CANBERT+: An Optimized Multi-Layer BERT-Based Intrusion Detection for CAN Bus

Jun Yeon Won, Qadeer Ahmed, and Zhiqiang Lin

VehicleSec 2026

Modern Vehicle



Modern Vehicle

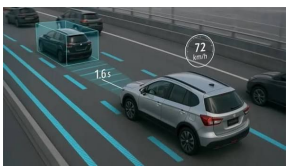




Modern Vehicle



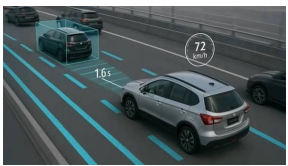


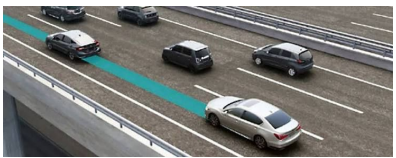


Modern Vehicle

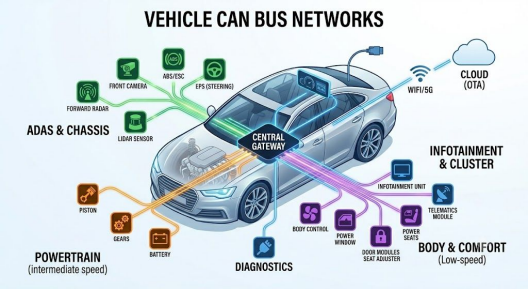




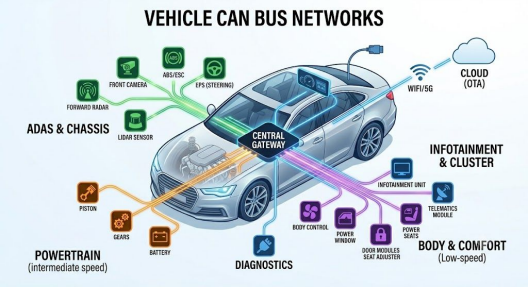




In-Vehicle Controller Area Network

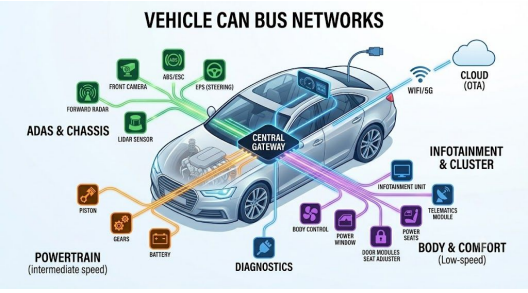


In-Vehicle Controller Area Network



SOF	**CAN ID**	RTR	IDE	DLC	**DATA (8 bytes)**	CRC	ACK	EOF

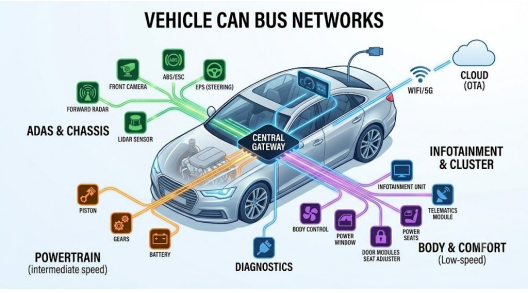
In-Vehicle Controller Area Network



CAN

- Massively Reduced Wiring

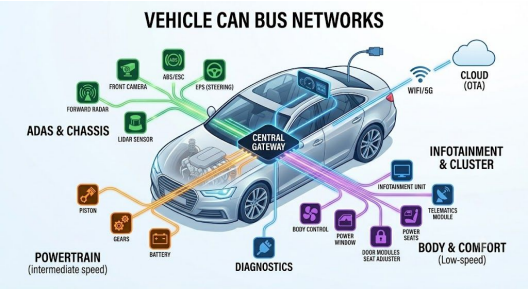
In-Vehicle Controller Area Network



CAN

- Massively Reduced Wiring
- Message Priority

In-Vehicle Controller Area Network



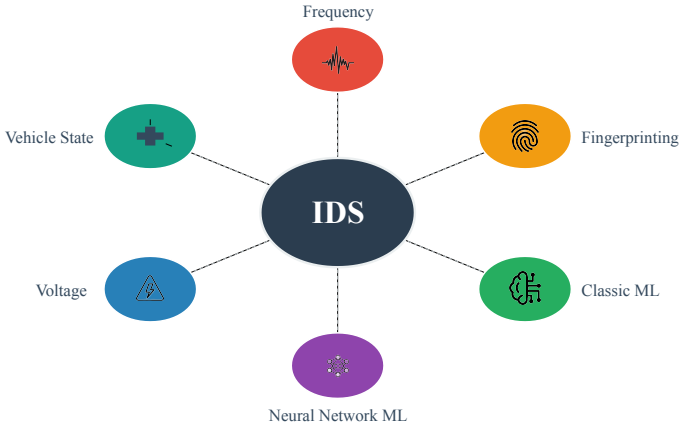
CAN

- Massively Reduced Wiring
