THESIS

TOWARD ROBUST EMBEDDED NETWORKS IN HEAVY VEHICLES - MACHINE LEARNING STRATEGIES FOR FAULT TOLERANCE

Submitted by

Chandrima Ghatak

Department of Computer Science

In partial fulfillment of the requirements

For the Degree of Master of Science

Colorado State University

Fort Collins, Colorado

Summer 2024

Master's Committee:

Advisor: Indrakshi Ray

Yashwant Malaiya
Piotr Kokoszka

ABSTRACT

TOWARD ROBUST EMBEDDED NETWORKS IN HEAVY VEHICLES - MACHINE LEARNING STRATEGIES FOR FAULT TOLERANCE

In the domain of critical infrastructure, Medium and Heavy Duty (MHD) vehicles play an integral role in both military and civilian operations. These vehicles are essential for the efficiency and reliability of modern logistics. The operations of modern MHD vehicles are heavily automated through embedded computers called Electronic Control Units (ECUs). These ECUs utilize arrays of sensors to control and optimize various vehicle functions and are critical to maintaining operational effectiveness. In most MHD vehicles, this sensor data is predominantly communicated using the Society of Automotive Engineering's (SAE) J1939 Protocol over Controller Area Networks (CAN) and is vital for the smooth functioning of MHD vehicles. The resilience of these communication networks is especially crucial in adversarial environments where sensor systems are susceptible to disruptions caused by physical (kinetic) or cyber-attacks.

This dissertation presents an innovative approach using predictive machine learning algorithms to forecast accurate sensor readings in scenarios where sensor systems become compromised. The study focuses on the SAE J1939 networks in MHD vehicles, utilizing real-world data from a Class 6 Kenworth T270 truck. Three distinct machine-learning methods are explored and evaluated for their effectiveness in predicting missing sensor data. The results demonstrate that these models can nearly accurately predict sensor data, which is essential in preventing the vehicle from entering engine protection or limp modes, thereby extending operational capacity under adverse conditions.

Overall, this research highlights the potential of machine learning in enhancing the resilience of networked cyber-physical systems, particularly in MHD vehicles. It underscores the significance of predictive algorithms in maintaining operational feasibility and contributes to the broader discussion on the resilience of critical infrastructure in hostile settings.

ACKNOWLEDGEMENTS

I extend my deepest gratitude to my advisor, Dr. Indrakshi Ray, whose guidance, patience, and invaluable advice have been the cornerstone of my academic journey. Special thanks go to my committee members, Yashwant Malaiya and Piotr Kokoszka, for their insightful feedback and invaluable contributions to my study. I owe a debt of gratitude to Hossein Shirazi and Saira Jabeen, whose assistance in the foundational study and experiments with machine learning was indispensable. Their collaboration has been a pivotal part of my academic growth. I would like to acknowledge all my educators, from my kindergarten teachers to my graduate tutors, who have been integral to my learning journey. Each one of them has played a significant role in shaping the knowledge and skills I possess today.

I am also profoundly grateful to Dr. Jeremy Daily, whose support was instrumental in funding my studies and providing me with an in-depth understanding of heavy vehicle networks.

I extend my gratitude to NSF for their support in partially funding my research under the Award Number ATD 2123761, CNS 1822118, ARL, Statnett, AMI, NewPush, and Cyber Risk Research.

A heartfelt thank you to my late grandparents, Dibyendu and Rama Ghatak. Their teachings and values have been a guiding light in my life, helping me grow into the person I am today.

To my parents, Jaydip and Debamita Ghatak, words cannot express my gratitude for your unwavering support and belief in me. Your love, encouragement, and sacrifices have been the foundation of all my achievements.

Lastly, to my husband, Rik Chatterjee, your love, support, and steadfast presence have been my greatest strength. Your unwavering belief in me and your unconditional love have been the bedrock of my success.

This journey would not have been possible without the collective support, encouragement, and inspiration provided by each one of you. I am forever grateful.

DEDICATION

This thesis is dedicated to my beloved husband, for his constant love and support...

TABLE OF CONTENTS

ABSTRACT ii
ACKNOWLEDGEMENTS iii
DEDICATION iv
LIST OF TABLES vii
LIST OF FIGURES viii

Chapter 1 Introduction 1

Chapter 2 Background on Heavy Vehicle Networks and the SAE J1939 Protocol 4
2.0.1 ISO/OSI Protocol Stack and SAE J1939 4
2.0.2 SAE J1939 Protocol Characteristics 5
2.0.3 SAE J1939 Frames 7
2.0.4 Implications for Sensor Data Communication 9

Chapter 3 Literature Review 10
3.0.1 Recent Developments 10
3.0.2 Limitations of Existing Approaches 11

Chapter 4 Proposed Approach 13
4.0.1 Our Unified Neural Network Approach 13
4.0.2 Innovation in Our Approach 16

Chapter 5 Data Analysis and Pre-processing 18
5.0.1 Data Collection and Initial Processing 18
5.0.2 Sensor Selection and Data Sampling 18
5.0.3 Training Strategies and Error Quantification 19
5.0.4 Dataset Structure and Model Training 20

Chapter 6 Experiments and Results 22
6.0.1 Multiple SPNs Missing for One-step in Future Predictions 22
6.0.2 One SPN Missing for Multiple Steps in Future Predictions 26

Chapter 7 Discussion 29
7.0.1 Experiment 1: Handling Missing SPNs 29
7.0.2 Experiment 2: Long-term Forecasting 29
7.0.3 Model-Specific Performance Analysis 29
7.0.4 Collective Insights and Implications 30

Chapter 8 Conclusion 31
8.0.1 Key Findings and Implications 31
8.0.2 Advancing Cyber-Physical Systems Resilience 31
8.0.3 Future Directions and Deployment Potential 32

8.0.4 Concluding Remarks 32

Chapter 9 Future Work and Research Enhancement 33
9.0.1 Deployment in Embedded Systems 33
9.0.2 Real-world Testing and Evaluation 33
9.0.3 Integration and Feedback Mechanism 34
9.0.4 Comprehensive Validation and Refinement 34
9.0.5 Conclusion 34

Bibliography 35

LIST OF TABLES

5.1 Mean Squared Error on Test Data 18

6.1 Average Errors for n-steps in future prediction (in %) for One-shot Method 26
6.2 Average Errors for n-steps in future prediction (in %) for Recursive Feeding Method . 26

LIST OF FIGURES

1.1 Image depicting Dependency on a Sensor Value to execute other functions in the Vehicle 2
1.2 Kenworth T270 Research Truck that broke down on I80 near Ogallala, NE 3

2.1 Illustration of SAE J1939 Protocol Structure in Heavy Vehicle Networks 7

6.1 Error comparison at 4 different values of k, where k represents the number of missing SPNs 23
6.2 Error comparison between proposed approaches when correlated SPNs are missing . . 24
6.3 Average prediction error of DBT per SPN 24

Chapter 1

Introduction

Medium and Heavy Duty (MHD) vehicles, crucial to modern infrastructure, rely heavily on sophisticated sensor systems controlled by Electronic Control Units (ECUs) for performing critical functions. In the United States, the predominant standard for ECU communication is the Society of Automotive Engineering's (SAE) J1939 protocol [1], which is based on the Controller Area Networks (CAN) [2]. Despite their notable fault tolerance, these protocols are less effective in mitigating risks during sensor malfunctions or cyber-attacks, which can lead to significant operational disruptions. The challenge is amplified by the interdependency of sensors; if one sensor (Sensor A) malfunctions, another sensor (Sensor B) relying on Sensor A's data might also malfunction, as illustrated in Figure 1.1. This cascade effect highlights the need for more resilient solutions.

Traditional fault-tolerance mechanisms in MHD vehicles often fall short in adaptability and real-time responsiveness, requiring manual intervention and updates. This is particularly problematic in the face of evolving sensor failures and emerging cyber threats. Machine learning offers a dynamic and automated alternative. By employing machine learning models, we can recognize complex sensor data patterns, predict future values, and detect anomalies, thereby enhancing the resilience of these vehicles. Our research explores whether a single, generalized machine learning model can provide real-time, robust predictions for compromised or missing sensor values in MHD vehicles, aiming to develop a framework that simplifies system complexity and broadens applicability compared to designing separate models for each sensor type.

