniques for each exploitation involved in those attacks. Ideally, if the framework could accurately predict the potential attack paths and associated attack techniques, an attacker could utilize any of those means to compromise the target asset. Choosing a widely discussed real-world cyberattack incident, like the Colonial Pipeline, is useful for this evaluation since different public and private security organizations have published several reports highlighting their insights related to the incident, which helps us better understand the actual cyberattack incident and compare the predictions of our framework with the actual incident. In addition, to identify the relevant attack techniques for each data flow, the LLM fine-tuned for our first evaluation case is utilized without any additional training, which enables us to assess the transferability of the TraCR-TMF. In the following sections, we provide a brief description of how the attack took place to the best of our knowledge, state our assumptions for this evaluation, present the threat modeling evaluation conducted based on the TraCR-TMF, and discuss the evaluation results.

1) COLONIAL PIPELINE RANSOMWARE: KNOWN FACTS

Colonial Pipeline is among the major movers of gasoline and other refined fuels on the East Coast of the U.S. In May 2021, an attacker group utilized the Darkside malware, a ransomware-as-a-service (RaaS) provided by a hacker group, to attack the Colonial Pipeline network [58], [59]. About a week after the attack began, the attacker group demanded a \$4.4 million ransom in cryptocurrency from the Colonial Pipeline authority. Within this period, the attacker exfiltrated over 100 gigabytes of data outside the Colonial Pipeline’s information technology (IT) network before encrypting the company’s valuable data [60].

According to multiple sources [58], [60], [61], the attackers accessed the IT network of the Colonial Pipeline through a compromised virtual private network (VPN) account near the end of April 2021. The account credentials might have been leaked as part of a separate data breach, or the attackers might have blackmailed a legitimate account holder, or the attackers might even guess the credentials, which could not be confirmed based on the publicly available resources. In addition, the VPN account did not have multi-factor authentication (MFA) enabled, which helped the attackers gain access to the VPN using the compromised VPN account credentials only [60], [61]. Once gained access, the attackers moved laterally within the IT network, gradually exfiltrated over 100 gigabytes of the company’s valuable data, and installed ransomware to encrypt the files. Finally, on May 7, the attackers left a note demanding ransom on one of the company’s computers [60].

2) ASSUMPTIONS FOR COLONIAL PIPELINE EVALUATION

To apply the TraCR-TMF for carrying out a threat modeling of the Colonial Pipeline’s industrial control systems, we need its systems architecture first. However, the actual architecture of the pipeline’s industrial control systems is a confidential information that is not publicly available. To this end, we assume that the Colonial Pipeline follows a network architecture similar to the widely known Purdue Model for industrial control systems networks [62] for this evaluation. The Purdue model provides a hierarchical design framework that segregates the software and hardware parts of the network into six different levels, as depicted in Fig. 7. In this architecture, levels 0 through 3 comprise the operational technology (OT) network, and levels 4 through 5 represent the IT network. During the actual ransomware attack, the attackers compromised levels 4 and 5 (i.e., the company’s IT network), and there was no evidence indicating the attackers got into the OT network [58]. Once the ransomware incident

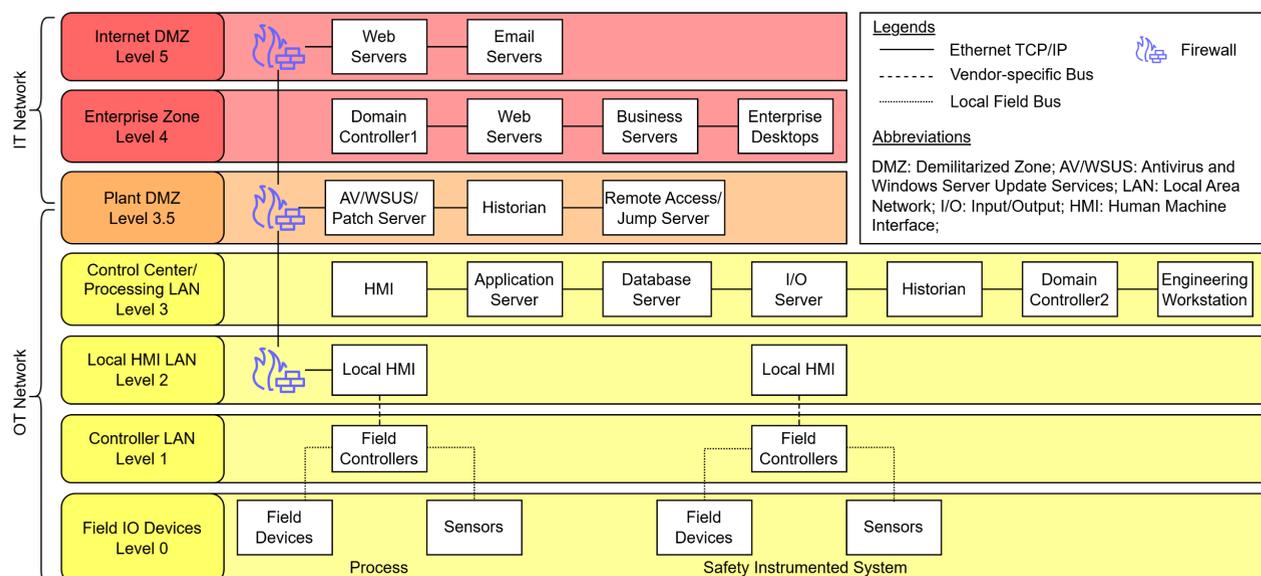


FIGURE 7. Purdue model for industrial control systems (adopted from [62]).