- Message Priority
- Security Flaws

In-Vehicle Controller Area Network



Intrusion Detection System



BERT-based IDS

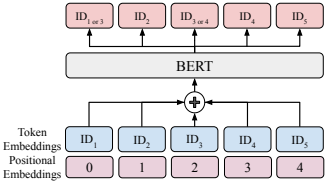


Figure: CAN-BERT [AMGD22]

BERT-based IDS

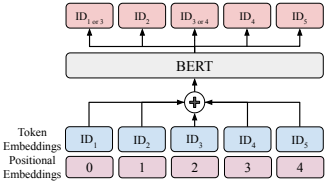


Figure: CAN-BERT [AMGD22]

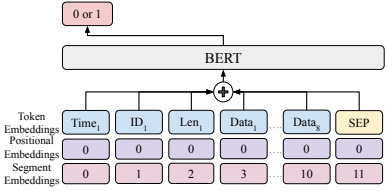


Figure: CANBERT [NO22]

BERT-based IDS

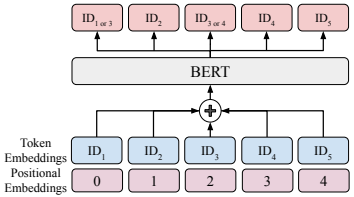


Figure: CAN-BERT [AMGD22]

Limitation

BERT-based IDS

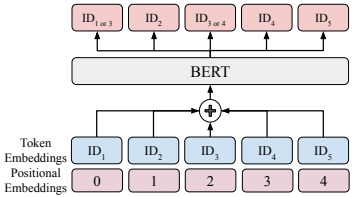


Figure: CAN-BERT [AMGD22]

Limitation

- Lack of context.

BERT-based IDS

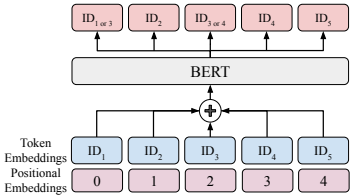


Figure: CAN-BERT [AMGD22]

Limitation

- Lack of context.
- Post-processing.

BERT-based IDS

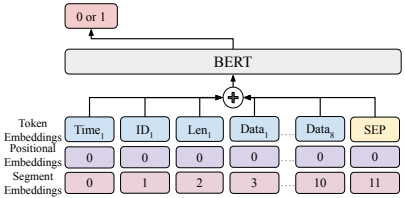


Figure: CANBERT [NO22]

Limitation

BERT-based IDS

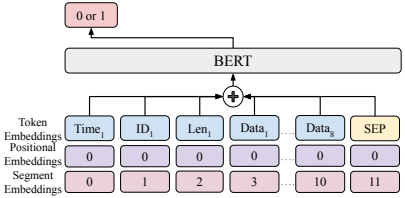


Figure: CANBERT [NO22]

Limitation

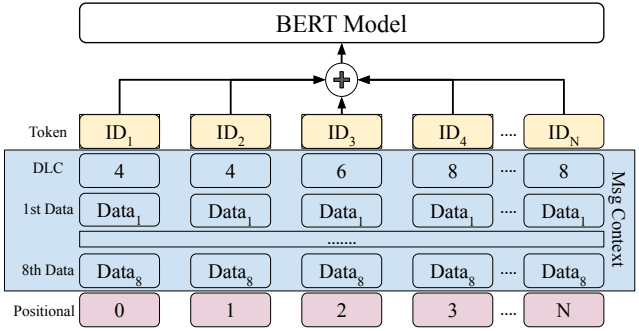
- Computationally inefficient.