A real-world incident in June 2023, involving our 2014 Kenworth T270 research truck, further underscores the practical importance of our research. Returning from the CyberTruck Challenge [3], the truck broke down on US Interstate 80 near Ogallala, Nebraska due to a fuel injector malfunction, as seen in Figure 1.2. This breakdown, though caused by a single malfunctioning component, led to a complete halt, necessitating a tow back to the Colorado State University garage for repairs. Had our machine learning model been implemented, it is plausible that the

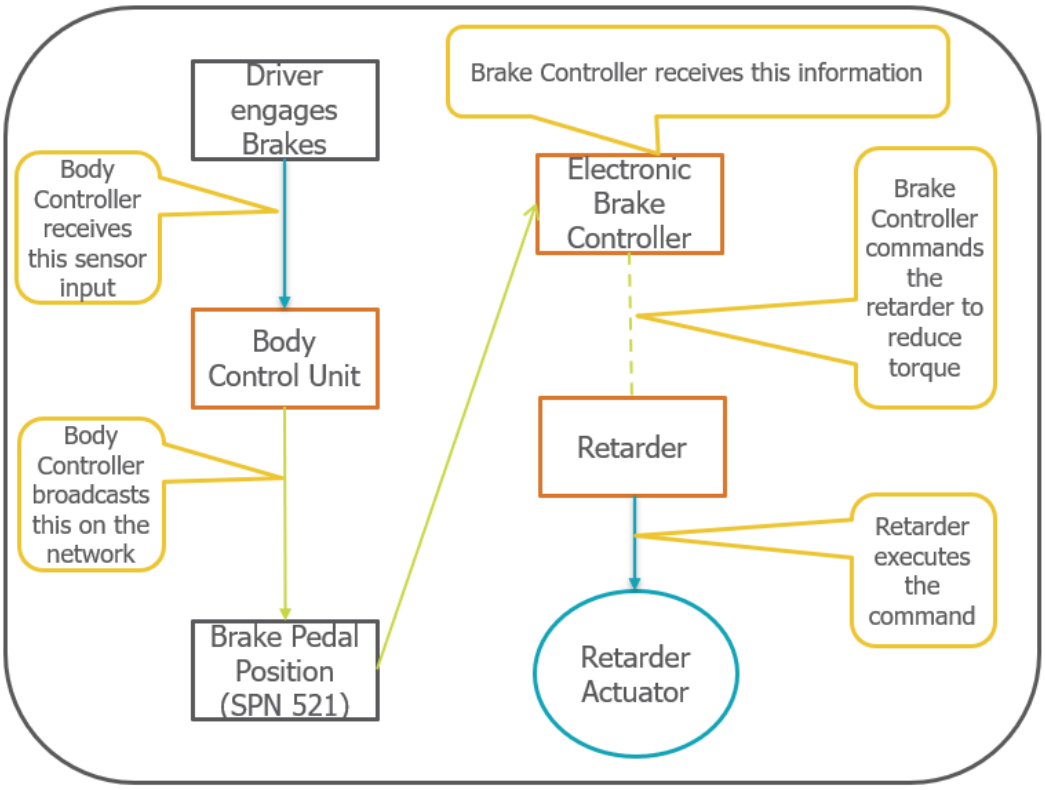

Figure 1.1: Image depicting Dependency on a Sensor Value to execute other functions in the Vehicle

truck could have continued operating, reaching a repair shop safely without the need for a tow. This incident vividly demonstrates the real-world impact and benefits of our research in preventing such breakdowns and enhancing vehicle resilience.

In validating our approach, we experimented with three machine learning algorithms: Dense Binary Transformer (DBT) [4, 5], Sparse Binary Transformer (SBT) [6], and Long Short-Term Memory (LSTM) [7]. Each algorithm was chosen for its strengths in processing time-series data, computational efficiency, and accuracy. Our initial findings indicate that the DBT model performs exceptionally well in various scenarios, with LSTM and SBT also showing considerable effectiveness.

This study contributes to vehicular technology by demonstrating the potential of a robust machine-learning framework in significantly enhancing the resilience of sensor systems in MHD vehicles. This advancement is crucial in counteracting mechanical and cyber vulnerabilities within the rapidly evolving digital landscape of modern vehicular technology.

Figure 1.2: Kenworth T270 Research Truck that broke down on I80 near Ogallala, NE

Chapter 2

Background on Heavy Vehicle Networks and the SAE J1939 Protocol

Modern MHD vehicles have evolved significantly from being purely mechanically driven machines. Their operations heavily rely on multiple embedded computers called Electronic Control Units (ECUs). ECUs are sophisticated embedded computers that manage various vehicle functions, enhancing performance, safety, and efficiency. They process data from numerous sensors to monitor and control systems such as engine performance, fuel efficiency, braking, and stability. This automation not only improves the reliability and precision of MHD vehicles but also reduces the need for manual intervention, leading to more streamlined and effective operations in critical infrastructure. These added functionalities are driven by inter-ECU communication. The Controller Area Network (CAN) [2] is the standard choice of network communication between ECUs. In the United States and growing all over the world, the SAE J1939 protocol [1], built on top of the physical layer CAN specifications, is common industry practice for managing communication amongst Electronic Control Units (ECUs). This protocol follows a layered approach, crafted based on the ISO/OSI model, ensuring a structured approach to data transmission and system interoperability.

2.0.1 ISO/OSI Protocol Stack and SAE J1939

The SAE J1939 protocol leverages four layers of the ISO/OSI stack: the application, network, data link, and physical layers. This stratification allows for a clear delineation of responsibilities and processes across the communication framework.

- The **application layer** handles the end-to-end management of the data, defining the messages and tasks for the system's applications. The application layer is responsible for defining the specific data exchanged between electronic control units (ECUs) within a vehicle. It

specifies the format, size, and meaning of messages and the data within these messages. Additionally, it manages the prioritization of messages to ensure critical data is transmitted first, which is essential in systems with limited bandwidth. It supports diagnostic functions and configuration parameters that allow for troubleshooting, maintenance, and customization of vehicle systems.

- The **network layer** facilitates the switching and routing of data, allowing messages to traverse the network infrastructure. It manages the segmentation and reassembly of messages. The network layer implements the mechanics of address management, ensuring each ECU has a unique address on the network. This includes handling address conflicts and managing dynamic address assignments. It is also responsible for detecting and managing transmission errors, ensuring data integrity through error checking and retransmission mechanisms.

- The **data link layer** assures error-free transmission between network nodes through framing and flow control. The data link layer formats the data into frames suitable for transmission. This includes adding necessary headers and trailers that contain control information. The data link layer manages the acknowledgment of received frames to confirm successful transmission and reception, ensuring reliable communication. It regulates the flow of data to prevent overwhelming the receiver, maintaining smooth and efficient communication.

- The **physical layer** governs the electrical and physical specifications for the devices and media to carry the bits across. This is done using the CAN protocol specifications. The data is encoded using the Non-Return-to-Zero (NRZ) encoding scheme, which ensures a high level of signal integrity and efficient use of bandwidth. The physical layer also specifies the connectors, voltage levels, and timing parameters, ensuring compatibility and reliable communication between different components of the network.

2.0.2 SAE J1939 Protocol Characteristics

The SAE J1939 protocol is a higher-layer protocol that operates on the Controller Area Network (CAN) bus, widely used in heavy-duty and commercial vehicles for in-vehicle networking. It

facilitates communication between Electronic Control Units (ECUs) and other devices, providing a standardized method for data exchange.

Key characteristics of the SAE J1939 protocol include:

- **Use of Protocol Data Units (PDUs):** The protocol employs Protocol Data Units (PDUs) for message transmission. PDUs are structured data packets that encapsulate the necessary information for communication between networked devices. Each PDU contains a 29-bit identifier and a data field, ensuring detailed and organized message handling.

- **29-bit Identifier:** The identifier in each PDU is 29 bits long, which is significantly larger than the 11-bit identifier used in standard CAN protocols. This extended identifier allows for more detailed message categorization and prioritization, supporting complex networking needs.

- **Parameter Group Numbers (PGNs):** Each message in the SAE J1939 protocol is associated with a Parameter Group Number (PGN), which indicates the type and purpose of the data being transmitted. PGNs enable efficient message routing and interpretation by defining standardized data types and their respective functions.

- **Suspect Parameter Numbers (SPNs):** Within the data field of a PDU, individual pieces of data are identified using Suspect Parameter Numbers (SPNs). SPNs represent specific sensor readings or control parameters, allowing for precise monitoring and control of vehicle systems.