was identified, the Colonial Pipeline authority disconnected its OT network from the IT network to contain the attack from spreading into the OT network and the systems connected to it [58], [59].

3) COLONIAL PIPELINE THREAT MODELING

First, we develop the DFD of the Colonial Pipeline industrial control systems in the MS SDL threat modeling tool based on the Purdue model architecture presented in Fig. 7. Second, the STRIDE-based threat report generated by the MS SDL threat modeling tool is used to construct the basic input shown in Fig. 2. Third, the basic input is processed by the supervised LLM trained for the first evaluation case to identify relevant attack techniques from the MITRE ATT&CK matrix. Fourth, the asset-centric threat modeling prompt (shown in Fig. 5) is constructed based on the identified attack techniques. In this prompt, we assign the Business Servers (shown in Fig. 7) as the target asset to be compromised and a human user as the starting point of the attack paths to be identified. Finally, the prompt is processed by the customized cybersecurity analyst GPT model to identify potential attack paths and attack techniques associated with each step of these attack paths that could be utilized by attackers to compromise the Business Servers within the Colonial Pipeline IT network.

4) COLONIAL PIPELINE EVALUATION RESULTS AND DISCUSSIONS

Table V presents three different potential attack paths and associated attack techniques identified by the customized cybersecurity analyst GPT model in our TraCR-TMF. The model predicted that any of these three attack paths could be followed by an attacker to perform an attack compromising the Business Servers in the Colonial Pipeline IT network. To assess whether these attack paths and their associated attack techniques from the MITRE ATT&CK matrix are feasible, we again consulted the cybersecurity experts. Their expert opinions validating the attack paths and their associated attack

techniques from the MITRE ATT&CK matrix are summarized below:

- Potential Attack Path #1:** T1552 (Unsecured Credentials) [63] is not a directly relevant technique for this attack path as this technique is about obtaining unsecured stored credentials, which the Colonial Pipeline attacker already possessed. On the other hand, T1059 (Command and Scripting Interpreter) [64] is relevant because the initial access payload must have included some remote execution-related commands and scripts that helped the adversary get within the systems. T1105 (Ingress Tool Transfer) [65] goes beyond maintaining presence and is used to transfer/spread tools between victim devices within a compromised environment; therefore, it is not relevant in this context. On the other hand, T1040 (Network Sniffing) [66] is relevant because it is a basic approach to capture network traffic to realize the network configurations. Similarly, T1557 (Adversary-in-the-Middle) [67] is among the open-source techniques commonly used by the ALPHV Blackcat group [68], believed to be associated with DarkSide, to obtain credentials and session cookies that help attackers move laterally within a network; therefore, it is a relevant technique for this attack path. T1565 (Data Manipulation) [69] is also relevant because the attackers could delete some critical data after exfiltrating it outside the company's network. However, T1495 (Firmware Corruption) [70] is not relevant as there is no evidence that indicates the attackers performed firmware modification during the Colonial Pipeline incident.
- Potential Attack Path #2:** T1059 (Command and Scripting Interpreter), T1557 (Adversary-in-the-Middle), and T1040 (Network Sniffing) are relevant for this attack path as well, while T1105 (Ingress Tool Transfer) and T1495 (Firmware Corruption) are not

TABLE V
PREDICTED PATHS FOR COLONIAL PIPELINE RANSOMWARE ATTACK

#	Predicted Attack Path	Executions Steps using ATT&CK Techniques
1	Human User → VPN Server → Business Servers	<ol style="list-style-type: none"> The attacker compromises the VPN server using stolen credentials (T1552) or remote command execution (T1059). Deploys a malicious payload (T1105) to maintain persistence. Uses Network Sniffing (T1040) to gather intelligence about internal traffic. Moves laterally to Business Servers using Man-in-the-Middle (T1557) or privilege escalation techniques. Executes data manipulation (T1565) or firmware corruption (T1495) on Business Servers.
2	Human User → VPN Server → Domain Controller1 → Business Servers	<ol style="list-style-type: none"> The attacker gains access to VPN server using stolen credentials (T1552) or by executing remote commands (T1059). Moves to Domain Controller1 using Ingress Tool Transfer (T1105) or Man-in-the-Middle (T1557). Gains domain admin privileges and escalates access (T1548). Moves laterally to Business Server using Unsecured Credentials (T1552) or Network Sniffing (T1040). Executes final payload, exfiltrates data (T1020), or corrupts the system (T1495).
3	Human User → VPN Server → Remote Access/Jump Server → Historian1 → Business Servers	<ol style="list-style-type: none"> The attacker exploits VPN server using credential theft (T1552) or executes malicious scripts (T1059). Moves to Remote Access/Jump Server via Man-in-the-Middle (T1557) or Network Sniffing (T1040). Gains access to Historian1 by using credential access (T1087) or executing remote commands (T1059). Uses Historian1 to pivot towards Business Servers via privilege escalation (T1548) or lateral movement techniques. Data exfiltration (T1020), encryption for ransom (T1486), or system manipulation (T1565).

