

Foundation CAN LM: A Pretrained Language Model For Automotive CAN Data

Akiharu Esashi^{1,3}, Pawissanutt Lertponggrujikorn¹, Justin Makino², Yuibi Fujimoto³, Mohsen Amini Salehi¹

¹HPCC Lab, University of North Texas, Denton, Texas, USA

{akiharu.esashi, pawissanutt.lertponggrujikorn, mohsen.aminisalehi}@unt.edu

²Connected Analytic Services, Plano, Texas, USA

justin.makino@connectedanalyticsservices.com

³Toyota Insurance Management Solutions, Plano, Texas, USA

yuibi.fujimoto@toyota-ims.com

Abstract—The Controller Area Network (CAN) bus provides a rich source of vehicular signals increasingly leveraged for applications in automotive and auto insurance domains, including collision detection, predictive maintenance, and driver risk modeling. Despite this potential, existing pipelines largely train isolated task-specific models on raw CAN data, with only limited efforts exploring decoded signals. Such fragmentation prevents shared representation learning and limits cross-task generalization. By contrast, natural language processing (NLP) and computer vision (CV) have been transformed by the foundation model paradigm: large-scale pretraining followed by task-specific adaptation. In this work, we introduce the *foundation CAN model* that demonstrates multi-objective downstream generalization using a single pretrained backbone. Our approach treats CAN data as a language: we pretrain on large-scale, unlabeled decoded CAN signals and fine-tune across heterogeneous auto insurance tasks. To enable this, we propose a unified tokenization scheme for mixed discrete-continuous signals and address challenges of temporal complexity and trip-specific variability. Our results show that one pretrained CAN model can adapt effectively to diverse predictive tasks, validating that the foundation modeling paradigm, proven in NLP and CV, also holds for CAN data. This establishes a new direction for generalizable representation learning in automotive AI.

I. INTRODUCTION

The Controller Area Network (CAN) bus [1] is the backbone of in-vehicle communication, enabling Electronic Control Units (ECUs) to exchange information such as speed, braking signals, and engine status [2]. With proper decoding, CAN data exposes rich signals that are increasingly leveraged by both the automotive and auto insurance industries. Applications span collision detection and avoidance, predictive maintenance, and autonomous driving on the automotive side, as well as driver behavior scoring, point-of-impact detection, and total loss assessment in the auto insurance domain [3]–[5]. Despite this growing interest, current AI/ML applications of CAN data remain fragmented. Most existing pipelines train isolated, task-specific models directly on raw CAN traffic, with a few recent efforts exploring decoded signals [6]–[9]. This fragmented approach prevents shared representation learning, incurs redundant data preparation and training costs, and limits

generalization across tasks (Fig. 1, top). In contrast, fields such as natural language processing (NLP) and computer vision (CV) have been transformed by the emergence of *foundation models* [10], [11]: large-scale pre-training followed by task-specific post-training, yielding general-purpose backbones that can be fine-tuned efficiently for diverse downstream applications [12]. Yet, no comparable framework has demonstrated such generalization for CAN data.

We identify this gap as the central research problem: while prior work has shown that CAN traffic can be modeled with transformer-based pretraining [13] for a *single objective*, there is no evidence of a generalizable pretraining paradigm for CAN signals that transfers effectively across *multiple* automotive tasks. To address this, we propose treating CAN data as a language and applying the Large Language Model (LLM) paradigm (Fig. 1, bottom). In this approach, a single foundation model is pretrained on large-scale, unlabeled decoded CAN data and subsequently adapted via fine-tuning to heterogeneous downstream objectives. This enables shared representation learning, reduces redundancy, and improves cross-task generalization, overcoming the limitations of the current isolated-model paradigm.

Developing a foundation CAN model, however, presents unique challenges beyond those faced in text or vision domains. First, tokenization is non-trivial: unlike text with discrete symbols, CAN signals contain mixed discrete-continuous values, requiring a robust, reproducible scheme. Second, CAN data is inherently a multi-scale time series, with dependencies ranging from millisecond-level sensor dynamics to long-horizon trip-level patterns. Third, each trip introduces unique contextual factors (vehicle state, driver, environment), demanding representations that generalize across trips while capturing within-trip dynamics.

