

Reducing Non-IID Effects in Federated Autonomous Driving with Contrastive Divergence Loss

Tuong Do^{1,3}, Binh X. Nguyen¹, Quang D. Tran¹, Hien Nguyen¹, Erman Tjiputra¹,
Te-Chuan Chiu², Anh Nguyen³

Abstract—Federated learning has been widely applied in autonomous driving since it enables training a learning model among vehicles without sharing users’ data. However, data from autonomous vehicles usually suffer from the non-independent-and-identically-distributed (non-IID) problem, which may cause negative effects on the convergence of the learning process. In this paper, we propose a new contrastive divergence loss to address the non-IID problem in autonomous driving by reducing the impact of divergence factors from transmitted models during the local learning process of each silo. We also analyze the effects of contrastive divergence in various autonomous driving scenarios, under multiple network infrastructures, and with different centralized/distributed learning schemes. Our intensive experiments on three datasets demonstrate that our proposed contrastive divergence loss significantly improves the performance over current state-of-the-art approaches. Our source code is available at <https://github.com/aioz-ai/CDL>.

I. INTRODUCTION

Autonomous driving is an emerging field that enables vehicles to operate without a human driver by using a combination of vision, learning, and control algorithms to observe and respond to changes in the environment [1]. Recently, many works have been proposed to address different problems in autonomous driving [2]–[4]. While significant progress has been made in the field, traditional works utilize supervised learning methods and require data collection to train the model [5], [6]. Although collecting data is necessary to improve the system’s accuracy, it strongly violates user privacy since the users’ data are shared with third parties. Recent works have adapted federated learning (FL) as a new learning mechanism to overcome this limitation. Federated learning allows multiple parties to collaboratively train a model without sharing their data [7]–[10]. In practice, federated learning enables autonomous vehicles to learn a shared prediction model together, involving more diverse data while preserving users’ privacy [11].

Typically, there are two main federated learning scenarios [12]: Server-based Federated Learning (SFL) which has a central node to orchestrate the training process and receive the contributions of all clients, and Decentralized Federated Learning (DFL) which utilizes a fully peer-to-peer (P2P) setup between data silos using a predefined topology. Although SFL can enhance data privacy since only model weights are transmitted, having orchestration nodes




Fig. 1. Sample viewpoints over three vehicles and their steering angle distributions in Carla dataset [7]. The visualization shows that the three vehicles have differences in visual input as well as steering angle distribution.

potentially represents a bottleneck of the system since most of the data transmission has to go through the central nodes. On the other hand, DFL does not require a server and uses a fully distributed network. In autonomous driving, several works have explored both DFL and SFL to address different problems such as collision avoidance [8], [13], trajectory prediction [14], and steering prediction [7], [9].

In practice, while SFL or DFL approaches have their own advantages and limitations, both of them suffer from the non-IID problem in federated learning. According to [15], [16], the *non-IID* (independently and identically distributed) problem occurs when data partitioning across silos has a significant *distribution shift*. Although the non-IID problem occurred in many contexts, it is an immense problem in autonomous driving and causes difficulties when the accumulation process for all vehicle silos is conducted [17]. In practice, each vehicle has its unique driving patterns, weather conditions, and road types, which can cause differences in the data distribution. For example, data collected from a car driving on a highway may be different from data collected from a car driving in a city center. When the accumulation process for all vehicles is conducted, the non-IID problem can cause difficulties in building robust and accurate learning models, e.g., if the model is trained mainly on data from highways, it may perform poorly in urban environments. Figure 1 illustrates different scenarios when we have the non-IID problem in autonomous driving.

Prior works that address the non-IID problem by optimizing the accumulation step [17]–[19], adding normalization to the global model [20], fine-tuning global model weights through distillation [21], or aligning the distribution between the global model and local ones [22], [23]. In general,

¹AIOZ, Singapore tuong.khanh-long.do@aioz.io

²Department of Computer Science, National Tsing Hua University, Taiwan theochoiu@cs.nthu.edu.tw

³Department of Computer Science, University of Liverpool, UK anh.nguyen@liverpool.ac.uk

these solutions try to adapt to divergence factors by refining the global model weights. However, this process can be intricate for optimization since the global model receives weights learned from diverse distributions of distinct local datasets [24]. Consequently, these methods encounter unresolved hurdles such as identifying the ideal accumulation step size tailored to each local silo, dealing with disconnection and bottleneck in the network topology, or ensuring the consistency between local and global objectives.

In this paper, we propose a new Contrastive Divergence Loss (CDL) to address the non-IID problem in autonomous driving. Instead of *waiting for the model* to be learned locally and then dealing with the Non-IID problem during the global aggregating process as other methods [17], [25], [26], we *directly mitigate the impact of divergence factors* throughout the learning process of each local silo using our proposed CDL loss. Compared with dealing Non-IID problem in the global model [19], [25]–[27], facilitating the adaptation of distribution between neighboring silos is less complex. In this way, each local model is more robust to changes in the data distribution and we can use typical accumulation methods (e.g., FedAvg [28]) to conduct global weights without concerning the Non-IID effects on the convergence. The intensive experiments on three autonomous driving datasets verify our observation and show significant improvements over state-of-the-art methods.