relevant, as explained for the first predicted attack path. T1548 (Abuse Elevation Control Mechanism) [71] is relevant, because once the attackers gain access to the Domain Controller1, they could elevate their privileges to help move to the Business Servers. Although T1552 (Unsecured Credentials) is not relevant for step 1 of this attack path (shown in Table V), it is relevant for step 4 to obtain the credentials for Business Servers. On the other hand, T1020 (Automated Exfiltration) is directly relevant because the attackers exfiltrated 100 gigabytes of data during the Colonial Pipeline ransomware attack.

- **Potential Attack Path #3:** Like the first two attack paths, T1059 (Command and Scripting Interpreter), T1557 (Adversary-in-the-Middle), T1040 (Network Sniffing), T1548 (Abuse Elevation Control Mechanism), T1020 (Automated Exfiltration), and T1565 (Data Manipulation) are all relevant for this attack path as well. T1552 (Unsecured Credentials) is not relevant for step 1 of this attack path, as explained for the first attack path. On the other hand, T1087 (Account Discovery) is relevant to discovering credentials for accessing Historian1. Similarly, T486 (Data Encrypted for Impact) is directly relevant since attackers encrypted data to demand ransom during the Colonial Pipeline incident.

Based on this discussion, we observe that the TraCR-TMF was able to identify potential attack techniques from the MITRE ATT&CK matrix and devise attack paths based on them, which could have led to the compromised Business Servers during the actual Colonial Pipeline incident. The cybersecurity experts' validation showed us that most of the attack techniques associated with the steps for each predicted attack path presented in Table V are contextually reasonable, although the supervised LLM used here was not trained for the Colonial Pipeline evaluation case.

However, there is no way of verifying these predicted attack paths and the associated attack techniques against the actual attack path and associated attack techniques utilized by the attackers during their Colonial Pipeline ransomware attack, because such information is not publicly available. Nevertheless, we can assess whether our framework is able to identify some key exploitations of the actual Colonial Pipeline cyberattack incident. Table VI presents what exploitations we know about the incident with certainty based on publicly available resources, and the relevance of our framework's

prediction to those exploitations, for example, data exfiltration and ransomware. We observe from Table VI that the TraCR-TMF was able to capture the known exploitations of the actual Colonial Pipeline cyberattack incident through its predicted attack paths and associated attack techniques. This demonstrates the efficacy and transferability of our TraCR-TMF for transportation CPS' threat modeling by validating with a real-world cyberattack incident that falls outside the training domain of the framework's constituent models.

VII. CONCLUSIONS

This study presents TraCR-TMF, an LLM-supported framework for threat modeling of transportation CPS, which requires limited cybersecurity expert intervention. TraCR-TMF leverages open-source tools, databases, and LLMs to identify vulnerabilities in transportation CPS, corresponding attack techniques, and potential attack paths that could lead to the compromise of critical assets. By mapping attack techniques to specific vulnerabilities within transportation CPS, the framework empowers cybersecurity professionals dealing with transportation CPS to critically evaluate both proactive and reactive defense strategies. In addition, equipped with insights into relevant attack techniques and paths leading to the compromise of critical transportation CPS assets, cybersecurity professionals can proactively tailor their cyberattack countermeasures. Notably, TraCR-TMF draws from the MITRE ATT&CK matrix to identify applicable attack techniques, providing users with immediate access to established detection and mitigation strategies, which is a valuable starting point for implementing cybersecurity measures.

TraCR-TMF supports three LLM-based alternative strategies. Of these, the RAG-based approach demands the least cybersecurity expert intervention but is also the least effective in identifying relevant attack techniques. In contrast, both the ICL-based and supervised fine-tuning approaches demonstrate improved performance but require varying degrees of input from cybersecurity experts. While the supervised fine-tuning approach yielded the best performance in this study, benefiting from targeted retraining on the attack identification task with expert cybersecurity input, the performance gap between this method and the other two could narrow as LLMs continue to evolve. In particular, the emergence of models with more advanced reasoning capabilities, such as the anticipated advent of artificial general

TABLE VI
COMPARISON OF TRACR-TMF'S PREDICTIONS WITH THE COLONIAL PIPELINE CYBERATTACK INCIDENT FACTS

Confirmed Events Based on Publicly Available Knowledge	Relevance of TraCR-TMF's Predictions
The attackers moved laterally within the Colonial Pipeline IT network	All three predicted attack paths show that an attacker could move laterally within the IT network with the help of different relevant attack techniques
The attackers exfiltrated valuable data outside the company's network	Predicted attack paths #2 and #3 include data exfiltration potentials using T1020 technique from the MITRE ATT&CK matrix
The attackers demanded ransom after encrypting the company's valuable data	Encryption for ransom or impact using the MITRE ATT&CK technique T1486 is included in the predicted attack path #3

intelligence (AGI), holds promise for enhancing the effectiveness of less cybersecurity expert-intervention-intensive approaches. In future work, we will explore the integration of such advanced LLMs into the TraCR-TMF framework.