In this work, we make the following contributions:

- 1) **Multi-Objective Foundation CAN Model:** We introduce the pretraining–post-training framework that demonstrates cross-task generalization from a single CAN backbone, enabling diverse automotive tasks to benefit from shared representations.
- 2) **Unified Tokenization for Mixed CAN Signals:** We design a reproducible tokenization scheme tailored to

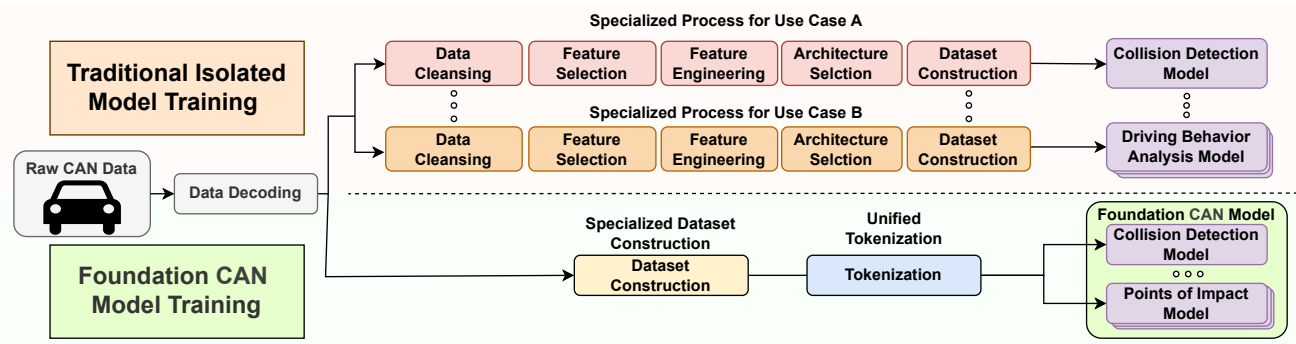


Fig. 1: Traditional task-specific pipelines (**top**) require isolated preprocessing and models for each task. Our foundation model approach (**bottom**) uses unified tokenization and a single pretrained backbone that adapts to multiple downstream tasks through task-specific fine-tuning.

CAN data, integrating scaling and quantization to handle both discrete and continuous signals for large-scale pretraining.

- 3) **Cross-Task Evaluation Study:** We provide a systematic study of a pretrained CAN model across heterogeneous downstream tasks, validating that the foundation modeling paradigm, proven in NLP and CV, also holds for CAN data.

II. RELATED WORK

Task-specific modeling of CAN signals: Most learning pipelines over CAN data train isolated models for a single objective; e.g., intrusion detection system (IDS) [14], driver identification [15], or predictive maintenance [5], directly from raw CAN traffic. While effective within narrow domains, this fragmentation prevents shared representation learning, duplicates training effort across tasks, and limits generalization beyond the target application.

Transformer-based modeling of raw CAN traffic: Several works adapt Transformer architectures and language-modeling objectives to the in-vehicle setting. CAN-BERT [6] employs masked language modeling (MLM) over raw CAN messages for IDS, and Nam *et al.* [7] propose a bi-directional GPT-style model for the same task. Nguyen *et al.* [8] introduce an Attention Network that improves IDS efficiency and shows partial cross-vehicle transfer. While these works validate that raw CAN streams can be modeled with transformers, they remain confined to *single-task* objectives, and do not explore decoded signals or multi-task generalization.

Decoded CAN signal modeling: Beyond raw CAN data, a few recent efforts explore decoded CAN signals [3], [9]. The thesis *CAN you speak Forklift?* [9] demonstrates GPT-style pretraining on decoded industrial-vehicle signals for activity recognition. This work tokenizes forklift CAN signals by extracting state-transition events from *four signals* and models them autoregressively to predict operational sequences. However, the approach remains domain-specific (industrial forklifts vs. consumer vehicles), relies on event-based preprocessing that discards intra-event dynamics, operates on a limited signal set (4 vs. our 44 decoded features), and critically, does not

evaluate cross-task transfer. Our work addresses these limitations by pretraining on a comprehensive set of decoded CAN signals with unified tokenization and demonstrating adaptation across heterogeneous tasks.