II. RELATED WORKS

Autonomous Driving. Autonomous driving is an emerging field that has attracted significant research interests in recent years. Several works have focused on using deep learning for autonomous vehicle control, such as object detection and tracking [3], [29], trajectory prediction [4], autonomous braking and steering [30]. Recently, Xin *et al.* [31] proposed a recursive backstepping steering controller that effectively links yaw-rate-based path following commands to the steering angle. Xiong *et al.* [32] analyzed the nonlinear dynamics behavior using a proportional control law. Yi *et al.* [33] presented an algorithm to select the instantaneous center of rotation within the self-reconfigurable robot’s area and perform static rotation to adjust its heading angle during waypoint navigation while avoiding collisions. Recently, Yin *et al.* [14] combined model predictive control with covariance steering theory to obtain a robust controller for general nonlinear autonomous driving systems.

Federated Learning for Autonomous Driving. Federated learning allows many participants to train a machine learning model cooperatively without disclosing their local data [11]. Federated learning offers a privacy-aware solution to many automotive systems, such as cooperative autonomous driving and intelligent transport systems, which require efficient communication, computation, and storage [34]. Recently, many works have been proposed to address different problems in autonomous driving using federated learning [35]–[39]. Liang *et al.* [8] presented an online federated reinforcement transfer learning process where all the vehicles make actions from the knowledge learned by others. The

authors in [9] proposed an autonomous driving system to preserve vehicle privacy by using a server to store shared training models between vehicles. Zhang *et al.* [35] proposed an approach that included a unique asynchronous model aggregation mechanism. The authors in [7] introduced a deep federated network for steering angle prediction. Recently, Doomra *et al.* [40] used federated learning to predict the turning signal.

Although applying federated learning in autonomous driving is trendy, there are several open challenges. Besides the crucial disadvantages correlated to the availability and quality of the local computing devices, the variability in data distribution across the participants is also challenging. Specifically, the data of various vehicles are usually non-IID. While non-IID is a critical problem, most of the recent works focus on addressing it through the accumulation process [17]–[19], [41] and do not focus on local silos optimization.

Contrastive Divergence. Contrastive models have been studied in FL recently for dealing with the heterogeneity of local data distribution across parties [9]. This type of model has been applied to vision datasets [19], language datasets [42], and signal datasets [43]. Indeed, contrastive learning is a valuable candidate for dealing with the non-IID problem in federated autonomous driving, which is mostly caused by local data heterogeneity. It is worth noting that most contrastive-related works focus on building optimized frameworks in server-based [25], [44], modifying topology in P2P federated learning [19], [45], or adjusting the accumulation process [22], [46], [47]. However, as far as our knowledge, no works have considered the contrastive behavior of the loss function to reduce the effect of the Non-IID data in the federated autonomous driving scenario.

III. PRELIMINARY

A. Notation

We summarize the notations of our paper in Table I.

TABLE I
MATHEMATICAL NOTATIONS.

Not.	Description	Not.	Description
ξ	Data in a mini-batch	i, j	Workers (silos)
α	Learning rate	m	Mini-batch size
\mathbf{A}	Consensus matrix	\mathcal{L}	Loss
k	One specific iteration	$\vartheta \in \{0, 1\}$	Aggregation status
x	Input image of ξ in k -th iteration	x'	Input image of ξ in $(k - 1)$ -th iteration
\mathcal{H}	Kullback-Leibler Divergence	θ	Learnable parameters
β	Pulling control parameter for Kullback-Leibler Divergence	\mathcal{N}_i^+	In-neighbors of silo i
\mathbf{K}	Transition kernel	\mathbf{k}_1	1-st transitional kernel constraint

B. Federated Learning for Autonomous Driving

We consider each autonomous vehicle as a data silo. Our goal is to collaboratively train a global driving policy θ from N silos by aggregating all local learnable weights θ_i of each

silo. Each silo computes the current gradient of its local loss function and then updates the model parameter using an optimizer. Mathematically, in the local update stage, at each silo i , in each iteration k , the model weights of the local silo can be computed as:

$$\theta_i(k+1) = \theta_i(k) - \alpha_k \frac{1}{m} \sum_{h=1}^m \nabla \mathcal{L}_{lr}(\theta_i(k), \xi_i^h(k)) \quad (1)$$

where \mathcal{L}_{lr} is the local regression loss for autonomous steering. To update the global model, each silo interacts with the associated ones through a predefined topology:

$$\theta(k+1) = \sum_{i=0}^{N-1} \vartheta_i \theta_i(k) \quad (2)$$

In practice, the local model in each silo is a deep network that takes the RGB images as inputs and predicts the associated steering angles [7].

IV. FEDERATED AUTONOMOUS DRIVING WITH CONTRASTIVE DIVERGENCE LOSS

A. Overview

Motivation. Due to the non-IID problem, the federated learning algorithms for autonomous driving only achieve good results when two conditions are met: *i*) the local silo i can effectively learn from its local data, and *ii*) the synchronization between nearby silos is sufficient to minimize the effect of the non-IID problem. Currently, most recent works address both of these problems by optimizing the accumulation progress and optimizer [17], [38], [39], proposing new topologies [36], or utilizing high-performed deep networks that are robust to the non-identical characteristic of distributed data [18], [48]. However, according to [49], naively adopting high-performed deep architecture on centralized local learning and its corresponding optimizations into federated scenarios can increase the weight variance of local silo weights during the weights accumulation between silos. As a consequence, it affects the convergence ability of the model and may also cause divergence.