- **Network Management:** The SAE J1939 protocol includes mechanisms for network management, such as address claiming and message prioritization. These features ensure that each device has a unique address and that critical messages are transmitted with higher priority, maintaining the integrity and efficiency of the network.

- **Diagnostic Services:** The protocol supports comprehensive diagnostic services, enabling the detection, reporting, and resolution of faults within the vehicle network. This is crucial for maintenance and troubleshooting, ensuring the reliability and safety of vehicle operations.
- **Scalability and Flexibility:** SAE J1939 is designed to be scalable and flexible, accommodating a wide range of applications from simple sensor data transmission to complex control commands. This makes it suitable for various types of vehicles and operating environments.
- **Robustness and Reliability:** Built on the robust CAN bus, the SAE J1939 protocol inherits its reliability and fault tolerance. This ensures stable communication even in harsh automotive environments, characterized by electrical noise and extreme temperatures.

These characteristics make the SAE J1939 protocol a robust and versatile solution for in-vehicle networking, enabling efficient and reliable communication between various vehicle systems and components.

2.0.3 SAE J1939 Frames

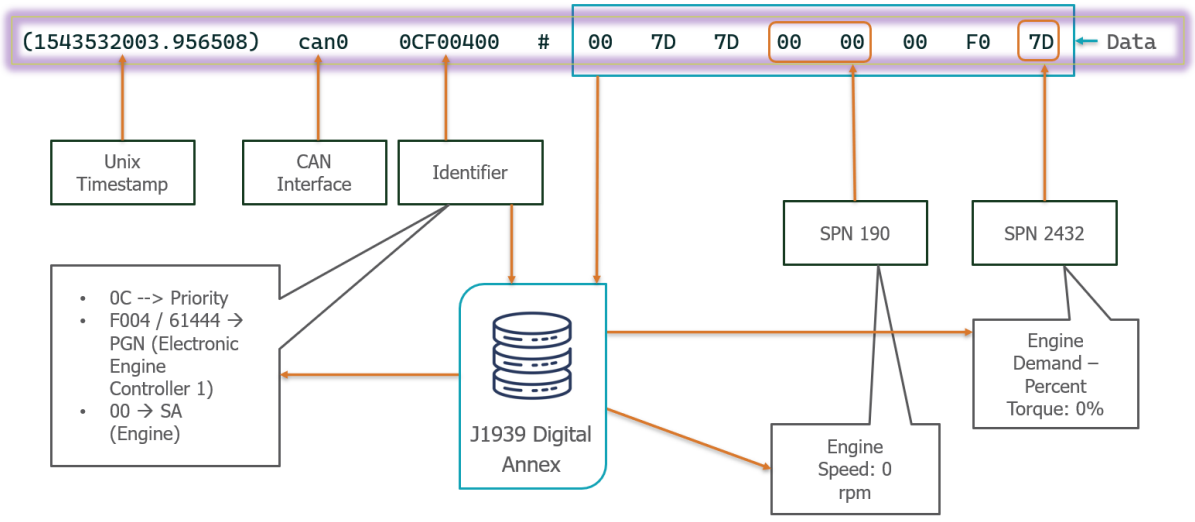

Figure 2.1: Illustration of SAE J1939 Protocol Structure in Heavy Vehicle Networks

The SAE J1939 standard defines the communication frames used in the networks of heavy vehicles. Each frame, fundamental to the protocol, is composed of two main parts: an Identifier and a variable length Data Field as can be seen in Figure 2.1 captured using a Linux socket.

- **Identifier Field:**
  - **Priority:**The 3-bit priority field is used as a base for the CAN arbitration scheme. Priorities can vary from 000_2 (0) to 111_2 (7). The J1939 standard assigns a default priority of 011_2 (3) to vehicle control messages and 110_2 (6) to all other messages. The priority is ultimately specified by the original equipment manufacturer (OEM).
  - **Extended Data Page (EDP):**The Extended Data Page (EDP) bit is used to differentiate between the standard data page and an extended data page. This bit is critical for expanding the number of available parameter groups, thus allowing for greater flexibility and categorization in message management. When the EDP bit is set to 0, it indicates a standard data page, while a value of 1 indicates an extended data page.
  - **Data Page (DP):**The Data Page (DP) bit specifies which data page the message belongs to, facilitating the categorization of parameter groups. A value of 0 indicates the default data page, while a value of 1 indicates the alternate data page, thereby enabling the network to handle a broader set of parameter groups.
  - **PDU Format and PDU Specific:**
    - **PDU Format (PF):** The PDU Format field is a byte that defines the format of the Protocol Data Unit (PDU). It indicates whether the message is intended for a specific address (destination-specific) or a broadcast (global). If the PF is greater than 240_{10}, it specifies the message is broadcast, if it is less than 240_{10}, the message is to a specific address.
    - **PDU Specific (PS):** The PDU Specific field is another byte in length and works in conjunction with the PDU Format to provide additional addressing information.

When the PDU Format indicates a broadcast, the PDU Specific field identifies the particular parameter group, otherwise, it specifies the destination address.

  - **Source Address (SA):**The Source Address (SA) field is a byte and uniquely identifies the sender of the message within the network. Each device on the network must have a unique source address to ensure proper message routing and identification. This 8-bit field provides up to 256 unique addresses, accommodating a wide range of devices within the network.

- **Data Field:**The Data Field contains the actual payload of the message, encapsulating the specific information or command being communicated. The structure and content of the Data Field vary depending on the type of message and its intended function within the vehicle's network system. This field can contain up to 8 bytes of data, providing sufficient space for a wide variety of parameters and commands. The data in the data field contains parameters in terms of encoded sensor data or commands in chunks called Suspect Parameter Numbers (SPNs). These are defined in the SAE J1939 protocol specifications. Figure 2.1, shows some of the decoding of SPNs to obtain relevant sensor data.

2.0.4 Implications for Sensor Data Communication

The efficiency and standardization afforded by the SAE J1939 protocol are critical for heavy vehicle network operations, where timely and reliable sensor data communication is essential. Through the use of PDUs and the layered structure of the ISO/OSI stack, the protocol supports the seamless exchange of operational data, a foundation upon which our machine-learning models can predict and ensure vehicular functionality and safety.

Chapter 3

Literature Review

In this section, we will talk about the recent work in this field and some of the limitations of existing approaches.

3.0.1 Recent Developments

Recent advancements in machine learning have heralded a new era in sensor data reconstruction, particularly when the original data has been compromised by mechanical failure or cyber-attacks. Extensive security analyses of the CAN protocol, commonly employed in passenger vehicles, have been conducted [8]. It has been shown that injecting messages into a vehicle's internal network can significantly impede its operations [9]. The security landscape for vehicular networks took on a new level of complexity with the revelation that remote access could be used to exploit communication networks [10] [11]. Additionally, attacks on sensor systems have also increased over the years [12–14].

Initial inquiries into the cyber-security of medium and heavy vehicles revealed that they are not immune to the attack vectors prevalent in their lighter counterparts [15,16]. The accessibility of the SAE J1939 specification to potential attackers compounds the risk, as it allows for the exploitation of protocol flaws [17–21]. Attacks on application, network, and data link layers of the SAE J1939 have been documented, highlighting the need for robust defense mechanisms. The potential for an attacker to manipulate physical signals on the CAN bus after compromising a number of ECUs has been a recent alarming development [22], along with exploits to commercial vehicle Electronic Logging Devices and the first demonstration of a Truck-to-truck worm [23].

While no cyber-attacks on medium and heavy vehicles have been reported, the vulnerabilities uncovered by research underscore the urgency of bolstering security measures. Proposed countermeasures have traditionally included intrusion detection and prevention systems, as well as cryp-

tographic methods [15, 16]. However, the latter may be too resource-intensive for the constraints of vehicular networks [24].

Machine learning has been posited as a viable alternative, with systems designed to predict erroneous sensor values. [25] suggest a database of normal and attack-simulated inputs to train classifiers for anomaly detection. In a similar vein, multiple neural networks—one for each sensor—have been used to predict correct sensor values, assuming all other sensor values are accurate [26]. LSTM autoencoders have shown potential in reconstructing compromised sensor data, leveraging legitimate sensor values to predict the correct values [26].