Research gap and positioning: The literature demonstrates that (i) transformers can model CAN sequences effectively, (ii) decoded signals enable richer semantic modeling than raw traffic, and (iii) foundation-style pretraining benefits adjacent automotive domains. Yet, to the best of our knowledge, *no prior work has demonstrated that a single pretrained CAN backbone can adapt across multiple heterogeneous automotive tasks*. Existing CAN models target single objectives, and the limited decoded-signal work does not establish cross-task transfer. Our work bridges this gap by (i) pretraining an LM with MLM over a unified mixed discrete–continuous vocabulary spanning 44 decoded automotive sensor signals, and (ii) demonstrating adaptation to heterogeneous downstream tasks. This transition from task-specific modeling to a *CAN foundation model* establishes the viability of multi-objective downstream generalization for automotive and auto insurance applications (aka downstream tasks).

III. METHODOLOGY

Our methodology follows a two-stage paradigm: (1) large-scale pretraining on unlabeled decoded CAN signals using MLM, and (2) task-specific fine-tuning for heterogeneous downstream objectives. Central to this approach is a unified tokenization framework that bridges the gap between discrete linguistic representations and mixed discrete–continuous CAN signals. We first describe the dataset and preprocessing (§III-A), then detail our tokenization strategy (§III-B), and the pretraining procedure (§III-C).

A. Dataset and Preprocessing

The pretraining dataset comprises fully anonymized, non-identifiable CAN data, decoded from $\sim 10,000$ vehicles over a 90-day period from a leading global automotive OEM, with a subset of 9 days ($\sim 19\text{B}$ tokens) used for initial training. Each trip log is decoded into 44 interpretable features in four domains: (i) vehicle dynamics (e.g., speed, acceleration), (ii) driver behavior (e.g., brake, gear), (iii) safety indicators (e.g.,

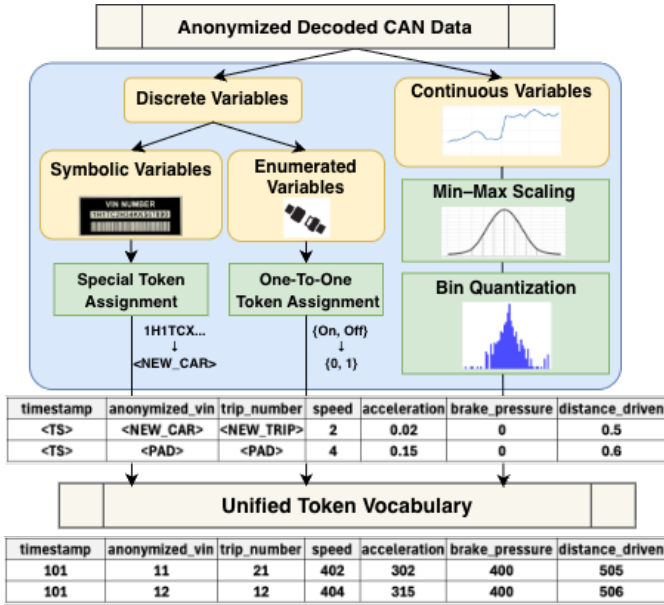


Fig. 2: Unified tokenization pipeline for mixed-mode CAN signals. Decoded CAN data is split into discrete variables (symbolic and enumerated) and continuous variables, processed through type-specific strategies, then merged into a unified token vocabulary for pretraining.

collision warning), and (iv) vehicle state and context (e.g., doors, occupancy). Decoded CAN data is synchronized to a 1 Hz frequency, normalized within empirically determined operating ranges, and segmented into 10-second windows, producing 450 tokens per sequence after tokenization.