Although the evaluation scenarios in this study focused on transportation CPS and the LLMs were tailored to that domain, the TraCR-TMF framework is broadly applicable to CPS across other domains. The tools and knowledge bases employed in TraCR-TMF, such as the MS SDL threat modeling tool and the MITRE ATT&CK matrix, are not specific to transportation CPS and can be applied across various domains. Moreover, the three LLM-supported approaches introduced in this study for mapping existing threats to specific adversarial techniques require varying levels of cybersecurity expert intervention, offering flexible adaptation options to suit the needs of different domains.

ACKNOWLEDGMENT

The authors would like to acknowledge Dr. Bhavani Thuraisingham, Professor of Computer Science, the University of Texas at Dallas, and Dr. Amjad Ali, Professor of Computer Science, Morgan State University, for their input to this article.

REFERENCES

- [1] K. N. Qureshi and A. H. Abdullah, "A survey on intelligent transportation systems," *Middle-East Journal of Scientific Research*, vol. 15, no. 5, pp. 629–642, 2013.
- [2] A. Lamssaggad, N. Benamar, A. S. Hafid, and M. Msahli, "A Survey on the Current Security Landscape of Intelligent Transportation Systems," *IEEE Access*, vol. 9, pp. 9180–9208, 2021, doi: 10.1109/ACCESS.2021.3050038.
- [3] D. Hahn, A. Munir, and V. Behzadan, "Security and Privacy Issues in Intelligent Transportation Systems: Classification and Challenges," *IEEE Intelligent Transportation Systems Magazine*, vol. 13, no. 1, pp. 181–196, 2021, doi: 10.1109/MITS.2019.2898973.
- [4] "Under Attack: Trucking industry increasingly at risk of cyberattacks," *Truck News*. Accessed: Mar. 25, 2025. [Online]. Available: <https://www.trucknews.com/features/under-attack-trucking-industry-increasingly-at-risk-of-cyberattacks/>
- [5] "Aviation is facing a rising wave of cyber-attacks in the wake of COVID." Accessed: Mar. 25, 2025. [Online]. Available: <https://www.stephenharwood.com/insights/aviation-is-facing-a-rising-wave-of-cyber-attacks-in-the-wake-of-covid?06082024160610&26032025015047>
- [6] "IAPH Cybersecurity Guidelines for Ports and Port Facilities," International Association of Ports and Harbors. Accessed: Mar. 25, 2025. [Online]. Available: https://sustainableworldports.org/wp-content/uploads/IAPH-Cybersecurity-Guidelines-version-I_0.pdf
- [7] W. Xiong and R. Lagerström, "Threat modeling – A systematic literature review," *Computers & Security*, vol. 84, pp. 53–69, Jul. 2019, doi: 10.1016/j.cose.2019.03.010.
- [8] F. Swiderski and W. Snyder, *Threat modeling*. Microsoft Press, 2004.
- [9] X. He *et al.*, "DNS Rebinding Threat Modeling and Security Analysis for Local Area Network of Maritime Transportation Systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 2, pp. 2643–2655, Feb. 2023, doi: 10.1109/TITS.2021.3135197.
- [10] W. Xiong, F. Krantz, and R. Lagerström, "Threat Modeling and Attack Simulations of Connected Vehicles: Proof of Concept," in *Information Systems Security and Privacy*, P. Mori, S. Furnell, and O. Camp, Eds., Cham: Springer International Publishing, 2020, pp. 272–287. doi: 10.1007/978-3-030-49443-8_13.
- [11] M. A. Ramazanzadeh, B. Barzegar, and H. Motameni, "ASATM: Automated security assistant of threat models in intelligent transportation systems," *IET Computers & Digital Techniques*, vol. 16, no. 5–6, pp. 141–158, 2022, doi: 10.1049/cdt2.12045.
- [12] J. Steven, "Threat Modeling - Perhaps It's Time," *IEEE Security & Privacy*, vol. 8, no. 3, pp. 83–86, May 2010, doi: 10.1109/MSP.2010.110.