These findings are promising, yet they do not account for scenarios where the system lacks foreknowledge of which sensor might fail. Our current research addresses this gap by employing a single neural network model that can accurately predict any sensor value, regardless of which sensors might fail or be compromised. This approach simplifies the prediction process and is better suited to the real-world constraints of heavy vehicle networks. Our goal is to offer a practical solution for reconstructing correct sensor values, to replace those that are compromised or missing, thereby enhancing the safety and resilience of medium and heavy vehicles.

3.0.2 Limitations of Existing Approaches

In the landscape of machine learning applications within vehicular networks, previous research endeavors, such as those by Shirazi et al. [26], have often employed multiple neural networks to predict accurate sensor values. Each neural network was tasked with reconstructing the value for a specific sensor, predicated on the assumption that all other sensor values are available and accurate. While this methodology may be effective under certain conditions, it introduces substantial complexity to the system architecture. Specifically, it necessitates the deployment and management of multiple neural networks, which increases the resource overhead and may not be viable in computational environments where resources are constrained.

Additionally, these approaches presuppose the existence of only a single compromised Suspect Parameter Number (SPN) within the network. This implies that for each neural network, a

combination of uncompromised sensor values is employed to predict the value of the compromised sensor. The underlying challenge arises when it becomes necessary to first identify which sensor data is compromised before initiating the prediction process, thereby complicating the operational workflow. In a real-world scenario, where quick adaptation to sensor failures is critical, this approach may not be efficient. To address these limitations, we advocate for a streamlined methodology that leverages a single, well-trained neural network capable of predicting all sensor values within the network.

Chapter 4

Proposed Approach

In this section, we will discuss the limitations of the existing approach and how our approach improves upon that to provide a more feasible solution.

4.0.1 Our Unified Neural Network Approach

We propose a unified approach that employs a single generalized neural network model, which is capable of predicting any sensor value within the network, regardless of whether it is compromised or missing. Our approach is designed to provide several key advantages over existing methods:

- **Generalization:** The proposed generalized neural network is trained on a comprehensive dataset encompassing a wide array of sensor data configurations. This broad training base enables the model to be sensor-agnostic, meaning it is not limited to predicting values for a predefined set of sensors. Consequently, this eliminates the necessity for multiple sensor-specific networks and provides the flexibility to predict any sensor value within the network.

- **Real-time Prediction:** The ability to deliver predictions in real time is paramount for vehicular networks, where sensor data is continuously generated and consumed by various vehicular subsystems. Our model ensures that sensor dependencies are maintained, allowing other systems that rely on the compromised or missing sensor data to operate without interruption.

- **Reduced System Complexity:** Simplifying the system architecture to a single neural network model reduces the number of potential failure points and decreases the computational overhead associated with managing multiple models. This streamlined approach not only improves system reliability but also optimizes resource utilization.

The proposed model is trained to discern the intricate relationships and patterns among various sensor inputs, facilitating the prediction of sensor values with high accuracy. The training process involves exposing the network to diverse scenarios, both normal and anomalous, to ensure that the model can generalize its predictive capabilities to any sensor data, irrespective of the presence or absence of specific sensors.

In our research, we explore a suite of machine learning algorithms that are theoretically sound and demonstrate promising results in sensor data prediction tasks:

1. Long Short-term Memory (LSTM) [7]
2. Dense Binary Transformer (DBT) [4, 5]
3. Sparse Binary Transformer (SBT) [6]

Dense Binary Transformer (DBT)

The Dense Binary Transformer (DBT) serves as an advanced variant of the well-established Transformer model, which is an architecture originally designed for natural language processing tasks. The Transformer architecture consists of two primary components: an encoder and a decoder. These are responsible for transforming the input data into a format that can be utilized for tasks like classification or prediction. A noteworthy feature of the Transformer model is the self-attention mechanism, which allows the model to weigh different parts of the input data based on their relevance to the task at hand. DBT employs a specialized form of the self-attention mechanism known as ProbSparse Self-Attention. Unlike traditional self-attention, which involves each input unit (commonly referred to as a 'key') interacting with all other units (referred to as 'queries'), the ProbSparse mechanism limits these interactions. Specifically, each key is permitted to interact only with a subset of queries, effectively reducing computational time and memory usage. This is particularly advantageous when the model has to handle large and complex datasets. In the context of machine learning, an encoder is a component of a neural network responsible for transforming raw input data into a condensed, machine-interpretable form. DBT optimizes the conventional Transformer encoder for increased efficiency in processing long sequences of data.

To achieve this, it integrates Convolutional 1D layers—which are generally employed for spatial feature extraction in image data—and MaxPooling layers that reduce data dimensions while preserving essential features. These layers work in tandem with traditional self-attention blocks to extract vital features from large datasets.DBT is specially engineered to capture intricate patterns in large and complex data structures, making it highly effective for tasks requiring predictive accuracy. In particular, it excels in real-time sensor value prediction. The architecture of DBT is tailored to optimize both computational efficiency and the capability to recognize complex patterns, positioning it as an ideal choice for performance-critical applications.

Sparse Binary Transformer (SBT)

The SBT algorithm builds on the principles of the Transformer model, focusing on computational efficiency and scalability. By employing a specialized self-attention mechanism that reduces computational demands, SBT is designed for real-time applications that necessitate immediate sensor value predictions, such as those found in vehicular networks.The Sparse Binary Transformer (SBT) is another derivative of the Transformer model, which, similar to DBT, addresses specific challenges in computational efficiency and scalability. As discussed under DBT, the Transformer architecture is bifurcated into an encoder and a decoder, with self-attention mechanisms playing a crucial role in data transformation and task-specific learning. SBT distinguishes itself by adopting an even more computationally efficient self-attention mechanism. The mechanism is optimized to reduce the amount of computational resources required for processing, making it suitable for applications that need real-time response, such as sensor value prediction. The encoder in SBT is engineered to minimize memory usage and computational time. It employs a selection of specialized layers designed to reduce the dimensionality of the input data effectively while maintaining essential features. The layers include variants of traditional self-attention blocks but are optimized to be computationally less demanding. The SBT model is ideal for scenarios that require real-time sensor value prediction. Its architecture focuses on achieving high predictive performance while being computationally efficient, a crucial factor for ensuring optimal system performance in real-time applications.

Long Short-term Memory (LSTM)

Long Short-Term Memory (LSTM) is a type of Recurrent Neural Network (RNN) architecture, specifically designed for sequence prediction problems. RNNs are neural networks where connections between nodes form a directed graph along a temporal sequence. This allows them to maintain a 'memory' of previous inputs, making them well-suited for tasks involving sequential data, such as time series prediction. LSTM networks include memory cells that allow them to store and recall information over long sequences effectively. Unlike standard RNNs, which often suffer from the 'vanishing gradient' problem, LSTMs are capable of learning long-term dependencies in the data. The LSTM architecture described employs a two-stage process—an encoder that processes the input sequence and captures its information in a 'context vector', and a decoder that generates the output sequence based on this context vector. The architecture includes special layers like 'RepeatVector' for replicating the context vector and 'TimeDistributed' for generating the output sequence. LSTM is particularly advantageous for making inferences from sensor data that exhibit temporal correlations. It is adept at capturing long-range dependencies in time-series data, making it a potentially effective algorithm for applications that require high predictive accuracy over extended periods.

Each of these algorithms is selected based on its intrinsic strengths in processing time-series data, computational efficiency, and predictive accuracy. Through an extensive experimental framework, we aim to rigorously evaluate the performance of these models in a variety of operational scenarios, thus empirically validating their applicability and effectiveness in the context of vehicular networks.

4.0.2 Innovation in Our Approach

Our approach represents a significant innovation over existing methods by unifying the prediction process into a single neural network model as we found in [27]. This model is not only trained to handle diverse data configurations but also to accommodate the dynamic nature of vehicular networks, where sensor configurations may vary over time. By adopting a single model, we greatly

simplify the predictive process, enhance the system's adaptability, and ensure that the network can maintain its operational integrity even in the face of sensor anomalies.

Furthermore, we extend our investigation to include the exploration of newer machine learning algorithms that may offer even more robustness and reduced complexity for our application scenario. These algorithms, discussed in Section 4, have shown potential in recent research to be more effective for the challenges at hand. Our research is therefore not only a step forward in predictive modeling for vehicular networks but also an exploration into the future of machine learning applications in this field.

Chapter 5

Data Analysis and Pre-processing

In this section, we will discuss how we prepared our datasets, for training and testing our algorithms.