B. Unified Tokenization of Mixed Signals

Unlike natural language, where words belong to a discrete vocabulary, CAN signals are composed of mixed discrete and continuous variables. Each decoded CAN frame consists of three signal types: 1) **Continuous variables** (e.g., speed, steering angle): numeric, real-valued, and smoothly varying over time. 2) **Discrete variables** (e.g., gear, door status): finite, enumerated categorical values. 3) **Symbolic identifiers** (e.g., anonymized vehicle identification number (VIN), trip ID): identifiers representing context boundaries between trips and vehicles.

Tokenization is therefore non-trivial and central to enabling foundation-style pretraining. A naïve one-to-one mapping of continuous values would produce an unbounded, non-reproducible vocabulary. To address this, we develop a *unified tokenization framework* for mixed-mode CAN signals that ensures semantic fidelity, reproducibility, and transferability across datasets (Fig. 2). This framework employs a fixed, predefined binning and enumeration schema to provide interpretable mappings between feature values and token IDs, emphasizing reproducibility and cross-dataset consistency.

Continuous-Variable Tokenization. For continuous signals, we employ an empirically calibrated, feature-specific quan-

tization strategy to ensure reproducible and scale-aware tokenization:

- 1) **Outlier Handling:** Each feature’s value is compared against pre-defined minimum and maximum thresholds obtained from the static sensor range limits. Values outside these bounds are replaced with a dedicated <OUTLIER> token, predefined invalid entries are replaced with <ERROR>, and <NULL> for missing values. Separating these cases enables the model to differentiate between fundamentally different failure modes and operational situations.
- 2) **Normalization:** Valid values are linearly scaled to the range [0,1] using *empirically determined* min–max bounds specific to each feature. Unlike OEM-specified sensor ranges, which are often significantly broader than values observed in real driving, these empirical limits are estimated from nine days of CAN data collected across ~10,000 vehicles. This captures realistic operating intervals while excluding synthetic or rare extreme values that distort scaling. Such calibration reduces token sparsity by preventing overly fine quantization of unoccupied value regions and ensures reproducible normalization even when OEM scaling or sensor characteristics differ across datasets.
- 3) **Discretization:** Each continuous feature is discretized into a fixed number of uniform bins (e.g., 128, 256) assigned from a preconfigured mapping based on its normalized temporal variation r_i (Eq. 1). To estimate Δ_i , approximately 100 vehicles \times 15 trips were sampled; each sequence was ordered by timestamp, clipped to valid empirical ranges, and the upper and lower 0.5% of $|x_{t+1} - x_t|$ values were excluded before averaging to suppress transient spikes. Features with smaller r_i receive finer bins, while more dynamic signals use coarser ones based on preset thresholds, ensuring reproducible yet scale-aware tokenization across datasets.

$$r_i = \frac{\Delta_i}{\max(x_i) - \min(x_i)}, \quad \Delta_i = \mathbb{E}[|x_{t+1} - x_t|], \quad (1)$$

This fixed-bin discretization produces reproducible token boundaries and maintains a consistent representation scheme across all datasets.

Discrete-Variable Tokenization. Discrete variables are further divided into two subtypes:

- 1) **Enumerated states:** low-cardinality categorical variables (e.g., turn signal, cruise control = {On, Off}) are mapped one-to-one to categorical tokens.
- 2) **Symbolic identifiers:** Context markers such as anonymized VIN and trip ID are abstracted into meta-tokens <NEW_CAR> and <NEW_TRIP> to denote context shifts without inflating the vocabulary. Analogous to special tokens in NLP (e.g., <BOS> or <SEP>) [13], these meta-tokens serve as boundary markers indicating transitions between distinct contexts, such as a new vehicle instance or trip.

During tokenization, a meta-token is inserted whenever the VIN or trip ID changes. For all subsequent samples within the same vehicle or trip context, the corresponding identifier field is filled with a <PAD> token, indicating that no context change has occurred.

This distinction allows the model to learn intra-trip temporal dependencies while maintaining generalization across trips and vehicles.