BERT-based IDS

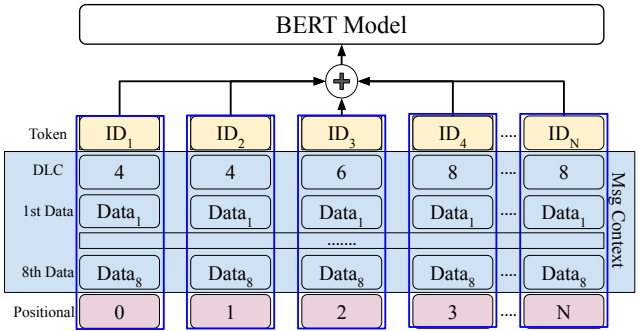
	Context	# Tokens Per Msg	Model Update
CAN-BERT [AMGD22]	✗	1	✗
CANBERT [NO22]	✓	11	✗

Table: **Comparison with Other BERT Model IDS.**

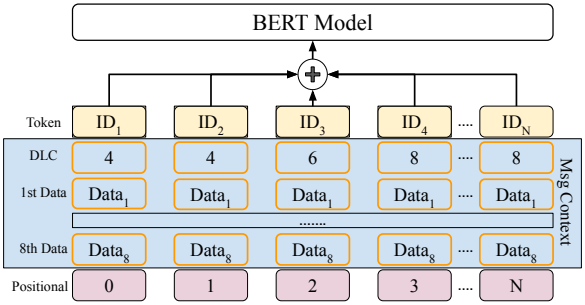
CanBert+


BERT Model
Token
DLC
1st Data
8th Data
Positional
ID1
ID2
ID3
ID4
....
IDN
4
4
6
8
....
8
Data1
Data1
Data1
Data1
....
Data1
.....
Data8
Data8
Data8
Data8
....
Data8
0
1
2
3
....
N
Msg Context

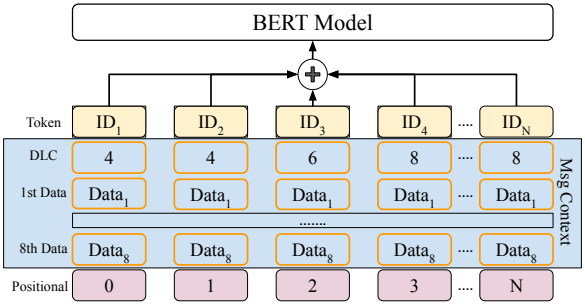
Multi-Layer Input Embeddings



Data Differentiation

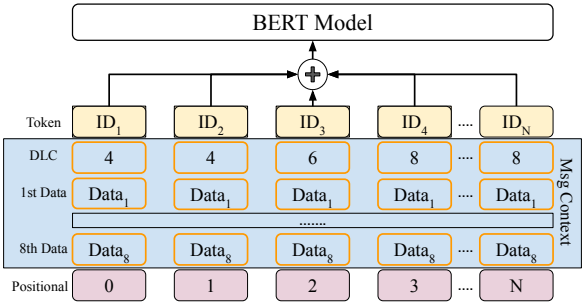


Data Differentiation



Potential Issues

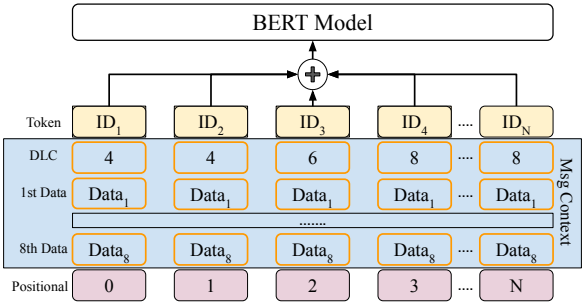
Data Differentiation



Potential Issues

- Data byte semantics

Data Differentiation



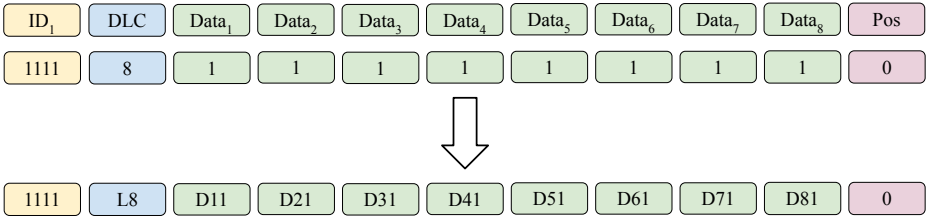
Potential Issues

- Data byte semantics
- DLC and data

Data Differentiation

ID_1	DLC	$Data_1$	$Data_2$	$Data_3$	$Data_4$	$Data_5$	$Data_6$	$Data_7$	$Data_8$	Pos
1111	8	1	1	1	1	1	1	1	1	0