5.0.1 Data Collection and Initial Processing

Our data collection initiative was centered around a 2014 Kenworth T270 research truck, which undertook a cross-country journey from Fort Collins to Detroit in 2018 [28]. The resulting dataset was initially captured in the 'candump' format, which provided a raw stream of time-stamped ASCII values alongside CAN identifiers and data fields. Utilizing the SAE-J1939 standard, we meticulously extracted and decoded relevant Parameter Group Numbers (PGNs) to translate them into actual engineering values that would be instrumental for further analysis.

5.0.2 Sensor Selection and Data Sampling

From the extensive array of sensor readings available on the CANbus, we identified 52 sensors with dynamic, non-static measurements that significantly contribute to the machine learning model's training and testing phases, ultimately enhancing the model's capacity to predict missing sensor values with high accuracy. It was these 52 Suspect Parameter Numbers (SPNs) that were earmarked for inclusion in our experiments.

Table 5.1: Mean Squared Error on Test Data

Window Size	Training with all data			Training with 5% of missing data		
	DBT	SBT	LSTM	DBT	SBT	LSTM
10	**0.2446**	**0.5901**	**0.4810**	**0.2432**	**0.4969**	**0.4952**
50	0.2767	0.6391	0.5710	0.2661	0.5886	0.6721
100	0.3174	0.6784	0.6312	0.2638	0.6453	0.7611
200	0.4850	0.7537	0.7005	0.4209	0.6892	0.9023

Considering the varied periodicities inherent in CAN message transmissions, we established a sampling strategy that entailed recording the most recent sensor data at consistent intervals of 500 milliseconds. This approach facilitated the creation of a time-series dataset, offering a structured temporal framework to the sensor data. Subsequent to this, we undertook the normalization of the dataset, scaling the sensor values to a uniform range between 0 and 1, a necessary step to standardize the data for the neural network training and testing processes.

5.0.3 Training Strategies and Error Quantification

Our experimentation focused on optimizing the machine learning algorithms' performance by subjecting them to a range of data conditions, particularly the occurrence of missing values—a common and practical challenge within vehicular networks. To evaluate the models' precision, we adopted the Average Percentage Error (err) as our metric of choice, defined by the following equation:

$$err = \frac{1}{n} \sum_{\mathrm{i}=1}^{n} \frac{|e_{\mathrm{i}}|}{z} \times 100, \tag{5.1}$$

where $|e_{\mathrm{i}}|$ is the absolute deviation between the actual sensor reading and the model's prediction, z denotes the full range of the SPN, and n represents the total number of data points within the specific SPN's dataset. This normalized error metric afforded us a consistent basis for comparing the performance across the diverse array of SPNs.

In alignment with our commitment to rigorous testing, we bifurcated our experiments into two distinct training strategies:

Full-data Training Under this regime, we trained our models using a comprehensive dataset inclusive of all sensor values. This strategy established a performance benchmark, delineating the upper limits of the models' capabilities when provided with an unimpaired dataset.

Training with Missing Values To mirror the realities of operational conditions, where sensor data can intermittently be missing or corrupted, we introduced controlled missing values within

our training data. This method was intended to test the models' ability to adapt and maintain predictive accuracy despite incomplete training data.

5.0.4 Dataset Structure and Model Training

The dataset comprised a total of 16,000 instances, which were partitioned into a training set and a testing set with a split of 75% and 25%, respectively. The models, including the Dense Binary Transformer (DBT), Sparse Binary Transformer (SBT), and Long Short-Term Memory (LSTM), were subjected to a training regimen spanning 100 epochs. This process was facilitated by the computational power of an NVIDIA TITAN V graphics card, boasting 12GB of memory.

Table 5.1 delineates the mean squared error (MSE) outcomes obtained from employing the two aforementioned training strategies across various window sizes. The MSE statistics revealed an optimal performance when the window size was configured to 10, thereby suggesting that an increase in window size does not yield proportionate performance gains. Furthermore, the results substantiated the proficiency of models trained on datasets with missing values, showcasing their capability to effectively generalize in the presence of incomplete data.

A randomized search algorithm was utilized to ascertain the optimal hyperparameters for each model. This search encompassed a spectrum of hyperparameters, such as the choice of optimizer, learning rate, batch size, and normalization techniques, within the layers.

The following are the architectural specifications of the best-performing models discovered through this process, with DBT and SBT sharing similar configurations:

- Positional encoding size of 64 units, enabling the model to discern the sequential order of data points.
- Four encoder layers to transform the input data into a higher-level representation.
- Four self-attention heads within each encoder layer, facilitating the model's focus on different parts of the input sequence.

- Hidden layers comprise 128 units each, providing the necessary computational depth for learning complex patterns.

The architecture for the LSTM model was distinct and included:

- An input layer shaped to accommodate the number of past observations and the corresponding features within each.
- Two encoder LSTM layers, each with a specified number of nodes, capable of returning sequences and their respective internal states.
- A `RepeatVector` layer to duplicate the encoder's output, thereby concentrating the decoder's efforts on pivotal elements of the input sequence.
- Two decoder LSTM layers tasked with constructing the output sequence based on the encoder's internal states.
- A `TimeDistributed` wrapper applied to a `Dense` layer, which ensures a dedicated output per time step in the sequence.

This comprehensive data preparation and model optimization framework underscores our commitment to developing a robust and reliable machine-learning solution for real-time sensor value prediction in medium and heavy-duty vehicles.

Chapter 6

Experiments and Results

In this section, we will discuss the different types of experiments we performed on the datasets and the results we achieved from these experiments.

6.0.1 Multiple SPNs Missing for One-step in Future Predictions

In this subsection, we evaluate the performance of different predictive models including LSTM, SBT, and DBT in predicting sensor values under various test cases.

Test Case 1: Missing one or more SPNs

The performance evaluation of LSTM, SBT, and DBT models in this test case is particularly critical in understanding their resilience to data sparsity. The results are a direct reflection of each model's ability to handle incomplete data sets, a common occurrence in sensor networks due to factors like sensor failures or communication errors.

In this scenario, represented in Figure 6.1, we explore situations where k SPNs are missing, along with their historical data within a specified time window. This setup is designed to simulate the practical challenges faced in sensor data prediction when multiple data points are absent. For each value of k, we randomly select $k - 1$ SPNs from a pool of 51, targeting the prediction of the kth SPN. This methodical approach allows us to gauge the models' efficiency in dealing with varying degrees of missing data.

The analysis reveals a consistent pattern where LSTM exhibits the highest prediction error among the three models across all k values. This outcome implies that LSTM's prediction accuracy is significantly reliant on the availability of a larger set of SPN values. In contrast, DBT showcases the lowest error rates, suggesting its superior ability to extract relevant features from a limited data context. Interestingly, we observe that an increase in the number of missing SPNs does not linearly