Unified Vocabulary Construction. All tokens derived from continuous and discrete variables are merged into a unified vocabulary of approximately 1,420 unique tokens, including other general special tokens such as <MASK> and <CLS>. In addition, infrastructure-level special tokens are included to encode sequence structure. <TS> (timestamp marker) is added to indicate the start of each signal sequence, thereby explicitly encoding temporal boundaries in the serialized token stream.

Sequencing of Tokens Across Time. After each signal value is converted into its corresponding token, we serialize the tokens into a fixed ordering that preserves temporal structure (see the bottom part of Fig. 2). At every timestamp (1 Hz), a <TS> and all 44 decoded CAN features are tokenized independently and then concatenated in a consistent, predetermined feature order to form a single time-step block of tokens. <TS> token is inserted at the beginning of each time-step block to mark temporal boundaries and allow the model to learn cross-feature dependencies within a moment as well as transitions across time. For a 10-second window, this produces a sequence of 10 <TS> markers followed by their associated feature tokens, resulting in 450 total tokens. This ordering ensures that the Transformer processes CAN data as a coherent, structured sequence analogous to a sentence, where temporal progression is expressed through repeated feature-wise token groups.

C. Foundation Model Pretraining

We pretrain a Transformer-based foundation model on large-scale, unlabeled decoded CAN signals to learn shared representations transferable across heterogeneous automotive tasks. The model is trained using a *MLM* objective, treating CAN signals as structured sequences analogous to natural language.

Model Architecture Tokenized sequences are used to train a Bidirectional Encoder Representations from Transformers (BERT) style Transformer encoder with an *MLM* objective without Next Sentence Prediction (NSP). Following the original pre-training setup [13], the model minimizes the cross-entropy loss between the predicted token distributions and the original masked tokens. Fifteen percent of tokens in each sequence are randomly selected for masking: 80% are replaced with the special <MASK> token, 10% with random tokens, and the remaining 10% are left unchanged. The model learns to reconstruct masked tokens using bidirectional context, capturing dependencies across temporal and semantic dimensions.

Pretraining Validation Model learning progress is evaluated via *MLM* loss convergence and internal t-SNE visualization of learned embeddings. Tokens corresponding to semantically related signals (e.g., speed and acceleration) form coherent

clusters as well as moderate intra-clusters (e.g., speed 0 - speed 100), indicating successful contextual learning. Due to proprietary constraints, these visualizations are not included. The results confirm that the proposed tokenization and pre-training pipeline effectively capture shared structure across heterogeneous CAN signals.

IV. PERFORMANCE EVALUATION

We evaluate the proposed foundation CAN model to verify its capacity for cross-task generalization and downstream adaptability. The evaluation focuses on how a large-scale pretrained representation, trained solely on decoded CAN data, transfers to specialized tasks with different label structures and class distributions. Importantly, the objective is not to outperform task-optimized or heavily engineered baselines on individual tasks, but to assess whether a single pretrained backbone can adapt across heterogeneous objectives. Rather than emphasizing absolute scores, we aim to demonstrate the robustness of the evaluation design in assessing the model’s adaptability to diverse objectives.

A. Experimental Setup

Our experiments utilize a single pretrained model comprising 9 transformer encoder layers with a hidden size of 670, 10 self-attention heads, and an intermediate feed-forward dimension of 2,680, totaling ~50M parameters, serving as a unified backbone across all downstream tasks. We adopt full-parameter fine-tuning, updating all encoder weights alongside a newly initialized classification head, and apply class-weighted loss to address severe class imbalance. This design isolates the transferability of pretrained representations from task-specific architectural choices, directly evaluating foundation model generalization in the automotive CAN domain.

To ensure that the evaluation reflects realistic deployment conditions, we compare against established baselines currently deployed in operational contexts. These include Generalized Linear Models (GLM) with a logit link [16] and 1D Convolutional Neural Networks (CNN) [17], both of which represent widely adopted and well-validated approaches in automotive and insurance applications. These baselines provide a meaningful benchmark precisely because they incorporate domain expertise and have demonstrated effectiveness in real-world settings, making them representative of the performance standards our foundation model to compare for practical adoption.