Siamese Network. In this work, we propose to address the non-IID problem directly from each local silo by solving the learning good feature problem and synchronization challenge separately. Consequently, we want to have *two networks in each silo*. One network is responsible for learning meaningful features from the local image data, and the other network is responsible for minimizing the distribution gap between the current model weight and other nearby weights. To this end, the Siamese Network [50] which includes two networks is well-fitted to our needs. In particular, the Siamese network has two branches. We consider the first branch (*backbone network*) of the Siamese to have the local regression loss \mathcal{L}_{lr} to learn the local image features for autonomous steering and the positive contrastive divergence loss \mathcal{L}_{cd+} to learn knowledge transmitted from neighbor silos. The second branch (*sub-network*) is utilized to control the divergence factors

from knowledge of the backbone through the contrastive regularizer term \mathcal{L}_{cd-} (See Fig. 2).

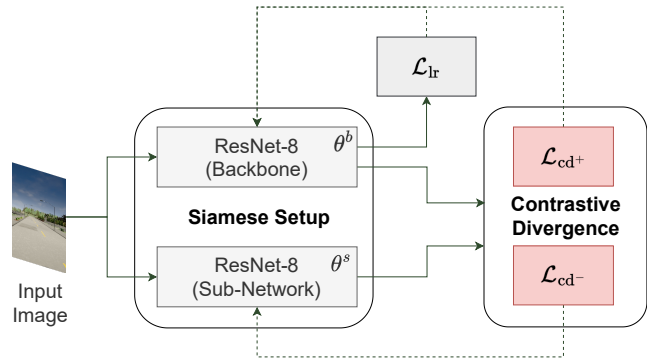


Fig. 2. The Siamese setup when our CDL is applied for training federated autonomous driving model. ResNet-8 is used in the backbone and sub-network in the Siamese setup. During inference, the sub-network will be removed. Dotted lines represent the backward process. Our CDL has two components: the positive contrastive divergence loss \mathcal{L}_{cd+} and the negative regularize term \mathcal{L}_{cd-} . The local regression loss \mathcal{L}_{lr} for automatic steering prediction is calculated only from the backbone network.

In practice, the sub-network shares the same weights with the backbone during the first communication round. However, from the next communication round, after the backbone is accumulated using Equation 1, the local model of each silo is trained by the contrastive divergence loss. The sub-network outputs support features that have the same size as the output features of the backbone. During the training, we expect that the weights between the backbone and the sub-network should not have significant differences while applying the contrastive divergence loss. The weights of all silos are synchronized whenever gradients from the learning process of the backbone and the sub-network of all silos are not significantly different.

B. Contrastive Divergence Loss

In practice, we observe that the early stages of federated learning mostly have poor accumulated models. Different from other works that deal with the non-IID problem by optimizing the accumulation step whenever silos transmit their models, we directly decrease the effect of divergence factors during the local learning process of each silo. To achieve that, we reduce the distance between distribution from accumulated weights θ_i^b at silo i in the backbone network, which contains information from other silos known as divergence factors, and its i -th silo weights θ_i^s in the sub-network, which only contains knowledge learned from local data. When the distribution between silos has been synchronized at an acceptable rate, we lower the effectiveness of the sub-network and focus more on the steering angle prediction task. Our proposed Contrastive Divergence Loss is motivated by the contrastive loss from the original Siamese Network [50] and is defined as:

$$\begin{aligned} \mathcal{L}_{cd} &= \beta \mathcal{L}_{cd+} + (1 - \beta) \mathcal{L}_{cd-} \\ &= \beta \mathcal{H}(\theta_i^b, \theta_i^s) + (1 - \beta) \mathcal{H}(\theta_i^s, \theta_i^b) \end{aligned} \quad (3)$$

where $\mathcal{L}_{\text{cd}+}$ is the positive contrastive divergence term and $\mathcal{L}_{\text{cd}-}$ is the negative regularizer term; \mathcal{H} is the Kullback-Leibler Divergence loss function [51]:

$$\mathcal{H}(\hat{y}, y) = \sum \mathbf{f}(\hat{y}) \log \left(\frac{\mathbf{f}(\hat{y})}{\mathbf{f}(y)} \right) \quad (4)$$

where \hat{y} is the predicted representation, y is dynamic soft label.

Consider $\mathcal{L}_{\text{cd}+}$ in Equation 3 as a Bayesian statistical inference task, our goal is to estimate the model parameters θ^{b*} by minimizing the Kullback-Leibler divergence $\mathcal{H}(\theta_i^b, \theta_i^s)$ between the measured regression probability distribution of the observed local silo $P_0(x|\theta_i^s)$ and the accumulated model $P(x|\theta_i^b)$. Hence, we can assume that the model distribution has a form of $P(x|\theta_i^b) = e^{-E(x, \theta_i^b)} / Z(\theta_i^b)$, where $Z(\theta_i^b)$ is the normalization term. However, evaluating the normalization term $Z(\theta_i^b)$ is not trivial, which leads to risks of getting stuck in a local minimum. Inspired by Hinton [52], we use samples obtained through a Markov Chain Monte Carlo (MCMC) procedure with a specific initialization strategy to deal with the mentioned problem. Additionally inferred from Equation 1, the $\mathcal{L}_{\text{cd}+}$ can be expressed under the SGD algorithm in a local silo by setting:

$$\mathcal{L}_{\text{cd}+} = - \sum_x P_0(x|\theta_i^s) \frac{\partial E(x; \theta_i^b)}{\partial \theta_i^b} + \sum_x Q_{\theta_i^b}(x|\theta_i^s) \frac{\partial E(x; \theta_i^b)}{\partial \theta_i^b} \quad (5)$$

where $Q_{\theta_i^b}(x|\theta_i^s)$ is the measured probability distribution on the samples obtained by initializing the chain at $P_0(x|\theta_i^s)$ and running the Markov chain forward for a defined step.