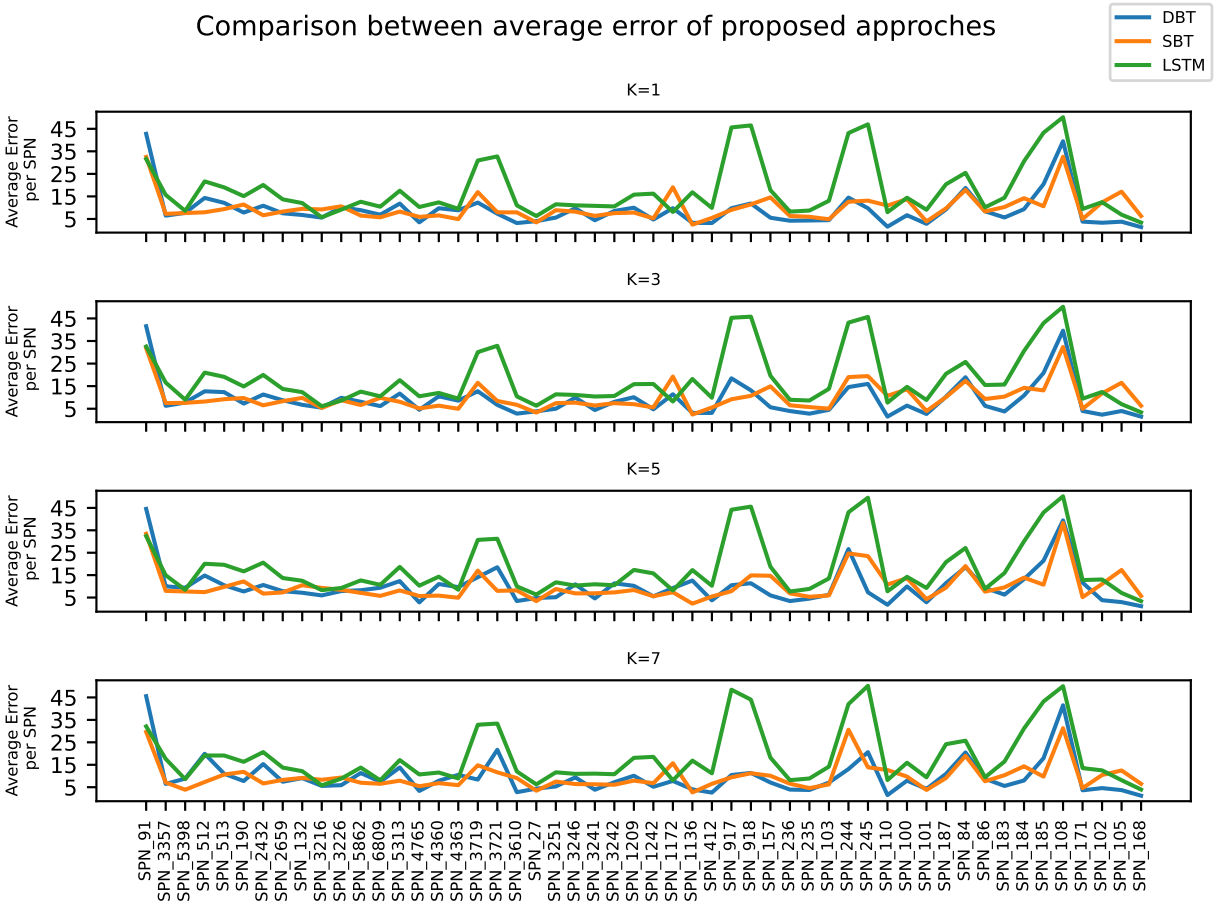

Figure 6.1: Error comparison at 4 different values of k, where k represents the number of missing SPNs

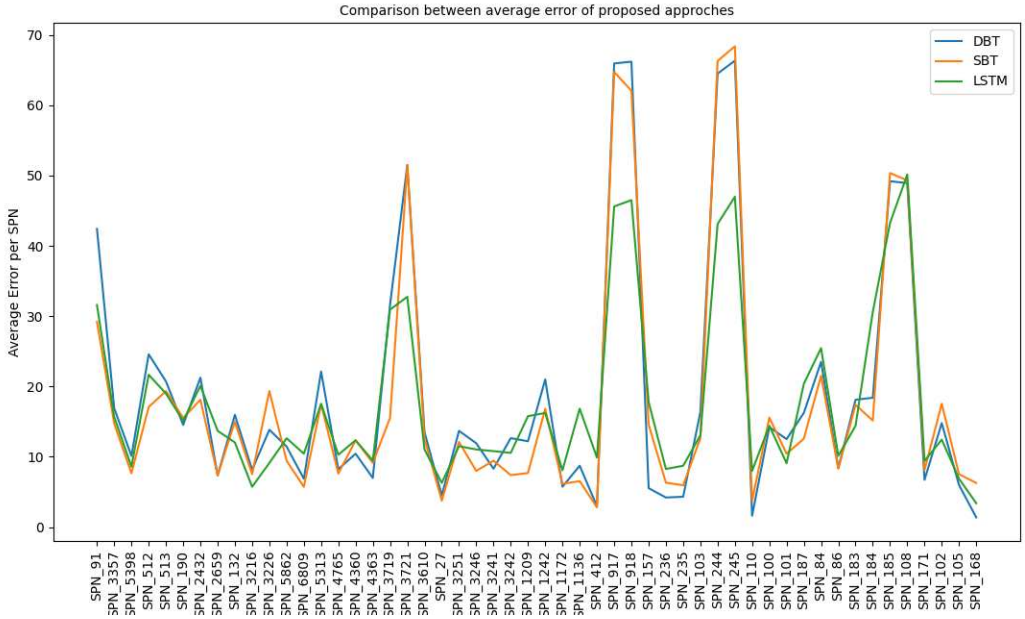

Figure 6.2: Error comparison between proposed approaches when correlated SPNs are missing

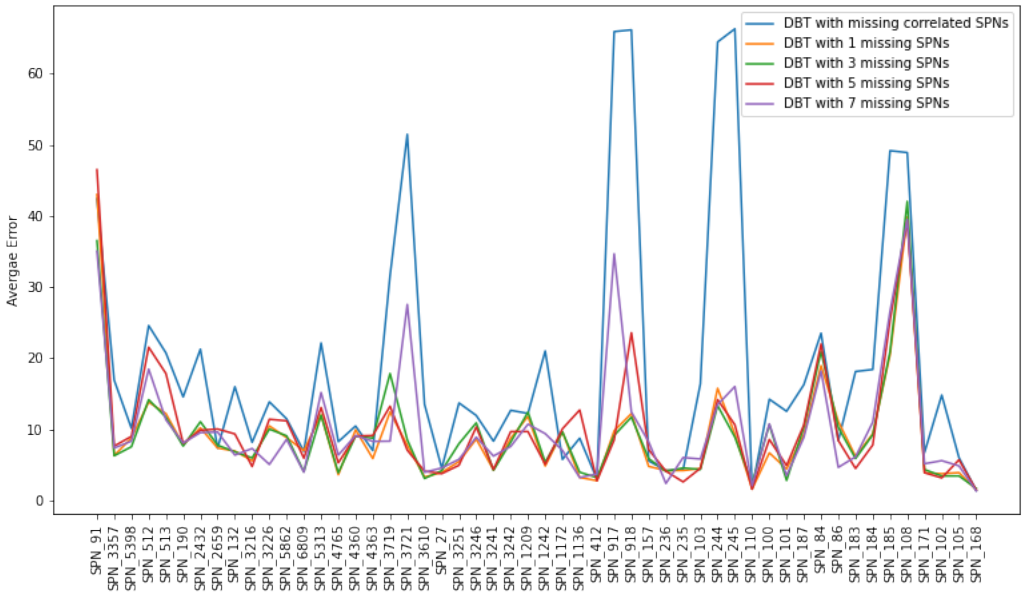

Figure 6.3: Average prediction error of DBT per SPN

escalate the prediction error for most SPNs, indicating a certain level of robustness in the models against data scarcity.

Test Case 2: Missing SPN and its Correlated SPNs

This test case delves into scenarios where not only is a specific SPN missing, but also SPNs that are correlated to it. We identify these correlations using Pearson Correlation, setting a threshold value of 0.5 to discern significant relationships. This approach is crucial in evaluating the models' capabilities in predicting sensor values when inter-sensor dependencies are disrupted.

Figure 6.2 provides a comparative analysis of the LSTM, SBT, and DBT models, each configured with a window size of 10. In this setting, the error rates spike notably when SPNs correlated to the target SPN are missing, offering valuable insights into each model's handling of related sensor data. The transformer models, particularly DBT, demonstrate an adeptness in learning from the interdependencies between SPNs. Conversely, LSTM's lower error rates in this case can be attributed to its focus on temporal sequences over inter-sensor relationships, highlighting its distinct methodological advantage.

Comprehensive Analysis of DBT Performance

To further elucidate the models' performance, Figure 6.3 presents a detailed view of DBT's average prediction error for each SPN under two distinct conditions: one with a varying number of randomly selected missing SPNs (1, 3, 5, 7), and another with missing SPNs that exhibit correlations to the target SPN. This analysis reinforces our previous findings, showing that while the prediction error increases with a higher number of missing SPNs, the increase is not disproportionately large. However, the error rates are more pronounced when correlated SPNs are absent, underscoring the transformer models' proficiency in exploiting inter-sensor dependencies for accurate predictions.