All tasks employ decoded CAN data, signals are uniformly resampled at 1 Hz to match the pretraining configuration, with each segment 10-second sequences representing the driving history corresponding to a 450-token input sequence. For supervised learning, collision labels are derived through temporal alignment between insurance records, GPS traces, and verified police accident reports. The resulting dataset captures both normal driving dynamics and rare collision scenarios, providing an ideal benchmark for evaluating representation robustness under severe class imbalance.

Ratio	Model	Precision	Recall	F1
10:1	GLM	96.9%	69.8%	81.1%
	Foundation CAN	94.7%	63.7%	76.1%
100:1	GLM	76.8%	56.8%	65.3%
	Foundation CAN	60.9%	48.9%	54.2%

TABLE I: Evaluation results for binary classification across imbalance ratios.

Model	Precision		Recall		F1	
	Macro	Weighted	Macro	Weighted	Macro	Weighted
CNN	30.4%	34.6%	27.8%	30.9%	24.3%	27.4%
Foundation CAN	28.3%	33.1%	28.2%	34.9%	27.0%	32.6%

TABLE II: Evaluation results for multiclass classification. Macro averages treat all classes equally, while weighted averages account for class imbalance by giving more weight to classes with more samples.

B. Downstream Tasks

We assess the generalization capability of the pretrained model across two representative downstream tasks:

Binary Class Collision Detection: The model predicts whether a collision event occurs within a given time window. This task evaluates sensitivity to rare and transient patterns. We compare against a statistical baseline GLM that reflects current practice, and is consistent with prior statistical approaches to collision detection using acceleration-based sensor data [18]. Although GLM is simple, its input variables are carefully engineered using domain knowledge and years of insurance industry experience; these handcrafted features represent a mature modeling pipeline, making it a realistic benchmark in practice. The dataset comprises 112,312 sequences, including 111,200 non-collision and 1,112 collision samples. To evaluate the adaptability of this homogeneous task under increasingly severe label imbalance, we further assess performance using synthetic negative-to-positive ratios of 10:1 and 100:1. These configurations serve as controlled stress tests that reveal how well the pretrained representations adapt with limited task-specific tuning in extreme imbalance settings.

Multi Class Point of Impact Classification: The model identifies the physical collision location in 8 classes based on CAN signal trajectories during an accident. As a baseline, we employ a 1D CNN, a commonly adopted approach in sensor-based fault and anomaly detection tasks [17], [19]. The dataset contains 1,088 labeled distributed across eight classes: *rear* (21.7%), *front* (18.6%), *front-right* (15.3%), *front-left* (13.8%), *rear-left* (9.6%), *rear-right* (7.4%), *left* (7.0%), and *right* (6.6%).

C. Experimental Results and Analysis

We first analyze downstream generalization under the condition of data imbalance. In the 10:1 ratio collision-detection

		Predicted	
		No Accident	Accident
Actual	No Accident	22187	53
	Accident	117	106

Fig. 3: Foundation CAN Model (10:1 Ratio)

		Predicted	
		No Accident	Accident
Actual	No Accident	22203	38
	Accident	96	126

Fig. 4: Baseline GLM (10:1 Ratio)

setting (see Table I), the pretrained foundation model achieves an F1 score of 76.1%, that is $\sim 5\%$ below the GLM baseline (81.1%). The model retains moderate rare-event sensitivity, indicating that the pretrained representations encode meaningful CAN dynamics that transfer to the binary detection objective.

In the multi-class point-of-impact task (Table II), the foundation model attains a macro-F1 of 27.0% and a weighted-F1 of 32.6%, outperforming the CNN baseline by $\sim 3\%$ and $\sim 5\%$, respectively. Since this task differs substantially from binary collision detection, the improvement demonstrates that the learned representations generalize beyond a single problem formulation. These findings support the broader claim that the pretraining \rightarrow fine-tuning paradigm established in NLP also applies to properly tokenized, decoded-CAN data.