Consider $\mathcal{L}_{\text{cd}-}$ regularizer in Equation 3 as a Bayesian statistical inference task, we can calculate $\mathcal{L}_{\text{cd}-}$ as in Equation 5, however, the role of θ^s and θ^b is inverse:

$$\mathcal{L}_{\text{cd}-} = - \sum_x P_0(x|\theta_i^b) \frac{\partial E(x; \theta_i^s)}{\partial \theta_i^s} + \sum_x Q_{\theta_i^s}(x|\theta_i^b) \frac{\partial E(x; \theta_i^s)}{\partial \theta_i^s} \quad (6)$$

We note that although Equation 5 and Equation 6 share the same structure, the key difference is that while the weight θ_i^b of the backbone is updated by the accumulation process from Equation 2, the weight θ_i^s of the sub-network, instead, is not. This lead to different convergence behavior of contrastive divergence in $\mathcal{L}_{\text{cd}+}$ and $\mathcal{L}_{\text{cd}-}$. The negative regularizer term $\mathcal{L}_{\text{cd}-}$ will converge to state θ_i^{s*} provided $\frac{\partial E}{\partial \theta_i^s}$ is bounded:

$$g(x, \theta_i^s) = \frac{\partial E(x; \theta_i^s)}{\partial \theta_i^s} - \sum_x P_0(x|(\theta_i^b, \theta_i^s)) \frac{\partial E(x; \theta_i^s)}{\partial \theta_i^s} \quad (7)$$

and

$$(\theta_i^s - \theta_i^{s*}) \cdot \left\{ \sum_x P_0(x)g(x, \theta_i^s) - \sum_{x', x} P_0(x')\mathbf{K}_{\theta_i^s}^m(x', x)g(x, \theta_i^{s*}) \right\} \geq \mathbf{k}_1 |\theta_i^s - \theta_i^{s*}|^2 \quad (8)$$

for any \mathbf{k}_1 constraint. Note that, $\mathbf{K}_{\theta_i^s}^m$ is the transition kernel. The proof for the above result is analyzed in [53].

Note that the negative regularizer term $\mathcal{L}_{\text{cd}-}$ is only used in training models on local silos. Thus, it does not contribute to the accumulation process of federated training.

C. Total Training Loss

Local Regression Loss. We use mean square error (MAE) to compute loss for predicting the steering angle in each local silo. Note that, we only use features from the backbone for predicting steering angles.

$$\mathcal{L}_{\text{lr}} = \text{MAE}(\theta_i^b, \hat{\xi}_i) \quad (9)$$

where $\hat{\xi}_i$ is the ground-truth steering angle of the data sample ξ_i collected from silo i .

Local Silo Loss. The local silo loss computed in each communication round at each silo before applying the accumulation process is described as:

$$\mathcal{L}_{\text{final}} = \mathcal{L}_{\text{lr}} + \mathcal{L}_{\text{cd}} \quad (10)$$

In practice, we observe that both the contrastive divergence loss \mathcal{L}_{cd} to handle the non-IID problem and the local regression loss \mathcal{L}_{lr} for predicting the steering angle is equally important and indispensable. Hence, we do not set a parameter to control their contributions in Equation 10.

Combining all losses together, at each iteration k , the update in the backbone network is defined as:

$$\theta_i^b(k+1) = \begin{cases} \sum_{j \in \mathcal{N}_i^+ \cup \{i\}} \mathbf{A}_{i,j} \theta_j^b(k), & \text{if } k \equiv 0 \pmod{u+1}, \\ \theta_i^b(k) - \alpha_k \frac{1}{m} \sum_{h=1}^m \nabla \mathcal{L}_b(\theta_i^b(k), \xi_i^h(k)), & \text{otherwise.} \end{cases} \quad (11)$$

where $\mathcal{L}_b = \mathcal{L}_{\text{lr}} + \mathcal{L}_{\text{cd}+}$, u is the number of local updates.

In parallel, the update in the sub-network at each iteration k is described as:

$$\theta_i^s(k+1) = \theta_i^s(k) - \alpha_k \frac{1}{m} \sum_{h=1}^m \nabla \mathcal{L}_{\text{cd}-}(\theta_i^s(k), \xi_i^h(k)) \quad (12)$$

V. EXPERIMENT

A. Implementation

Dataset. We use three datasets (Table II) in our experiment: Udacity+ [54], Gazebo Indoor [7], and Carla Outdoor dataset [7]. Gazebo and Calar are non-IID datasets while Udacity+ is the non-IID version of the Udacity dataset.

TABLE II
THE STATISTIC OF DATASETS IN OUR EXPERIMENTS.

Dataset	Total samples	Average samples in each silo		
		Gaia [55] (11 silos)	NWS [56] (22 silos)	Exodus [55] (79 silos)
Udacity+	38,586	3,508	1,754	488
Gazebo	66,806	6,073	3,037	846
Carla	73,235	6,658	3,329	927

Network Topology. Following [12], we conduct experiments on three federated topologies: the Internet Topology Zoo [55] (Gaia), the North American data centers [56] (NWS), and the Zoo Exodus network (Exodus) [55]. We use Gaia topology in our main experiment and provide the comparison of two other topologies in our ablation study.