The evaluations presented in this section provide a comprehensive understanding of the capabilities and limitations of LSTM, SBT, and DBT models in scenarios involving missing sensor data. While transformer models excel in situations where understanding sensor interdependencies

Table 6.1: Average Errors for n-steps in future prediction (in %) for One-shot Method

Model	Average Errors for n-steps in future prediction (in %)								
	1	3	5	10	15	20	25	50	100
DBT	4.90	4.45	4.86	5.20	5.41	5.36	7.35	9.43	11.88
SBT	6.90	6.78	6.81	7.31	7.19	7.67	8.62	10.93	14.22
LSTM	5.66	5.45	5.56	5.86	5.93	7.62	9.32	12.93	15.86

Table 6.2: Average Errors for n-steps in future prediction (in %) for Recursive Feeding Method

Model	Average Errors for n-steps in future prediction (in %)								
	1	3	5	10	15	20	25	50	100
DBT	4.90	3.73	3.55	4.01	4.66	4.70	6.04	8.18	10.33
SBT	6.90	5.97	5.40	6.05	5.91	6.27	7.04	9.55	13.17
LSTM	5.66	4.21	4.79	4.34	4.57	6.11	8.13	11.50	14.31

is crucial, LSTM's strength lies in its ability to leverage temporal data patterns. This analysis underlines the importance of algorithm selection based on the specific nature of the sensor data and missing value patterns in sensor networks.

6.0.2 One SPN Missing for Multiple Steps in Future Predictions

To further our understanding of the capabilities of the LSTM, SBT, and DBT algorithms, we conducted a series of experiments aimed at evaluating their performance in predicting multiple future missing sensor values. This investigation is crucial for assessing the long-term predictive accuracy of these models, which is essential in scenarios where continuous sensor data prediction is required. Our experimental setup involved two distinct methods: the One-shot Method and the Recursive Feeding Method. Each of these methods offers a unique approach to future value prediction, and our goal was to understand how each algorithm performs under these different strategies.

Test Case 1: One-shot Method

The One-shot Method focuses on directly predicting the future value of a missing sensor at a specific time step n. This approach requires the models to be trained on the remaining 50 sensor values and output the predicted value of the missing sensor at the n-th future time step.

Table 6.1 presents the results of this method, showcasing the average prediction errors for each model across various future time steps ranging from 1 to 100. The DBT model exhibited remarkable consistency, maintaining the lowest prediction errors for most future steps. This highlights its robustness and predictive accuracy, particularly in scenarios requiring direct future value prediction. In contrast, the LSTM and SBT models demonstrated higher errors, especially for longer-term predictions. The increasing error rates at future steps like 50 and 100 suggest challenges in maintaining accuracy over extended prediction horizons.

Test Case 2: Recursive Feeding Method

In contrast to the One-shot Method, the Recursive Feeding Method involves initially predicting the missing sensor value one step ahead and then recursively using this predicted value, along with the other 50 sensor values, to predict the next step. This process continues until the prediction is made for the desired future time step n.

Table 6.2 details the average prediction errors for each model using the Recursive Feeding Method. Here, DBT again performs exceptionally well, showing lower errors across the range of future time steps. Interestingly, LSTM shows improved performance compared to the One-shot Method, particularly in the short-to-medium range of future steps. The SBT model, while showing higher errors, demonstrates a degree of improvement in long-term predictions.

Discussion and Implications

The findings from these experiments provide valuable insights into the capabilities of each predictive model in handling long-term future predictions. The DBT model's consistent performance across both test cases underscores its potential as a reliable tool for future value prediction in sensor networks. The LSTM model, with its improved performance in the Recursive Feeding Method, highlights its suitability for scenarios where intermediate forecasts are valuable. The SBT model, despite its relatively higher error rates, still offers useful predictive capabilities, particularly for longer prediction horizons.

The choice between the One-shot and Recursive Feeding methods should be informed by the specific requirements of the application, such as the need for computational efficiency, prediction accuracy, and the importance of intermediate forecasts. The results of these experiments thus not only demonstrate the strengths and limitations of each model but also provide guidance for selecting the appropriate prediction strategy based on the application's needs.

Chapter 7

Discussion

In our comprehensive evaluation of machine learning models for predicting sensor data in Medium and Heavy Duty (MHD) vehicles, we conducted two distinct types of experiments. These experiments were meticulously designed to address the multifaceted challenges encountered in real-world vehicular sensor data prediction. The first set of experiments was aimed at assessing the models' resilience to missing sensor data under various conditions, mirroring the uncertainties prevalent in real-world scenarios. The second set of experiments was focused on evaluating the models' capabilities in predicting sensor values over multiple future time steps, thus gauging their effectiveness in long-term forecasting.

7.0.1 Experiment 1: Handling Missing SPNs

The primary objective of this experiment was to test the models' ability to accurately predict sensor data in the presence of missing SPNs, a common occurrence in MHD vehicle networks due to factors like sensor malfunctions or communication disruptions. This experiment provided vital insights into each model's robustness and adaptability to incomplete data scenarios.

7.0.2 Experiment 2: Long-term Forecasting

In contrast, the second experiment was designed to understand the models' predictive performance over extended future time horizons. This is particularly crucial in applications where anticipating sensor values well into the future can significantly aid in proactive decision-making and system management.

7.0.3 Model-Specific Performance Analysis

- **DBT**: Demonstrating remarkable versatility, the Dense Binary Transformer (DBT) model exhibited superior performance in both experiments. It consistently provided high accuracy

predictions, even in scenarios with missing SPNs, and showcased its proficiency in long-term forecasting. These results underline DBT's robustness and efficiency, making it a suitable choice for a wide range of predictive tasks in MHD vehicle sensor networks.

- **LSTM**: The Long Short-Term Memory (LSTM) model excelled in scenarios that required an in-depth understanding of temporal patterns, particularly evident in its performance in the Recursive Feeding Method. This underscores LSTM's strength in capturing and utilizing temporal sequences, a crucial factor in many predictive scenarios, especially where historical sensor data trends play a significant role in forecasting.

- **SBT**: The Sparse Binary Transformer (SBT) model, while generally lagging slightly behind in terms of accuracy, showed notable potential in long-term prediction tasks. It also exhibited improved performance in handling missing SPN values. This suggests that SBT can be an effective option in scenarios where long-term forecasting is more critical than immediate accuracy or where data sparsity is a significant challenge.

7.0.4 Collective Insights and Implications

The findings from these experiments offer a nuanced understanding of each model's specific strengths and limitations. While the DBT model emerges as a robust all-rounder, suitable for a variety of predictive scenarios, both LSTM and SBT have distinct areas where they offer significant advantages. This layered insight is invaluable for selecting the most appropriate model based on the specific requirements and constraints of MHD vehicle sensor data prediction tasks. Understanding these dynamics is crucial for developing effective and efficient predictive systems in the automotive industry.

Chapter 8

Conclusion

This paper has explored the application of advanced machine learning techniques to address the critical challenge of predicting missing sensor data in Medium and Heavy Duty (MHD) vehicles. Our research contributes significantly to the field by demonstrating the effective use of a single neural network model, which simplifies the traditionally complex process of sensor data prediction. This approach marks a departure from previous methodologies that often relied on multiple models, thereby reducing the computational overhead and enhancing the system's efficiency.

8.0.1 Key Findings and Implications

The core findings of our research indicate that the newer algorithms, including DBT, SBT, and LSTM, can predict missing sensor data with high accuracy. These results are pivotal, as they suggest that the deployment of such models in real MHD vehicles could substantially mitigate the risks associated with sensor malfunctions or cyberattacks. In practical terms, this implies that mission-critical MHD vehicles experiencing sensor failures could avoid operational shutdowns or being forced into a limp mode, which are significant concerns in current vehicular systems.

8.0.2 Advancing Cyber-Physical Systems Resilience

The broader implication of our work is its contribution to enhancing the resilience of modern networked cyber-physical systems. By providing a reliable method for compensating for sensor data loss, our approach not only addresses a specific challenge within the realm of MHD vehicles but also sets a precedent for similar applications in other domains where sensor reliability is crucial. This research, therefore, extends beyond the automotive industry, offering potential applications in various sectors that depend on robust sensor networks.

8.0.3 Future Directions and Deployment Potential

Looking ahead, the practical deployment of these models in real-world settings presents an exciting avenue for further research and development. The successful implementation of these algorithms in operational MHD vehicles would not only validate our findings but also mark a significant step forward in the practical application of machine learning in vehicular technology. Such advancements could lead to substantial improvements in the safety, reliability, and efficiency of MHD vehicles, ultimately contributing to the broader goal of enhancing cyber-physical system resilience.

8.0.4 Concluding Remarks

In conclusion, this paper provides a valuable and feasible solution to a pressing problem in the field of vehicular technology. By leveraging modern machine learning algorithms, we offer a practical and efficient approach to enhance the resilience of MHD vehicles, potentially transforming their operational capabilities under adverse conditions. Our research paves the way for future innovations in the field and underscores the critical role of machine learning in advancing the safety and reliability of networked cyber-physical systems.