The extremely imbalanced 100:1 collision detection stress test reveals the current model’s limitations. Under this regime, the foundation model’s F1 score decreases to 54.2%, trailing the baseline’s 65.3% by $\sim 11\%$ (Table I). This degradation indicates that while the 50M-parameter encoder captures coarse-grained CAN signal structure, it does not yet model the subtle temporal cues required for ultra-rare event discrimination under standard fine-tuning. While future iterations will address these specific temporal cues, the current results validate the broader utility of our approach, which prioritizes transferable representations and cross-task generalization through a scalable, unified framework, with downstream benefits such as reduced redundant data preparation and training costs, over disjointed task-specific pipelines, as presented in Figure 1.

Overall, the evaluation confirms that a single pretrained foundation CAN model can transfer meaningfully across tasks. At the same time, the results reveal that the current encoder under-represents fine-grained dynamics needed for extremely rare-event detection. These limitations point to opportunities for future improvements, including larger-scale pretraining,

architectural refinement, or more specialized task-aware fine-tuning approaches.

V. ABLATION STUDY

We investigate the role of model capacity in CAN representation learning by scaling the foundation model from **20M** to **100M** parameters under identical training conditions. Despite a fivefold increase in size, all models converge at approximately **10k steps** and reach losses nearly identical to the 50M-parameter model. This indicates that the primary bottleneck is the complexity of the data rather than the capacity of the encoder.

This suggests that decoded CAN data, though information-rich, exhibits relatively low compositional entropy compared to natural language, as many normal driving behaviors produce near-repetitive signal patterns. Given the current input window of **10 seconds (450 tokens)**, model size alone does not yield proportional gains.

These findings imply that transformer-based architectures could benefit from richer temporal structure and higher-resolution tokenization. Future directions include (1) **extending the temporal window** to capture long-range dependencies and (2) **adopting finer-grained bin quantization** to encode subtle signal variations using a larger token inventory. Additionally, increasing the sampling rate from **1 Hz to 5 Hz** may reveal finer temporal dynamics, improving transferability across downstream tasks.

Beyond scaling model capacity, we also highlight the importance of diversifying the pretraining corpus across a broader range of driving behaviors, analogous to multi-domain corpora in language model pretraining, which may promote more robust representations and improved generalization across heterogeneous CAN environments.

VI. LIMITATION AND DISCUSSION

Although the model generalizes across heterogeneous tasks, the pretraining corpus is limited to nine days of driving data. This duration is sufficient for our primary objective, evaluating the feasibility of a foundation-style generalization framework, but it does not capture longer-term seasonal or environmental variability such as snow, rain, or regional driving conditions. Applications that depend on these contextual factors will require substantially broader and more diverse corpora. Extending pretraining across longer time horizons would reduce seasonality bias and better support realistic cross-context generalization. While the evaluated downstream tasks are heterogeneous in label structure, the study is limited to two event-centric objectives, and broader task coverage is left for future work.

Additionally, pretraining was stopped once MLM loss plateaued; however, scaling law analyses [20] indicate that in NLP, such plateaus can appear well before the model’s true capacity is reached, particularly when data diversity is constrained. Longer horizon training paired with more heterogeneous CAN corpora may continue to improve representation quality, especially for rare event tasks. While these extensions

fall outside the scope of this study, they represent a natural progression toward deploying foundation CAN models in specialized or production-grade settings.

VII. CONCLUSION AND FUTURE WORK

We presented **Foundation CAN LM**, a pretrained language model for decoded automotive CAN data. By treating CAN signals as a structured language, the model achieves cross-task generalization across heterogeneous objectives such as collision detection and point-of-impact prediction. Our unified tokenization framework establishes a transferable representation that bridges task-specific CAN modeling and foundation-style learning. These findings suggest that properly decoded CAN data can be linguistically modeled, enabling scalable transformer architectures for automotive intelligence. In the future, we plan to extend this work to longer temporal contexts, higher-frequency data, and more diverse driving corpora.