Data Differentiation

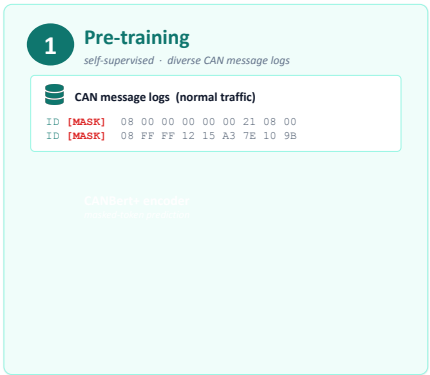


Training

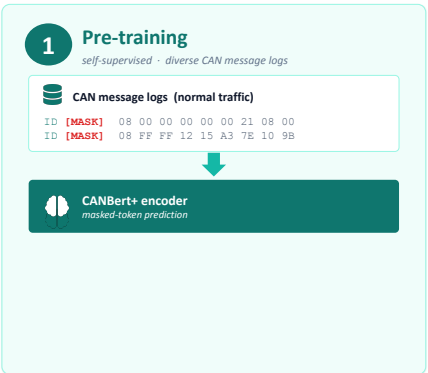
1 **Pre-training**

self-supervised · diverse CAN message logs

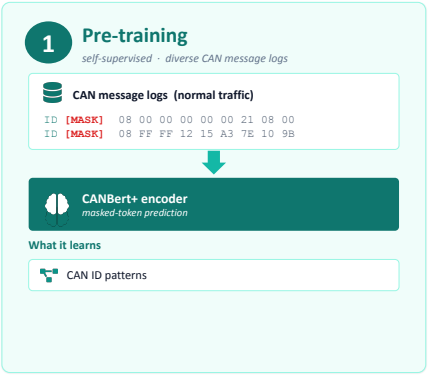
Training


1
Pre-training
self-supervised · diverse CAN message logs
CAN message logs (normal traffic)
ID [MASK] 08 00 00 00 00 00 21 08 00
ID [MASK] 08 FF FF 12 15 A3 7E 10 9B
CANBert+ encoder

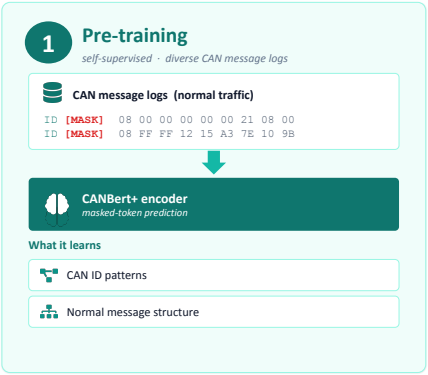
Training


1
Pre-training
self-supervised · diverse CAN message logs
CAN message logs (normal traffic)
ID [MASK] 08 00 00 00 00 00 21 08 00
ID [MASK] 08 FF FF 12 15 A3 7E 10 9B
CANBert+ encoder
masked-token prediction

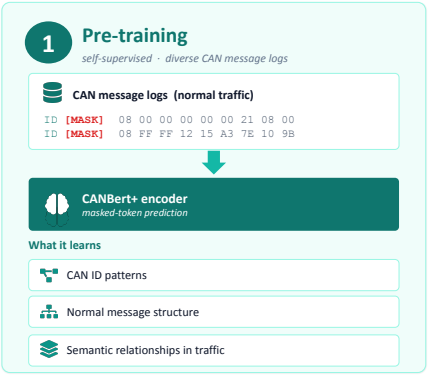
Training


1
Pre-training
self-supervised · diverse CAN message logs
CAN message logs (normal traffic)
ID [MASK] 08 00 00 00 00 00 21 08 00
ID [MASK] 08 FF FF 12 15 A3 7E 10 9B
CANBert+ encoder
masked-token prediction
What it learns
CAN ID patterns

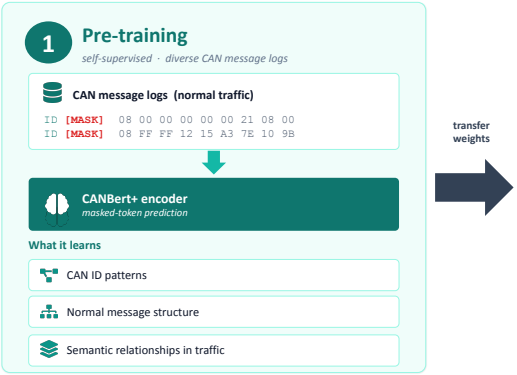
Training


1
Pre-training
self-supervised · diverse CAN message logs
CAN message logs (normal traffic)
ID [MASK] 08 00 00 00 00 00 21 08 00
ID [MASK] 08 FF FF 12 15 A3 7E 10 9B
CANBert+ encoder
masked-token prediction
What it learns
CAN ID patterns
Normal message structure