- [13] A. Konev, A. Shelupanov, M. Kataev, V. Ageeva, and A. Nabieva, "A Survey on Threat-Modeling Techniques: Protected Objects and Classification of Threats," *Symmetry*, vol. 14, no. 3, Art. no. 3, Mar. 2022, doi: 10.3390/sym14030549.
- [14] S. M. Khalil, H. Bahsi, and T. Korötko, "Threat modeling of industrial control systems: A systematic literature review," *Computers & Security*, vol. 136, p. 103543, Jan. 2024, doi: 10.1016/j.cose.2023.103543.
- [15] M. Aijaz, M. Nazir, and M. N. A. Mohammad, "Threat Modeling and Assessment Methods in the Healthcare-IT System: A Critical Review and Systematic Evaluation," *SN COMPUT. SCI.*, vol. 4, no. 6, p. 714, Sep. 2023, doi: 10.1007/s42979-023-02221-1.
- [16] M. Erbas, S. M. Khalil, and L. Tsiopoulos, "Systematic literature review of threat modeling and risk assessment in ship cybersecurity," *Ocean Engineering*, vol. 306, p. 118059, Aug. 2024, doi: 10.1016/j.oceaneng.2024.118059.
- [17] H. Cho and S. Kim, "Threat Modeling for the Defense Industry: Past, Present, and Future," *IEEE Access*, vol. 13, pp. 53276–53304, 2025, doi: 10.1109/ACCESS.2025.3550337.
- [18] S. Nagasundari, P. Manja, P. Mathur, and P. B. Honnavalli, "Extensive Review of Threat Models for DevSecOps," *IEEE Access*, vol. 13, pp. 45252–45271, 2025, doi: 10.1109/ACCESS.2025.3547932.
- [19] L. O. Nweke and S. Wolthusen, "A Review of Asset-Centric Threat Modelling Approaches," *I-6*, 2020, doi: 10.14569/IJACSA.2020.0110201.
- [20] E. N. Crothers, N. Japkowicz, and H. L. Viktor, "Machine-Generated Text: A Comprehensive Survey of Threat Models and Detection Methods," *IEEE Access*, vol. 11, pp. 70977–71002, 2023, doi: 10.1109/ACCESS.2023.3294090.
- [21] M. Kharma and A. Taweel, "Threat Modeling in Cloud Computing - A Literature Review," in *Ubiquitous Security*, G. Wang, K.-K. R. Choo, J. Wu, and E. Damiani, Eds., Singapore: Springer Nature, 2023, pp. 279–291. doi: 10.1007/978-981-99-0272-9_19.
- [22] F. Luo, Y. Jiang, Z. Zhang, Y. Ren, and S. Hou, "Threat Analysis and Risk Assessment for Connected Vehicles: A Survey," *Security and Communication Networks*, vol. 2021, no. 1, p. 1263820, 2021, doi: 10.1155/2021/1263820.
- [23] A. V. Uzunov and E. B. Fernandez, "An extensible pattern-based library and taxonomy of security threats for distributed systems," *Computer Standards & Interfaces*, vol. 36, no. 4, pp. 734–747, Jun. 2014, doi: 10.1016/j.csi.2013.12.008.
- [24] A. O. Baquero, A. Kornecki, and J. Zalewski, "Threat modeling for aviation computer security," *crossstalk*, vol. 28, no. 6, pp. 21–27, 2015.
- [25] W. Xiong, E. Legrand, O. Åberg, and R. Lagerström, "Cyber security threat modeling based on the MITRE Enterprise ATT&CK Matrix," *Softw Syst Model*, vol. 21, no. 1, pp. 157–177, Feb. 2022, doi: 10.1007/s10270-021-00898-7.
- [26] S. Al-Fedaghi and A. Alkandari, "On security development lifecycle: conceptual description of vulnerabilities, risks, and threats," *International Journal of Digital Content Technology and its Applications*, vol. 5, no. 5, pp. 296–306, 2011.
- [27] P. Torr, "Demystifying the threat modeling process," *IEEE Security & Privacy*, vol. 3, no. 5, pp. 66–70, Sep. 2005, doi: 10.1109/MSP.2005.119.
- [28] C. Alberts, A. Dorofee, J. Stevens, and C. Woody, "Introduction to the OCTAVE® Approach," Carnegie Mellon Software Engineering Institute, Aug. 2003. Accessed: Apr. 03, 2025. [Online]. Available: https://insights.sei.cmu.edu/documents/16/2003_012_001_51556.pdf
- [29] "Microsoft Threat Modeling Tool overview - Azure." Accessed: Mar. 18, 2025. [Online]. Available: <https://learn.microsoft.com/en-us/azure/security/develop/threat-modeling-tool>