REFERENCES

- [1] C. P. Szydlowski, “Can specification 2.0: Protocol and implementations,” SAE Technical Paper, Tech. Rep., 1992.
- [2] H. Oguma, A. Yoshioka, M. Nishikawa, R. Shigetomi, A. Otsuka, and H. Imai, “New attestation based security architecture for in-vehicle communication,” in *IEEE GLOBECOM 2008-2008 IEEE Global Telecommunications Conference*. IEEE, 2008, pp. 1–6.
- [3] U. Fugiglando, E. Massaro, P. Santi, S. Milardo, K. Abida, R. Stahlmann, F. Netter, and C. Ratti, “Driving behavior analysis through can bus data in an uncontrolled environment,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 2, pp. 737–748, 2018.
- [4] M. Hanselmann, T. Strauss, K. Dormann, and H. Ulmer, “Canet: An unsupervised intrusion detection system for high dimensional can bus data,” *Ieee Access*, vol. 8, pp. 58 194–58 205, 2020.
- [5] B. Abbache, M. Omar, and S. Bouchelaghem, “Toward optimized predictive maintenance for vehicle systems: Deep learning-based anomaly detection using can traffic,” in *14th International Conference on Pattern Recognition Applications and Methods*, 2025.
- [6] E. Nwafor and H. Olufowobi, “Canbert: A language-based intrusion detection model for in-vehicle networks,” in *2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2022, pp. 294–299.
- [7] M. Nam, S. Park, and D. S. Kim, “Intrusion detection method using bi-directional gpt for in-vehicle controller area networks,” *IEEE Access*, vol. 9, pp. 124931–124944, 2021.
- [8] T. P. Nguyen, H. Nam, and D. Kim, “Transformer-based attention network for in-vehicle intrusion detection,” *IEEE Access*, vol. 11, pp. 55 389–55 403, 2023.
- [9] A. Carlsson and B. Hans, “Can you speak forklift?: Exploring usage of large language models on forklift controller area network data for machine activity recognition,” 2024.
- [10] D. Myers, R. Mohawesh, V. I. Chellaboina, A. L. Sathvik, P. Venkatesh, Y.-H. Ho, H. Henshaw, M. Alhawawreh, D. Berdik, and Y. Jararweh, “Foundation and large language models: fundamentals, challenges, opportunities, and social impacts,” *Cluster Computing*, vol. 27, no. 1, pp. 1–26, 2024.
- [11] Y. Du, Z. Liu, J. Li, and W. X. Zhao, “A survey of vision-language pre-trained models,” *arXiv preprint arXiv:2202.10936*, 2022.
- [12] S. Qiu, B. Han, D. C. Maddix, S. Zhang, Y. Wang, and A. G. Wilson, “Transferring knowledge from large foundation models to small downstream models,” *arXiv preprint arXiv:2406.07337*, 2024.
- [13] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, 2019, pp. 4171–4186.
- [14] T.-N. Hoang, M. R. Islam, K. Yim, and D. Kim, “Canperfl: Improve in-vehicle intrusion detection performance by sharing knowledge,” *Applied Sciences*, vol. 13, no. 11, p. 6369, 2023.

- [15] W. Govers, A. Yurtman, T. Aslandere, N. Eikelenberg, W. Meert, and J. Davis, "Time-shifted transformers for driver identification using vehicle data," *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 5, pp. 3767–3776, 2024.
- [16] J. A. Nelder and R. W. Wedderburn, "Generalized linear models," *Journal of the Royal Statistical Society Series A: Statistics in Society*, vol. 135, no. 3, pp. 370–384, 1972.
- [17] K. O'shea and R. Nash, "An introduction to convolutional neural networks," *arXiv preprint arXiv:1511.08458*, 2015.
- [18] F. Becker and M. Ebner, "Collision detection for a mobile robot using logistic regression." in *ICINCO (2)*, 2019, pp. 167–173.
- [19] T. Wen and R. Keyes, "Time series anomaly detection using convolutional neural networks and transfer learning," *arXiv preprint arXiv:1905.13628*, 2019.
- [20] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, "Scaling laws for neural language models," *arXiv preprint arXiv:2001.08361*, 2020.