Training


1
Pre-training
self-supervised · diverse CAN message logs
CAN message logs (normal traffic)
ID [MASK] 08 00 00 00 00 00 21 08 00
ID [MASK] 08 FF FF 12 15 A3 7E 10 9B
CANBert+ encoder
masked-token prediction
What it learns
CAN ID patterns
Normal message structure
Semantic relationships in traffic

Training


1
Pre-training
self-supervised · diverse CAN message logs
CAN message logs (normal traffic)
ID [MASK] 08 00 00 00 00 00 21 08 00
ID [MASK] 08 FF FF 12 15 A3 7E 10 9B
CANBert+ encoder
masked-token prediction
What it learns
CAN ID patterns
Normal message structure
Semantic relationships in traffic
transfer weights

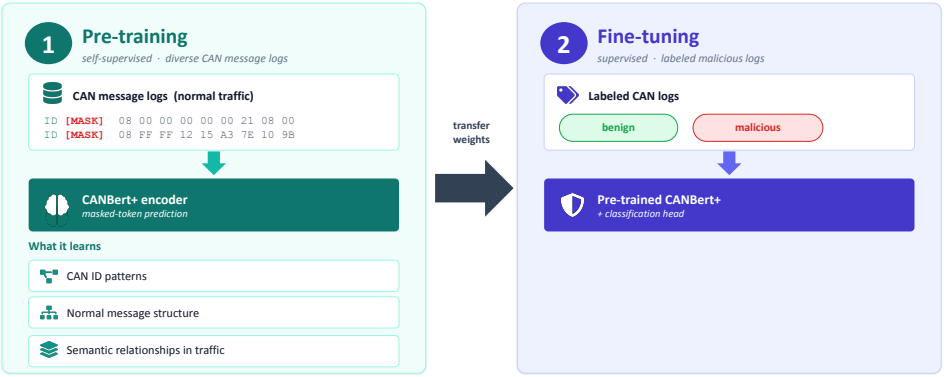
Training


1
Pre-training
self-supervised · diverse CAN message logs
CAN message logs (normal traffic)
ID [MASK] 08 00 00 00 00 00 21 08 00
ID [MASK] 08 FF FF 12 15 A3 7E 10 9B
CANBert+ encoder
masked-token prediction
What it learns
CAN ID patterns
Normal message structure
Semantic relationships in traffic
transfer weights
2
Fine-tuning
supervised · labeled malicious logs

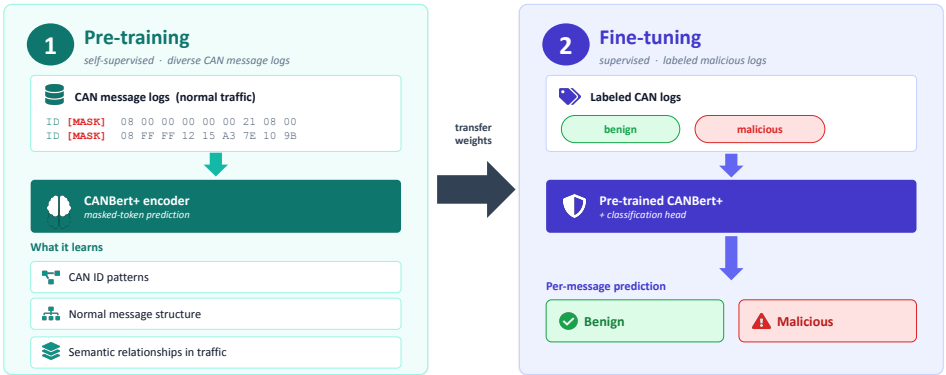
Training


1
Pre-training
self-supervised · diverse CAN message logs
CAN message logs (normal traffic)
ID [MASK] 08 00 00 00 00 00 21 08 00
ID [MASK] 08 FF FF 12 15 A3 7E 10 9B
CANBert+ encoder
masked-token prediction
What it learns
CAN ID patterns
Normal message structure
Semantic relationships in traffic
transfer weights
2
Fine-tuning
supervised · labeled malicious logs
Labeled CAN logs
benign
malicious

Training


1
Pre-training
self-supervised · diverse CAN message logs
CAN message logs (normal traffic)
ID [MASK] 08 00 00 00 00 00 21 08 00
ID [MASK] 08 FF FF 12 15 A3 7E 10 9B
CANBert+ encoder
masked-token prediction
What it learns
CAN ID patterns
Normal message structure
Semantic relationships in traffic
transfer weights
2
Fine-tuning
supervised · labeled malicious logs
Labeled CAN logs
benign
malicious
Pre-trained CANBert+
+ classification head