- [30] M. Deng, K. Wuyts, R. Scandariato, B. Preneel, and W. Joosen, "A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements," *Requirements Eng*, vol. 16, no. 1, pp. 3–32, Mar. 2011, doi: 10.1007/s00766-010-0115-7.
- [31] T. UcedaVelez and M. M. Morana, *Risk Centric Threat Modeling: Process for Attack Simulation and Threat Analysis* | Wiley, Wiley, 2015. Accessed: Apr. 07, 2025. [Online]. Available: <https://www.wiley.com/en-us/Risk+Centric+Threat+Modeling%3A+Process+for+Attack+Simulation+and+Threat+Analysis-p-9780470500965>
- [32] V. Saini, Q. Duan, and V. Paruchuri, "Threat modeling using attack trees," *Journal of Computing Sciences in Colleges*, vol. 23, no. 4, pp. 124–131, 2008.
- [33] B. Schneier, "Attack trees," *Dr. Dobbs' journal*, vol. 24, no. 12, pp. 21–29, 1999.
- [34] P. Kumar, G. P. Gupta, R. Tripathi, S. Garg, and M. M. Hassan, "DLTIF: Deep Learning-Driven Cyber Threat Intelligence Modeling and Identification Framework in IoT-Enabled Maritime Transportation Systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 2, pp. 2472–2481, Feb. 2023, doi: 10.1109/TITS.2021.3122368.
- [35] A. K. Subran, S. Sankaran, P. Sutraye, and A. K. Nishad, "Threat Modeling and Attack Simulation of Charging Infrastructures in Electric Vehicles," in *2024 15th International Conference on Computing Communication and Networking Technologies (ICCCNT)*, Jun. 2024, pp. 1–7. doi: 10.1109/ICCCNT61001.2024.10725968.
- [36] M. Hamad, M. Nolte, and V. Prevelakis, "Towards comprehensive threat modeling for vehicles," presented at the the 1st Workshop on Security and Dependability of Critical Embedded Real-Time Systems, 2016, p. 31.
- [37] M. B. Munir, Y. Cai, L. Khan, and B. Thuraisingham, "Leveraging Multimodal Retrieval-Augmented Generation for Cyber Attack Detection in Transit Systems," in *2024 IEEE 6th International Conference on Trust, Privacy and Security in Intelligent Systems, and Applications (TPS-ISA)*, Oct. 2024, pp. 341–350. doi: 10.1109/TPS-ISA62245.2024.00046.
- [38] R. Gabrys, M. Bilinski, S. Fugate, and D. Silva, "Using Natural Language Processing Tools to Infer Adversary Techniques and Tactics Under the Mitre ATT&CK Framework," in *2024 IEEE 14th Annual Computing and Communication Workshop and Conference (CCWC)*, Jan. 2024, pp. 0541–0547. doi: 10.1109/CCWC60891.2024.10427746.
- [39] I. Branescu, O. Grigorescu, and M. Dascalu, "Automated Mapping of Common Vulnerabilities and Exposures to MITRE ATT&CK Tactics," *Information*, vol. 15, no. 4, Art. no. 4, Apr. 2024, doi: 10.3390/info15040214.
- [40] D. Nir *et al.*, "Labeling Network Intrusion Detection System (NIDS) Rules with MITRE ATT&CK Techniques: Machine Learning vs. Large Language Models," *Big Data and Cognitive Computing*, vol. 9, no. 2, p. 23, 2025, doi: 10.3390/bdcc9020023.
- [41] "Threats - Microsoft Threat Modeling Tool - Azure." Accessed: Mar. 19, 2025. [Online]. Available: <https://learn.microsoft.com/en-us/azure/security/develop/threat-modeling-tool-threats>
- [42] A. Shostack, "Experiences Threat Modeling at Microsoft," *MODSEC@ MoDELS*, vol. 2008, p. 35, 2008.
- [43] "MITRE ATT&CK®." Accessed: Mar. 22, 2025. [Online]. Available: <https://attack.mitre.org/>
- [44] "Architecture Reference for Cooperative and Intelligent Transportation." Accessed: Mar. 22, 2025. [Online]. Available: <https://www.arc-it.net/index.html>
- [45] P. Lewis *et al.*, "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," in *Advances in Neural Information Processing Systems*, Curran Associates, Inc., 2020, pp. 9459–9474. Accessed: Mar. 21, 2025. [Online]. Available: <https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html>
- [46] OpenAI *et al.*, "GPT-4o System Card," Oct. 25, 2024, *arXiv:arXiv:2410.21276*. doi: 10.48550/arXiv.2410.21276.
- [47] J. Achiam *et al.*, "GPT-4 Technical Report," Mar. 04, 2024, *arXiv:arXiv:2303.08774*. doi: 10.48550/arXiv.2303.08774.