TABLE III
PERFORMANCE COMPARISON BETWEEN DIFFERENT METHODS. THE GAIA TOPOLOGY IS USED.

Model	Main Focus	Learning Method	RMSE			MAE			# Training Parameters	Avg. Cycle Time (ms)
			Udacity+	Gazebo	Carla	Udacity+	Gazebo	Carla		
Random [57]	-	-	0.358	0.117	0.464	0.265	0.087	0.361	-	-
Constant [57]	Statistical	-	0.311	0.092	0.348	0.209	0.067	0.232	-	-
Inception [58]	Architecture Design	CLL [57]	0.209	0.085	0.297	0.197	0.062	0.207	21,787,617	-
MobileNet [59]			0.193	0.083	0.286	0.176	0.057	0.200	2,225,153	-
VGG-16 [60]			0.190	0.083	0.316	0.161	0.050	0.184	7,501,587	-
DroNet [57]			0.183	0.082	0.333	0.150	0.053	0.218	314,657	-
FedAvg [61]	Aggregation Optimization	SFL [15]	0.212	0.094	0.269	0.185	0.064	0.222	314,657	152.4
FedProx [62]			0.152	0.077	0.226	0.118	0.063	0.151	314,657	111.5
STAR [15]			0.179	0.062	0.208	0.149	0.053	0.155	314,657	299.9
MATCHA [63]	Topology Design	DFL [7]	0.182	0.069	0.208	0.148	0.058	0.215	314,657	171.3
MBST [12], [64]			0.183	0.072	0.214	0.149	0.058	0.206	314,657	82.1
FADNet [7]			0.162	0.069	0.203	0.134	0.055	0.197	317,729	62.6
CDL (ours)	Loss Optimization	CLL [7]	0.169	0.074	0.266	0.149	0.053	0.172	629,314	-
		SFL [7]	0.150	0.060	0.208	0.104	0.052	0.150	629,314	102.2
		DFL [7]	0.141	0.062	0.183	0.083	0.052	0.147	629,314	72.7

Training. The model in a silo is trained with a batch size of 32 and a learning rate of 0.001 using Adam optimizer. In each communication round, the local training process is done in each silo before their models are transmitted and aggregated using Equation 2. The training process is conducted with 3,600 communication rounds. We apply the simulation as in [7] for training with an NVIDIA 1080 GPU.

Baselines. We compare our results with various recent methods in different learning scenarios, including Random baseline and the Constant baseline [57]. For the Centralized Local Learning (CLL) scenario, Inception-V3 [58], MobileNet-V2 [59], VGG-16 [60], and Dronet [57] are used as baselines. For the Server-based Federated Learning (SFL) scenario, we compare our method with FedAvg [61], FedProx [62], and STAR [15]. For the Decentralized Federated Learning (DFL), MATCHA [63], MBST [12], and FADNet [7] are used as baselines. We use the Root Mean Square Error (RMSE) and Mean Absolute Error (MAE) to evaluate the effectiveness of models. Besides, wall-clock time (ms) is used to calculate the time needed for training each method.

B. Qualitative Results

Table III summarises the performance of our method and recent state-of-the-art approaches. This table clearly shows the proposed CDL under Siamese setup with two ResNet-8 outperforms other methods by a large margin. In particular, our proposal significantly reduces both RMSE and MAE in all three datasets, including Udacity+, Carla, and Gazebo. As a loss function, CDL does not increase the number of parameters of the network. However, under the Siamese setup, the model size during the training is increased as the Siamese requires an additional sub-network. Moreover, we can see that our CDL with ResNet-8 outperforms other baselines by a large margin in the DFL learning scenario, and by a moderate value in SFL and CLL setup.

TABLE IV
PERFORMANCE UNDER DIFFERENT TOPOLOGIES.

Topology	Architecture	Dataset		
		Udacity+	Gazebo	Carla
Gaia (11 silos)	DroNet [57]	0.177(↓0.036)	0.073(↓0.011)	0.244(↓0.061)
	FADNet [7]	0.162(↓0.021)	0.069(↓0.007)	0.203(↓0.020)
	CDL (ours)	0.141	0.062	0.183
NWS (22 silos)	DroNet [57]	0.183(↓0.045)	0.075(↓0.017)	0.239(↓0.057)
	FADNet [7]	0.165(↓0.027)	0.070(↓0.012)	0.200(↓0.018)
	CDL (ours)	0.138	0.058	0.182
Exodus (79 silos)	DroNet [57]	0.448(↓0.310)	0.208(↓0.147)	0.556(↓0.380)
	FADNet [7]	0.179(↓0.041)	0.081(↓0.02)	0.238(↓0.062)
	CDL (ours)	0.138	0.061	0.176

C. Contrastive Divergence Loss Analysis

CDL on Different Topologies. In practice, it is usually more challenging to train federated algorithms when the topology has more vehicle data silos. To verify the effectiveness of our CDL, we train our method and compare it with other baselines on topologies with different numbers of silos. Table IV illustrates the performance of DroNet, FADNet, and our CDL with ResNet-8 backbone when we train them using DFL under three distributed network infrastructures with different numbers of silos: Gaia (11 silos), NWS (22 silos), and Exodus (79 silos). This table shows that our CDL clearly achieves the highest results in all topology setups, while DroNet meets divergence, and FADNet does not perform well in the Exodus topology which has 79 silos.