Training


1
Pre-training
self-supervised · diverse CAN message logs
CAN message logs (normal traffic)
ID [MASK] 08 00 00 00 00 00 21 08 00
ID [MASK] 08 FF FF 12 15 A3 7E 10 9B
CANBert+ encoder
masked-token prediction
What it learns
CAN ID patterns
Normal message structure
Semantic relationships in traffic
transfer weights
2
Fine-tuning
supervised · labeled malicious logs
Labeled CAN logs
benign
malicious
Pre-trained CANBert+
+ classification head
Per-message prediction
Benign
Malicious

Evaluation

Parameter	Value
Hidden size	768
Number of transformer layers	6
Attention heads	12
Dropout	0.1
Learning rate	0.001
Optimizer	AdamW
Adam beta	0.9, 0.999
Patience	10
Number of epochs	200
Batch size	128

Table: CanBert+ Parameter Configuration.

Evaluation

Dataset	# Logs	# Msgs	
		Benign	Attack
2019 Challenge [HCR18b]	16	4,829,571	1,086,989
CanML [cm23]	192	146,677,755	1,297,569
CSU [Cyb22]	4	82,724	14,089
Hacking [HCR18a]	4	14,237,958	2,331,517
Mcan [HCR22b]	2	1,635,080	500,000
Bcan [HCR22a]	2	5,019,845	500,000
Road [E22]	29	3,203,951	62,513
Survival [HCR18c]	9	1,105,903	183,314
Total	258	176,792,787	5,975,991

Table: List of Datasets and The Number of Logs in The Corresponding Dataset.

Evaluation

	SVM	RF	XGB	Reg	LSTM	B1	B2	OUR
	1st dataset							
File_1	0.93	0.99	0.99	0.94	0.99	0.46	0.99	0.99
File_2	0.81	0.99	0.42	0.80	0.93	0.45	0.99	0.99
Fuzzy_Sonata	0.93	0.99	0.99	0.94	0.98	0.56	0.99	0.99
Fuzzy_Soul	0.80	0.99	0.44	0.81	0.98	0.48	0.99	0.99
Malfunc_Sonata	0.00	1.00	0.80	0.00	0.66	0.00	1.00	1.00
Malfunc_Soul	0.00	0.65	0.00	0.00	0.00	0.00	0.66	1.00
	2nd dataset							
File_1	1.00	1.00	0.66	0.99	1.00	0.89	1.00	1.00
File_2	1.00	1.00	0.66	1.00	1.00	0.87	1.00	1.00
File_3	0.00	0.63	0.51	0.00	0.26	0.42	0.62	0.63
File_4	0.00	0.68	0.58	0.00	0.42	0.34	0.68	0.73
Malfunc_Sonata	1.00	1.00	0.00	1.00	1.00	0.00	1.00	1.00
Malfunc_Soul	1.00	1.00	0.99	0.99	1.00	0.00	1.00	1.00
Replay_Sonata	0.00	0.60	0.49	0.00	0.26	0.43	0.60	**0.72**
Replay_Soul	0.00	0.67	0.53	0.00	0.25	0.41	0.67	0.71

Table: F1 Score on 2019 Challenge Dataset.

Evaluation

	SVM	RF	XGB	Reg	LSTM	B1	B2	OUR
	1st dataset							
File_1	0.93	0.99	0.99	0.94	0.99	0.46	0.99	0.99
File_2	0.81	0.99	0.42	0.80	0.93	0.45	0.99	0.99
Fuzzy_Sonata	0.93	0.99	0.99	0.94	0.98	0.56	0.99	0.99
Fuzzy_Soul	0.80	0.99	0.44	0.81	0.98	0.48	0.99	0.99
Malfunc_Sonata	0.00	1.00	0.80	0.00	0.66	0.00	1.00	1.00
Malfunc_Soul	0.00	0.65	0.00	0.00	0.00	0.00	0.66	1.00
	2nd dataset							
File_1	1.00	1.00	0.66	0.99	1.00	0.89	1.00	1.00
File_2	1.00	1.00	0.66	1.00	1.00	0.87	1.00	1.00
File_3	0.00	0.63	0.51	0.00	0.26	0.42	0.62	0.63
File_4	0.00	0.68	0.58	0.00	0.42	0.34	0.68	0.73
Malfunc_Sonata	1.00	1.00	0.00	1.00	1.00	0.00	1.00	1.00
Malfunc_Soul	1.00	1.00	0.99	0.99	1.00	0.00	1.00	1.00
Replay_Sonata	0.00	0.60	0.49	0.00	0.26	0.43	0.60	**0.72**
Replay_Soul	0.00	0.67	0.53	0.00	0.25	0.41	0.67	0.71

Table: F1 Score on 2019 Challenge Dataset.