Chapter 9

Future Work and Research Enhancement

As we progress with our research endeavors, our future work is centered around several pivotal areas aimed at enhancing the practical applicability and effectiveness of our machine learning solutions in real-world scenarios. The primary objective is to transition from theoretical model development to practical, on-field applications, particularly within the context of heavy vehicle networks.

9.0.1 Deployment in Embedded Systems

A critical step in our future work involves the deployment of our trained predictive models into an embedded device integrated with our 2014 Kenworth T270 research truck. This move is designed to transition our research from controlled, simulated environments to real-world operational settings. By embedding the model directly into the truck's systems, we aim to gain firsthand insights into the model's performance under typical operational conditions. This will provide us with a valuable opportunity to evaluate the model's effectiveness and reliability in a live vehicular environment.

9.0.2 Real-world Testing and Evaluation

An essential component of our future work is the comprehensive testing and evaluation of the deployed model in live scenarios. This involves simulating various sensor malfunction situations and monitoring the model's response to these conditions. By artificially inducing sensor failures, we can closely observe how the model predicts missing sensor values and maintains vehicle functionality under these adverse conditions. This will enable us to assess the model's resilience and its ability to adapt to real-world sensor-related challenges.

9.0.3 Integration and Feedback Mechanism

Another significant aspect of our future work is the integration of the model's predicted sensor values back into the vehicle's controller area network. This step is crucial in determining whether our solution can effectively mitigate the impact of missing sensor values on the vehicle's overall operation. By feeding the estimated sensor values back into the network, we aim to test the model's capacity to maintain seamless vehicle operations and contribute to the system's overall resilience.

9.0.4 Comprehensive Validation and Refinement

The culmination of our future work will be a thorough validation and refinement process. This process is aimed at ensuring that our solution not only addresses the theoretical aspects of sensor data prediction but is also robust and reliable enough to handle the complexities and challenges inherent in heavy vehicle sensor networks. Through this extensive validation, we plan to fine-tune our models, making them more suited for deployment in a wide range of vehicular applications.

9.0.5 Conclusion

In summary, our future work plan is a comprehensive roadmap that encompasses the transition of our research from theoretical models to practical, real-world applications. By implementing these steps, we aim to significantly advance the field of sensor data prediction in heavy vehicles, providing solutions that are not only theoretically sound but also practically viable and effective in enhancing vehicular functionality and safety.

Bibliography

[1] SAE International. The SAE J1939 Standards Collection, 2023. Accessed: 2022, Dec 12.

[2] Bosch. Can specification, 2023. Accessed: 2022, Dec 12.

[3] cybertruckchallenge.org. Cybertruck challenge, 2023. Accessed: 2024-01-15.

[4] Shizhan Liu, Hang Yu, Cong Liao, Jianguo Li, Weiyao Lin, Alex X. Liu, and Schahram Dustdar. Pyraformer: Low-complexity pyramidal attention for long-range time series modeling and forecasting. In *International Conference on Learning Representations*, 2022.

[5] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting. *Proc. of AAAI Conf. on AI*, 2021.

[6] Matt Gorbett, Hossein Shirazi, and Indrakshi Ray. Sparse binary transformers for multivariate time series modeling. Proc. of ACM SIGKDD '23, 2023.

[7] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[8] Cesar Bernardini, Muhammad Rizwan Asghar, and Bruno Crispo. Security and privacy in vehicular communications: Challenges and opportunities. *Vehicular Communications*, 10:13–28, 2017.

[9] Marko Wolf, André Weimerskirch, and Christof Paar. *Secure In-Vehicle Communication*, pages 95–109. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.

[10] IJISC. Remote Exploitation of an Unaltered Passenger Vehicle., 2023. Accessed: 2022, Dec 12.

[11] Stephen Checkoway, Damon McCoy, Danny Anderson, Brian Kantor, Hovav Shacham, Stefan Savage, Karl Koscher, Alexei Czeskis, Franziska Roesner, and Tadayoshi Kohno. Comprehensive experimental analyses of automotive attack surfaces. In David Wagner, editor, *Proceedings USENIX Security 2011*. USENIX, August 2011.

[12] Brady W. O'Hanlon, Mark L. Psiaki, Jahshan A. Bhatti, Daniel P. Shepard, and Todd E. Humphreys. Real-time gps spoofing detection via correlation of encrypted signals. *NAVIGATION*, 60(4):267–278, 2013.

[13] Jonathan Petit, Bas Stottelaar, Michael Feiri, and Frank Kargl. Remote attacks on automated vehicles sensors: Experiments on camera and lidar. *Black Hat Europe*, 11(2015):995, 2015.

[14] Tencent KeenLab. Experimental security research of tesla autopilot, 2021. Accessed: 2023-08-31.

[15] United States Department of Transportation. Cybersecurity research considerations for heavy vehicles., 2023. Accessed: 2022, Dec 12.

[16] Marko Wolf and Robert Lambert. Hacking trucks - cybersecurity risks and effective cybersecurity protection for heavy duty vehicles. In Peter Dencker, Herbert Klenk, Hubert B. Keller, and Erhard Plödererder, editors, *Automotive - Safety Security 2017 - Sicherheit und Zuverlässigkeit für automobile Informationstechnik*, pages 45–60. Gesellschaft für Informatik, Bonn, 2017.

[17] Yelizaveta Burakova, Bill Hass, Leif Millar, and André Weimerskirch. Truck hacking: An experimental analysis of the sae j1939 standard. In *Workshop on Offensive Technologies*, 2016.

[18] Pal-Stefan Murvay and Bogdan Groza. Security shortcomings and countermeasures for the sae j1939 commercial vehicle bus protocol. *IEEE Transactions on Vehicular Technology*, 67:4325–4339, 2018.

[19] Subhojeet Mukherjee, Hossein Shirazi, Indrakshi Ray, Jeremy Daily, and Rose Gamble. Practical DoS attacks on embedded networks in commercial vehicles. In *Information Systems Security*. Springer International Publishing, 2016.

[20] Rik Chatterjee, Subhojeet Mukherjee, and Jeremy Daily. Exploiting transport protocol vulnerabilities in sae j1939 networks. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2023.

[21] Rik Chatterjee, Carson Green, and Jeremy Daily. Exploiting diagnostic protocol vulnerabilities on embedded networks in commercial vehicles. In *Symposium on Vehicles Security and Privacy (VehicleSec) 2024*, San Diego, CA, USA, February 2024. www.ndss-symposium.org.

[22] Abdullah Zubair Mohammed, Yanmao Man, Ryan Gerdes, Ming Li, and Z. Berkay Celik. Physical layer data manipulation attacks on the CAN bus. In *International Workshop on Automotive and Autonomous Vehicle Security (AutoSec), collocated with NDSS*, pages 1–5, 2022.

[23] Jake Jepson, Rik Chatterjee, and Jeremy Daily. Commercial vehicle electronic logging device security: Unmasking the risk of truck-to-truck cyber worms. In *Symposium on Vehicles Security and Privacy (VehicleSec) 2024*, San Diego, CA, USA, February 2024. www.ndss-symposium.org.

[24] Emad Aliwa, Omer Rana, Charith Perera, and Peter Burnap. Cyberattacks and countermeasures for in-vehicle networks. *ACM Comput. Surv.*, 54(1), mar 2021.

[25] Hossein Shirazi, Indrakshi Ray, and Charles Anderson. Using machine learning to detect anomalies in embedded networks in heavy vehicles. In *Proc. of 12th International Symposium on Foundations and Practice of Security*, volume 12056. Springer, 2019.

[26] Hossein Shirazi, William Pickard, Indrakshi Ray, and Haonan Wang. Towards resiliency of heavy vehicles through compromised sensor data reconstruction. In *Proceedings of the*

Twelfth ACM Conference on Data and Application Security and Privacy, CODASPY '22, page 276–287, New York, NY, USA, 2022. Association for Computing Machinery.

[27] Chandrima Ghatak, Saira Jabeen, Hossein Shirazi, and Indrakshi Ray. Improving the resiliency of embedded networks in heavy vehicles - towards fault tolerance. In *Proceedings of Ninth Annual Industrial Control System Security (ICSS) Workshop*. Annual Computer Security Applications Conference (ACSAC), 2023.

[28] Systems Engineering, Colorado State University. Heavy Vehicle CAN Data, 2023. "Accessed: 2022, Feb 02".