- [48] "Simple Directory Reader - LlamaIndex v0.10.10." Accessed: Mar. 26, 2025. [Online]. Available: https://llamaindexxx.readthedocs.io/en/latest/examples/data_connectors/simple_directory_reader.html
- [49] "Vector embeddings - OpenAI API." Accessed: Mar. 22, 2025. [Online]. Available: <https://platform.openai.com>
- [50] "New embedding models and API updates." Accessed: Mar. 25, 2025. [Online]. Available: <https://openai.com/index/new-embedding-models-and-api-updates/>
- [51] "VectorIndexRetriever," LlamaIndex.TS. Accessed: Apr. 09, 2025. [Online]. Available: <https://ts.llamaindex.ai/docs/api/classes/VectorIndexRetriever>
- [52] "Node Postprocessor Modules." Accessed: Mar. 26, 2025. [Online]. Available: https://docs.llamaindex.ai/en/stable/module_guides/querying/node_postprocessors/node_postprocessors/
- [53] B. Warner *et al.*, "Smarter, Better, Faster, Longer: A Modern Bidirectional Encoder for Fast, Memory Efficient, and Long Context Finetuning and Inference," Dec. 19, 2024, *arXiv:arXiv:2412.13663*. doi: 10.48550/arXiv.2412.13663.
- [54] "Creating a GPT | OpenAI Help Center." Accessed: Mar. 23, 2025. [Online]. Available: <https://help.openai.com/en/articles/8554397-creating-a-gpt>
- [55] "Electronic Clearance." Accessed: Mar. 24, 2025. [Online]. Available: <https://www.arc-it.net/html/servicepackages/sp168.html#tab-3>
- [56] P. Törnberg, "Large Language Models Outperform Expert Coders and Supervised Classifiers at Annotating Political Social Media Messages," *Social Science Computer Review*, p. 08944393241286471, Sep. 2024, doi: 10.1177/08944393241286471.
- [57] G. Perković, A. Drobniak, and I. Botički, "Hallucinations in LLMs: Understanding and Addressing Challenges," in *2024 47th MIPRO ICT and Electronics Convention (MIPRO)*, May 2024, pp. 2084–2088. doi: 10.1109/MIPRO60963.2024.10569238.
- [58] "DarkSide Ransomware: Best Practices for Preventing Business Disruption from Ransomware Attacks | CISA." Accessed: Mar. 21, 2025. [Online]. Available: <https://www.cisa.gov/news-events/cybersecurity-advisories/aa21-131a>
- [59] J. Chris and P. W. Parfomak, "Colonial Pipeline: The DarkSide Strikes," Congressional Research Service, legislation, May 2011. Accessed: Mar. 21, 2025. [Online]. Available: <https://www.congress.gov/crs-product/IN11667>
- [60] J. Beerman, D. Berent, Z. Falter, and S. Bhunia, "A Review of Colonial Pipeline Ransomware Attack," in *2023 IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing Workshops (CCGridW)*, May 2023, pp. 8–15. doi: 10.1109/CCGridW59191.2023.00017.
- [61] "Back to Basics: A Deeper Look at the Colonial Pipeline Hack," GovTech. Accessed: Apr. 07, 2025. [Online]. Available: <https://www.govtech.com/sponsored/back-to-basics-a-deeper-look-at-the-colonial-pipeline-hack>
- [62] D. Garton, "Purdue model framework for industrial control systems & cybersecurity segmentation," *US Department of Energy*, vol. 14, pp. 2022–10, 2022.
- [63] "Unsecured Credentials, Technique T1552 - Enterprise | MITRE ATT&CK®." Accessed: Apr. 07, 2025. [Online]. Available: <https://attack.mitre.org/techniques/T1552/>
- [64] "Command and Scripting Interpreter, Technique T1059 - Enterprise | MITRE ATT&CK®." Accessed: Apr. 07, 2025. [Online]. Available: <https://attack.mitre.org/techniques/T1059/>
- [65] "Ingress Tool Transfer, Technique T1105 - Enterprise | MITRE ATT&CK®." Accessed: Apr. 07, 2025. [Online]. Available: <https://attack.mitre.org/techniques/T1105/>
- [66] "Network Sniffing, Technique T1040 - Enterprise | MITRE ATT&CK®." Accessed: Apr. 07, 2025. [Online]. Available: <https://attack.mitre.org/techniques/T1040/>
- [67] "Adversary-in-the-Middle, Technique T1557 - Enterprise | MITRE ATT&CK®." Accessed: Apr. 07, 2025. [Online]. Available: <https://attack.mitre.org/techniques/T1557/>
- [68] "#StopRansomware: ALPHV Blackcat | CISA." Accessed: Apr. 07, 2025. [Online]. Available: <https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-353a>