Results

CanBert+ outperforms prior baselines, especially on replay attacks.

Evaluation

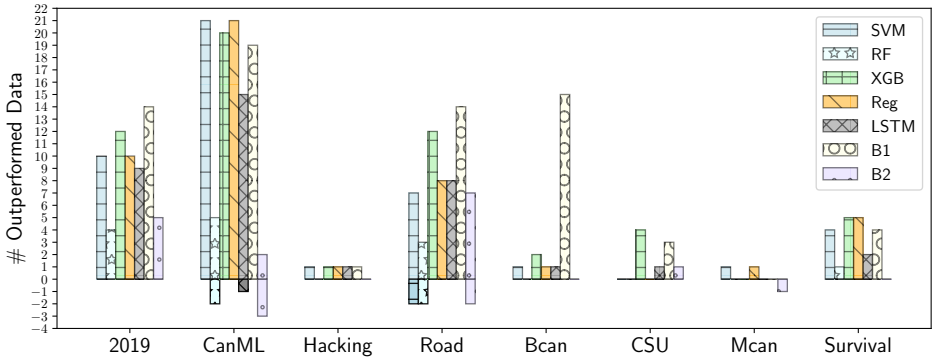


Figure: The Number of Datasets Each Model Outperforms Except for The Tie in Terms of F1 Score.

Evaluation

	Fine-tuning	Testing
CANBERT	$2,000s$/epoch	$1.3ms$/msg
CanBert+	$150s$/epoch	$0.2ms$/msg

Table: Execution Time on Fine-tuning and Testing the Messages.

Model Variants Evaluation

	# Hidden Dimensions	# Layers	# Attention Heads
CanBert+	768	6	12
s-CanBert+	60	2	12
t-CanBert+	12	2	4

Table: Configurations of Different Size Models.

Model Variants Evaluation

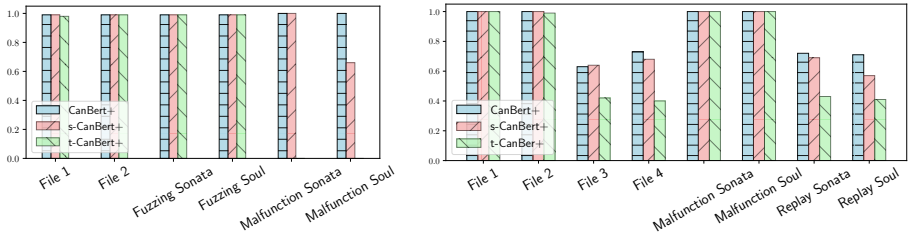
	# Hidden Dimensions	# Layers	# Attention Heads
CanBert+	768	6	12
s-CanBert+	60	2	12
t-CanBert+	12	2	4

Table: Configurations of Different Size Models.

	Param	Model	Ckpt	Mem	Time
CanBert+	61.7M	246M	708M	702.95MB	0.227ms/msg
s-CanBert+	1.6M	6.3M	50.7M	281.6M	0.199ms/msg
t-CanBert+	315K	1.2M	16.62M	274.5M	0.192ms/msg

Table: Statistics of Different Model Configurations.

Model Variants Evaluation



(a) 1st Dataset

(b) 2nd Dataset

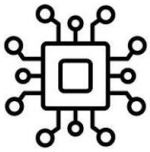
Figure: F1 Score on 2019 Challenge Dataset of Different Configuration Setups.