CDL on Different Backbones. Since our proposed CDL is a loss function, it can be applied to different networks under the Siamese setup to improve performance. Figure 3 illustrates the effectiveness of CDL when we change the network inside the Siamese to DroNet, FADNet, Incep-

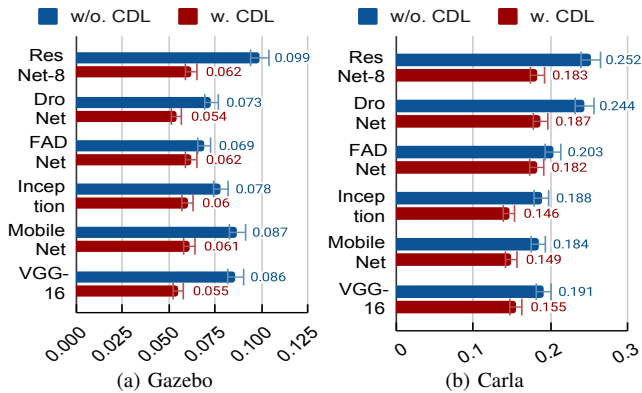


Fig. 3. Performance of CDL under different networks in Siamese setup.




Fig. 4. Performance of different methods on IID dataset (Udacity) and non-IID dataset (Udacity+).

tion, MobileNet, and VGG-16 under Gaia Network in the DFL scenario. The results show that CDL works well with different architectures to address the non-IID problem and consistently improving the performance.

CDL on IID Data. Figure 4 demonstrates the effectiveness of CDL in different data distributions. Although CDL is designed for dealing with the non-IID problem, it also slightly improves the performance of models trained on IID data distribution. Based on the Siamese setup, CDL inherits the behavior and characteristic of triplet loss form. Since triplet loss is proven to be effective in IID data [65], it is clear that CDL can also improve the performance of models when we train them with the IID data.

D. Ablation Study

Convergence Analysis. Figure 5 illustrates the training results in RMSE of our two baselines DroNet and FADNet as well as our proposed CDL. The results show the convergence ability of mentioned methods in three datasets (Udacity+, Gazebo, Carla) with NWS and Gaia topology. The results indicate that our proposed CDL can reach a better convergence point in comparison with the two baselines. While other methods (DroNet and FADNet) converged with difficulty or do not show good convergence trend, our proposed CDL can get over local optimal points better than other methods and also be less biased into any specific silo.

Robotic Demonstration. We deploy the global model aggregated from our CDL with ResNet-8 of each vehicle trained on the Non-IID Udacity+ dataset to a standard mobile robot platform. The robot is equipped with an 8-Core ARM v8.2 64-bit CPU and NVIDIA Jetson. Since the inference

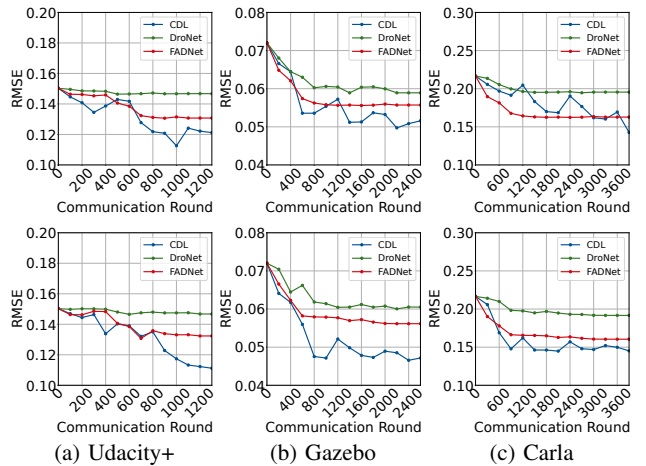


Fig. 5. The convergence ability of different methods under Gaia topology (top row) and NWS topology (bottom row).

time of our method is only 42 ms and the number of parameters is only approximately 630,000, it can predict steering angles for the robot in real time. More quantitative results can be found in our Demonstration Video.

E. Discussion

Although our method shows promising results, we see several improvement points that can be considered for future work: *i)* Our CDL is designed to mitigate the non-IID problem, thus, the approach’s effectiveness heavily relies on the existence of non-IID characteristics within the autonomous driving data. Consequently, in scenarios where the data exhibits relatively consistent distributions across vehicles or lacks significant non-IID factors, the proposed contrastive divergence loss only yields limited performance gains. *ii)* Our proposal has been validated on autonomous driving datasets, but not on real vehicles. A real-world study on different driving scenarios involving interactions with pedestrians, cyclists, and other vehicles can further validate our method. And *iii)* While the approach is tailored for autonomous driving, its principles might find applications in other domains where non-IID data is present. Exploring its efficacy in areas such as healthcare, IoT, and industrial settings could broaden its scope of impact.

VI. CONCLUSION

We presented a new method to address the non-IID problem in federated autonomous driving using contrastive divergence loss. Our method directly reduces the effect of divergence factors in the learning process of each silo. The experiments on three benchmarking datasets demonstrate that our proposed method performs substantially better than current state-of-the-art approaches. In the future, we plan to test our strategy with more data silos and deploy the trained model using an autonomous vehicle on roads.