- [69] "Data Manipulation, Technique T1565 - Enterprise | MITRE ATT&CK®." Accessed: Apr. 07, 2025. [Online]. Available: <https://attack.mitre.org/techniques/T1565/>
- [70] "Firmware Corruption, Technique T1495 - Enterprise | MITRE ATT&CK®." Accessed: Apr. 07, 2025. [Online]. Available: <https://attack.mitre.org/techniques/T1495/>
- [71] "Abuse Elevation Control Mechanism, Technique T1548 - Enterprise | MITRE ATT&CK®." Accessed: Apr. 07, 2025. [Online]. Available: <https://attack.mitre.org/techniques/T1548/>



M Sabbir Salek (Member, IEEE) received his Ph.D. and M.S. in civil engineering from Clemson University, Clemson, SC, USA, in 2023 and 2021, respectively. Dr. Salek received his bachelor's in mechanical engineering from Bangladesh University of Engineering & Technology (BUET), Dhaka, Bangladesh, in 2016.

Currently, he is working as a Senior Engineer for the USDOT-supported National Center for Transportation Cybersecurity and Resiliency (TraCR) at Greenville, SC, USA. He is also a Senior Research Engineer for the USDOT-supported Center for Regional and Rural Connected Communities (CR2C2), headquartered in Greensboro, NC, USA. He is an Adjunct Faculty with the Glenn Department of Civil Engineering at Clemson University, Clemson, SC, USA. His research interests include cybersecurity and resiliency of intelligent transportation systems (ITS), connected and autonomous vehicle technologies, and advanced, high-performance computing for ITS applications.



Mashrur Chowdhury (Senior Member, IEEE) completed his Ph.D. in civil engineering from the University of Virginia in 1995. Currently, Dr. Chowdhury is serving as the Eugene Douglas Mays Chair of Transportation in the Glenn Department of Civil Engineering at Clemson University. Dr. Chowdhury also holds a joint appointment from the School of Computing (courtesy appointment). Dr. Chowdhury serves as the founding director of the USDOT-sponsored

National Center for Transportation Cybersecurity and Resiliency (TraCR) and the USDOT-sponsored Center for Connected Multimodal Mobility (C²M²). He is also the director of the Complex Systems, Analytics, and Visualization Institute (CSAVI) and a co-associate director of the USDOT-sponsored Center for Regional and Rural Connected Communities (CR²C²). His current research focuses on the evolving realms of sensing, communications, computing, cybersecurity, and cyber-resiliency, all with the goal of establishing a secure and resilient IoT environment for smart cities and regions.



Muhaimin Bin Munir (Student Member, IEEE) is doing his PhD in computer science from the University of Texas at Dallas, Richardson, USA. He received his bachelor's in computer science and engineering from Military Institute of Science and Technology, Dhaka, Bangladesh, in 2020.

He is working as a Graduate Research Assistant at the University of Texas at Dallas, where he focuses on cyber threat modeling, multimodal and the application of large language models (LLMs) in cybersecurity. Previously he served as a lecturer at the Military Institute of Science and Technology. His research interests include multimodal LLMs, cyber security and blockchain.



Yuchen Cai received his M.S. in computer science from the University of Massachusetts Amherst, MA, USA, in 2020. Yuchen received his B.Mgt. in information management and information system from Beijing University of Technology, Beijing, China, in 2015.

Currently, he is a Ph.D. student in computer science and a Research Assistant at the University of Texas at Dallas, TX, USA. His research interests include Machine Learning, with a focus on Large Language Models (LLMs), Code Intelligence, and Intelligent Transportation Systems.



Mohammad Imtiaz Hasan (Student Member, IEEE) received his bachelor's in electrical and electronics engineering from Bangladesh University of Engineering & Technology (BUET), Dhaka, Bangladesh, in 2015. Currently, Mohammad is a Ph.D. student in the Glenn Department of Civil Engineering at Clemson University, SC, USA. His research interests include cyber physical system security, quantum computing and Artificial intelligence security.



Jean-Michel Tine (Student Member, IEEE) is a Ph.D. student at Clemson University, Clemson, South Carolina, USA. He received his M.S. in Civil Engineering from Clemson University from Clemson University in 2025 and a B.S. in Computer Science from Benedict College, Columbia, South Carolina, USA, in 2021.

Currently, he is working as a Cybersecurity Coordinator for the USDOT-supported National Center for Transportation Cybersecurity and Resiliency (TraCR) at Greenville, SC, USA. His research interests include transportation cyber-Physical systems (TCPS), cybersecurity and resiliency of intelligent transportation systems (ITS), connected and autonomous vehicle technologies, and advanced, high-performance computing for ITS applications.



Latifur Khan (Fellow, IEEE) is currently a professor (tenured) in the Computer Science department at the University of Texas at Dallas, USA where he has been teaching and conducting research since September 2000. He received his Ph.D. degree in Computer Science from the University of Southern California (USC) in August of 2000.

Dr. Khan is a fellow of IEEE, AAAS, IET, BCS, and an ACM Distinguished Scientist. He has received prestigious awards, including the IEEE Technical Achievement Award for Intelligence and Security Informatics, IEEE Big Data Security Award, and IBM Faculty Award (research) 2016. Dr. Khan has published over 300 papers in premier journals and prestigious conferences. Currently, Dr. Khan's research focuses on big data management and analytics, data mining and its application to cybersecurity, and complex data management, including geospatial data and multimedia data. His research has been supported by grants from NSF, DoT, NIH, the Air Force Office of Scientific Research (AFOSR), DOE, NSA, IBM, and HPE. More details can be found at www.utdallas.edu/~lkhan.



Mizanur Rahman (Senior Member, IEEE) received the M.Sc. and Ph.D. degrees in civil engineering (transportation systems) from Clemson University, Clemson, SC, USA, in 2013 and 2018, respectively. He is an assistant professor in the Department of Civil, Construction, and Environmental Engineering at the University of Alabama (UA), Tuscaloosa, AL, USA. His research focuses on cybersecurity and digital twins.