Limitations



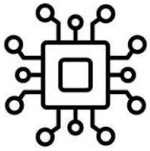
Limitations





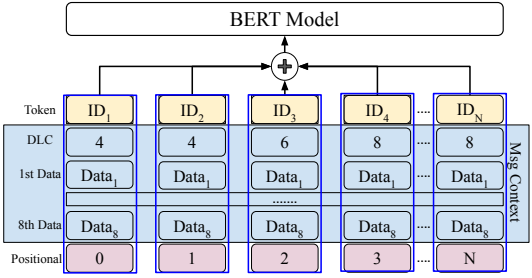
Limitations





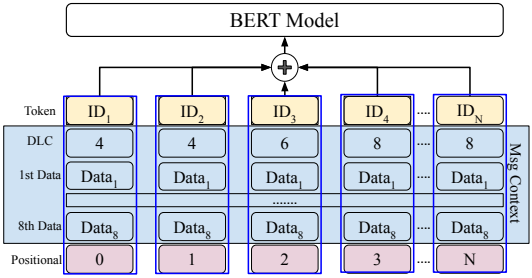


Takeaways



CanBert+

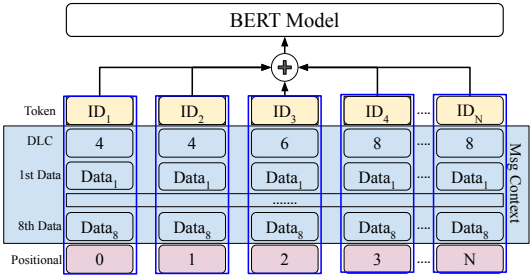
Takeaways



CanBert+

- We propose a BERT-based in-vehicle IDS using multi-layer input embeddings.

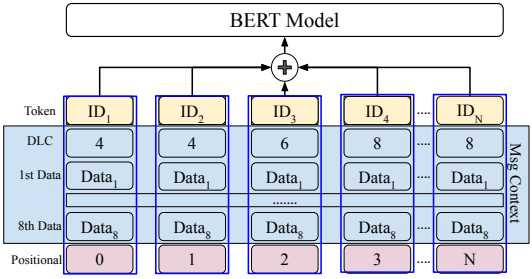
Takeaways



CanBert+

- We propose a BERT-based in-vehicle IDS using multi-layer input embeddings.
- Our approach outperforms prior IDS frameworks.

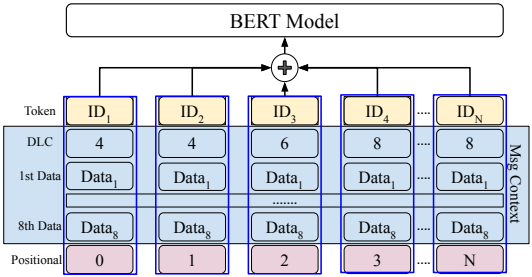
Takeaways



CanBert+

- We propose a BERT-based in-vehicle IDS using multi-layer input embeddings.
- Our approach outperforms prior IDS frameworks.
- We introduce lightweight CanBert+ variants suitable for resource-constrained deployments.

Takeaways



Source Code & Dataset

https://github.com/OSUSecLab/CanBERTplus

Thank You

CanBert+: An Optimized Multi-Layer BERT-Based Intrusion Detection for CAN Bus

Jun Yeon Won, Qadeer Ahmed, and Zhiqiang Lin

VehicleSec 2026

References I


Natasha Alkhatib, Maria Mushtaq, Hadi Ghauch, and Jean-Luc Danger, *Can-bert do it? controller area network intrusion detection system based on bert language model*, 2022 IEEE/ACS 19th International Conference on Computer Systems and Applications (AICCSA), IEEE, 2022, pp. 1–8.


can ml, *can-ml github*, https://bitbucket.org/brooke-lampe/can-ml/src/master/, 2023, Accessed: 2026-04-06.


Colorado State University System Cyber, *Colorado state university j1939 attack data*, https://projects-web.engr.colostate.edu/cybersystems/j1939-attacks/, 2022, Accessed: 2026-04-06.


Miki E, *Real ornl automotive dynamometer (road) can intrusion dataset*, https://0xsam.com/road/, 2022, Accessed: 2026-04-06.


HCRL, *Car-hacking dataset*, https://ocslab.hksecurity.net/Datasets/car-hacking-dataset, 2018, Accessed: 2026-04-06.


_____, *In-vehicle network intrusion detection challenge*, https://ocslab.hksecurity.net/Datasets/datachallenge2019/car, 2018, Accessed: 2026-04-06.


_____, *Survival analysis dataset for automobile ids*, https://ocslab.hksecurity.net/Datasets/survival-ids, 2018, Accessed: 2026-04-06.


_____, *B-can intrusion dataset*, https://ocslab.hksecurity.net/Datasets/b-can-intrusion-dataset, 2022, Accessed: 2026-04-06.


_____, *M-can intrusion dataset*, https://ocslab.hksecurity.net/Datasets/m-can-intrusion-dataset, 2022, Accessed: 2026-04-06.


Ebelechukwu Nwafor and Habeeb Olufowobi, *Canbert: A language-based intrusion detection model for in-vehicle networks*, 2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA), IEEE, 2022, pp. 294–299.