Acknowledgment. This work made use of the facilities of the N8 Centre of Excellence in Computationally Intensive Research (N8 CIR) provided and funded by the N8 research partnership and EPSRC (Grant No. EP/T022167/1).

REFERENCES

- [1] S. Grigorescu, B. Trasnea, T. Cocias, *et al.*, “A survey of deep learning techniques for autonomous driving,” *JFR*, 2020.
- [2] L. Bergamini, Y. Ye, O. Scheel, *et al.*, “Simnet: Learning reactive self-driving simulations from real-world observations,” in *ICRA*, 2021.
- [3] Y.-N. Chen, H. Dai, and Y. Ding, “Pseudo-stereo for monocular 3d object detection in autonomous driving,” in *CVPR*, 2022.
- [4] J. Wang, T. Ye, Z. Gu, *et al.*, “Ltp: Lane-based trajectory prediction for autonomous driving,” in *CVPR*, 2022.
- [5] K. Huang and Q. Hao, “Joint multi-object detection and tracking with camera-lidar fusion for autonomous driving,” in *IROS*, 2021.
- [6] M. Vitelli, Y. Chang, *et al.*, “Safetynet: Safe planning for real-world self-driving vehicles using machine-learned policies,” in *ICRA*, 2022.
- [7] A. Nguyen, T. Do, M. Tran, *et al.*, “Deep federated learning for autonomous driving,” in *IV*, 2022.
- [8] X. Liang, Y. Liu, T. Chen, *et al.*, “Federated transfer reinforcement learning for autonomous driving,” *arXiv*, 2019.
- [9] Y. Li, X. Tao, X. Zhang, *et al.*, “Privacy-preserved federated learning for autonomous driving,” *T-ITS*, 2021.
- [10] B. X. Nguyen, T. Do, *et al.*, “Reducing training time in cross-silo federated learning using multigraph topology,” in *ICCV*, 2023.
- [11] C. Zhang, Y. Xie, H. Bai, *et al.*, “A survey on federated learning,” *KBS*, 2021.
- [12] O. Marfoq, C. Xu, G. Neglia, *et al.*, “Throughput-optimal topology design for cross-silo federated learning,” in *NIPS*, 2020.
- [13] W. Hammedi, B. Brik, and S. M. Senouci, “Toward optimal mec-based collision avoidance system for cooperative inland vessels: a federated deep learning approach,” *T-ITS*, 2022.
- [14] J. Yin *et al.*, “Trajectory distribution control for model predictive path integral control using covariance steering,” in *ICRA*, 2022.
- [15] F. Sattler, S. Wiedemann, *et al.*, “Robust and communication-efficient federated learning from non-iid data,” *TNNLS*, 2019.
- [16] J. Wang and G. Joshi, “Cooperative sgd: A unified framework for the design and analysis of communication-efficient sgd algorithms,” in *ICLRW*, 2018.
- [17] B. Li, Y. Jiang, Q. Pei, *et al.*, “Feel: Federated end-to-end learning with non-iid data for vehicular ad hoc networks,” *T-ITS*, 2022.
- [18] H. Zhang, J. Bosch, and H. H. Olsson, “End-to-end federated learning for autonomous driving vehicles,” in *IJCNN*, 2021.
- [19] X. Mu, Y. Shen, K. Cheng, *et al.*, “Fedproc: Prototypical contrastive federated learning on non-iid data,” *FGCS*, 2023.
- [20] M. J. Idrissi, I. Berrada, and G. Noubir, “Fedbds: Learning on non-iid data in federated learning using batch normalization,” in *2021 IEEE 33rd International Conference on Tools with Artificial Intelligence (ICTAI)*. IEEE, 2021.
- [21] L. Zhang, L. Shen, L. Ding, D. Tao, and L.-Y. Duan, “Fine-tuning global model via data-free knowledge distillation for non-iid federated learning,” in *CVPR*, 2022.
- [22] L. Gao, H. Fu, L. Li, Y. Chen, M. Xu, and C.-Z. Xu, “Feddc: Federated learning with non-iid data via local drift decoupling and correction,” in *CVPR*, 2022.
- [23] L. Zhang, Y. Luo, Y. Bai, B. Du, and L.-Y. Duan, “Federated learning for non-iid data via unified feature learning and optimization objective alignment,” in *ICCV*, 2021.
- [24] H. Zhu, J. Xu, S. Liu, and Y. Jin, “Federated learning on non-iid data: A survey,” *Neurocomputing*, 2021.
- [25] Y. Tan, G. Long, J. Ma, L. Liu, T. Zhou, and J. Jiang, “Federated learning from pre-trained models: A contrastive learning approach,” *NIPS*, 2022.
- [26] X. You, X. Liu, N. Jiang, J. Cai, and Z. Ying, “Reschedule gradients: Temporal non-iid resilient federated learning,” *IoT-J*, 2022.
- [27] H. Yang, M. Fang, and J. Liu, “Achieving linear speedup with partial worker participation in non-iid federated learning,” *ICLR*, 2021.
- [28] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, “On the convergence of fedavg on non-iid data,” in *ICLR*, 2019.
- [29] H. Hu, Z. Liu, *et al.*, “Investigating the impact of multi-lidar placement on object detection for autonomous driving,” in *CVPR*, 2022.
- [30] N. Ijaz and Y. Wang, “Automatic steering angle and direction prediction for autonomous driving using deep learning,” in *ISCSIC*, 2021.
- [33] L. Yi, V. Le, *et al.*, “Anti-collision static rotation local planner for four independent steering drive self-reconfigurable robot,” in *ICRA*, 2022.
- [31] M. Xin *et al.*, “Slip-based nonlinear recursive backstepping path following controller for autonomous ground vehicles,” in *ICRA*, 2020.
- [32] J. Xiong, B. Li, R. Yu, *et al.*, “Reduced dynamics and control for an autonomous bicycle,” in *ICRA*, 2021.
- [34] Z. Du, C. Wu, T. Yoshinaga, *et al.*, “Federated learning for vehicular internet of things: Recent advances and open issues,” *OJ-CS*, 2020.
- [35] H. Zhang, J. Bosch, and H. H. Olsson, “Real-time end-to-end federated learning: An automotive case study,” *arXiv*, 2021.
- [36] Y. He, K. Huang, G. Zhang, *et al.*, “Bift: A blockchain-based federated learning system for connected and autonomous vehicles,” *IoT-J*, 2021.
- [37] M. A. Khan, H. El Sayed, *et al.*, “A journey towards fully autonomous driving-fueled by a smart communication system,” *VC*, 2022.
- [38] X. Liang, Y. Liu, T. Chen, *et al.*, “Federated transfer reinforcement learning for autonomous driving,” in *FTL*, 2022.
- [39] R. Parekh, N. Patel, *et al.*, “Gefli: Gradient encryption-aided privacy preserving federated learning for autonomous vehicles,” *Access*, 2023.
- [40] S. Doomra, N. Kohli, and S. Athavale, “Turn signal prediction: A federated learning case study,” *arXiv*, 2020.
- [41] Y. Li, A. S. Sani, D. Yuan, and W. Bao, “Enhancing federated learning robustness through clustering non-iid features,” in *ACCV*, 2022.
- [42] K. Georgiev and P. Nakov, “A non-iid framework for collaborative filtering with restricted boltzmann machines,” in *ICML*, 2013.
- [43] V. Tsouvalas, A. Saeed, and T. Ozcelebi, “Federated self-training for semi-supervised audio recognition,” *TECS*, 2022.
- [44] T. Qi, F. Wu, C. Wu, *et al.*, “Fairfl: A fair vertical federated learning framework with contrastive adversarial learning,” in *NIPS*, 2023.
- [45] A. Bellet, A.-M. Kermarrec, and E. Lavoie, “D-cliques: Compensating noniidness in decentralized federated learning with topology,” *arXiv*, 2021.
- [46] H. Kassem, D. Alapatt, P. Mascagni, *et al.*, “Federated cycling (fedcy): Semi-supervised federated learning of surgical phases,” *TMI*, 2022.
- [47] S. Hu, L. Feng, X. Yang, *et al.*, “Fedssc: Shared supervised-contrastive federated learning,” *arXiv*, 2023.
- [48] P. S. Bouzinis *et al.*, “Wireless federated learning (wfl) for 6g networks part i: Research challenges and future trends,” *COMML*, 2021.
- [49] M. Duan, D. Liu, X. Chen, *et al.*, “Self-balancing federated learning with global imbalanced data in mobile systems,” *TPDS*, 2020.
- [50] G. Koch, R. Zemel, R. Salakhutdinov, *et al.*, “Siamese neural networks for one-shot image recognition,” in *ICMLW*, 2015.
- [51] J. M. Joyce, “Kullback-leibler divergence,” in *International encyclopedia of statistical science*, 2011.
- [52] G. E. Hinton, “Training products of experts by minimizing contrastive divergence,” *Neural computation*, 2002.
- [53] A. L. Yuille, “The convergence of contrastive divergences,” *NIPS*, 2004.
- [54] Udacity, “An open source self-driving car,” 2016.
- [55] S. Knight, H. X. Nguyen, N. Falkner, *et al.*, “The internet topology zoo,” *J-SAC*, 2011.
- [56] F. P. Miller, A. F. Vandome, and J. McBrewster, “Amazon web services,” *Retrieved November*, 2011.
- [57] A. Loquercio, A. I. Maqueda, C. R. del Blanco, *et al.*, “Dronet: Learning to fly by driving,” *RA-L*, 2018.
- [58] C. Szegedy, V. Vanhoucke, S. Ioffe, *et al.*, “Rethinking the inception architecture for computer vision,” in *CVPR*, 2015.
- [59] M. Sandler, A. Howard, M. Zhu, *et al.*, “Mobilenetv2: Inverted residuals and linear bottlenecks,” in *CVPR*, 2018.
- [60] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv*, 2014.
- [61] B. McMahan, E. Moore, D. Ramage, *et al.*, “Communication-efficient learning of deep networks from decentralized data,” in *AISTATS*, 2017.
- [62] T. Li, A. K. Sahu, M. Zaheer, *et al.*, “Federated optimization in heterogeneous networks,” *arXiv*, 2018.
- [63] J. Wang, A. K. Sahu, Z. Yang, *et al.*, “Matcha: Speeding up decentralized sgd via matching decomposition sampling,” in *ICC*, 2019.
- [64] R. C. Prim, “Shortest connection networks and some generalizations,” *The Bell System Technical Journal*, 1957.
- [65] T. Wang *et al.*, “Understanding contrastive representation learning through alignment and uniformity on the hypersphere,” in *ICML*, 2020.