Popular Content Prediction Through Adversarial Autoencoder Using Anonymised Data

Deborah Vieira de Alencar Maia
University of Coimbra,
CISUC/LASI, DEI
and Federal Institute of Education,
Science, and Technology of RN
Email: dmaia@dei.uc.pt

Joao P Vilela
CISUC, CRACS/INESC TEC,
and Department of Computer Science,
University of Porto - Porto, Portugal.
Email: jvilela@fc.up.pt

Marilia Curado
University of Coimbra,
CISUC/LASI, DEI
Email: marilia@dei.uc.pt

***Abstract*—The increasing number of connected and autonomous vehicles generates an even greater demand for efficient content delivery in vehicular networks. Estimating the popularity of content is an important task to proactively cache and distribute content throughout the networks to add value to users' experiences and reduce network congestion. This paper presents a novel approach for predicting popular content on vehicular networks based on a Federated Learning-Adversarial Autoencoder model and anonymised data. Unlike prior works that relied on users' raw features, our model protects user privacy through data anonymisation. This allows us to learn from the hidden patterns of content popularity and deliver popular content without compromising user privacy. Experiments showed that our approach exceeded traditional collaborative filtering and deep learning methods in terms of accuracy and robustness, even with sparse data.**

***Index Terms*—Adversarial Autoencoder, Federated Learning, Data Anonymisation, Popular Content, Vehicular Network**

I. Introduction

Vehicular networks have emerged due to the rapid expansion of Internet-connected cars, allowing for the interchange of data and information between infrastructure and vehicles [1]. These networks are used for various purposes, including infotainment services, content distribution, traffic monitoring, and accident prediction [2]. Vehicular networks are dynamic and heterogeneous, which brings challenges involving high mobility, intermittent connectivity, and privacy concerns [3]. The complexity of this scenario calls for improvements in content prediction and dissemination methodologies. The efficiency of content distribution strategies can be greatly improved by accurately predicting the popularity of content in vehicular networks [4]. However, this task is challenging due to the dynamic nature of vehicular networks and the need to maintain user privacy.

Adversarial Autoencoder (AAE) models emerge as a solution due to the power of deep learning and adversarial networks to learn the underlying patterns in vehicular data and predict popular content. Kim et al. [5] demonstrated that adversarial autoencoder models achieve powerful and generalisable data representations, improving content prediction accuracy with user confidentiality. Yet, due to privacy concerns, Federated Learning (FL) techniques have come out to avoid sharing raw sensitive information. However, it is known that Federated Learning is also subject to a range of inference attacks [6], [7] against private data. Using Federated Learning with anonymised data can reduce privacy concerns. This allows the AAE model to be trained on anonymised data without centralising data collection.

Edge content caching is a frequent practice that can serve to relieve network traffic and yield a higher user quality of experience (QoE) [1], [8], [9]. Proactive caching previously caches the content based on the prediction of the trained model. Collaborative filtering and matrix factorisation approaches have been used to forecast content popularity [4]. Other methods have been proposed [10] that use the advantages of adversarial networks and autoencoder architectures to identify the patterns of content preference and acquire an effective representation of the data. Some works have shown that user information could improve content predictions [11]; however, in compliance with GDPR[1] and LGPD[2] regulations, user privacy must be guaranteed.

This work improves upon this Adversarial Autoencoder Model [10] by adding an anonymisation technique. The aim is to forecast popular content that most vehicle users find interesting and cache it at the vehicular network edge. The goal is to improve privacy without affecting system utility. Extensive experiments on a real-world dataset demonstrate the effectiveness of our proposed model compared to other state-of-the-art methods in terms of prediction accuracy, robustness, and privacy preservation.

The main contributions are summarised as follows:

1) Develop an Adversarial Autoencoder Model to predict popular content on vehicular networks using anonymised data and Federated Learning.
2) Assess the effectiveness of the proposed model concerning prediction accuracy, content diversity, and privacy preservation.
3) Investigate how adding user information affects the model's performance and the trade-offs between accu-

[1]General Data Protection Regulation - https://gdpr-info.eu/

[2]Brazil's General Data Protection Law - https://www.planalto.gov.br/ccivil_03/_ato2015-2018/2018/lei/l13709.htm

racy and privacy.

The rest of the article is organised as follows: the related works are presented in Section II. Section III and IV describe the model implementation and training in detail. The results found and their corresponding analyses are presented in Section V. Finally, the conclusions and the prospective directions for further research are stated in Section VI.

II. Related Works

Increasing demands for computing, communication, and storage in new edge devices have accelerated the development of Mobile Edge Computing (MEC) as a solution to such challenges [12]. In this same direction, Vehicular Edge Computing (VEC) arrives as a paradigm that combines MEC's features with vehicular networks to offer communication, computing, and caching resources closer to vehicular users, allowing its features to enhance road safety and traffic efficiency and value-added applications such as infotainment or path planning [1]. Vehicular applications collect significant data, including personal details like age, gender, and working hours, which can endanger user privacy [13]. Hence, ensuring user privacy in such situations is of the highest priority.

The VEC's architecture combines vehicular terminals on the user's layer, edge servers such as roadside units (RSU) on the MEC layer, and cloud servers on the cloud layer. Vehicles can offload their demanding computational tasks to the edge servers, reducing latency and energy consumption. In VEC contexts, several research works have suggested solutions for resource allocation, task offloading, and security assurance [5], [14]. In addition, content distribution and caching are other main concerns for VEC. Vehicles also often produce and consume a wide range of content related to traffic updates, accident reports, infotainment, and multimedia, which might cause congestion and add latency to the network.

Using a latent representation of user preferences and network characteristics, an AAE model may accurately forecast popular content in vehicular networks while maintaining resistance to possible attacks or privacy-preserving issues [15], [16]. However, despite this strategy, issues of privacy continue to exist. Federated Learning can help address data privacy and communication overhead challenges in vehicular networks [10], [17].

The Federated Learning paradigm has been used to improve data privacy in many decentralised systems, such as vehicular networks [10], [11]. This approach allows multiple clients (e.g., vehicles) to collaboratively train a shared model without centralising raw data, thereby maintaining privacy and reducing issues caused by the impact of intermittent connectivity. It minimises communication by only sharing model updates (gradients) rather than the raw data itself, significantly reducing the communication overhead. However, despite its use in ensuring data privacy, recent research suggests that Federated Learning has flaws that affect its privacy benefits [17]–[19], as training data may be reconstructed from the shared model updates using adversarial methods such as model inversion.

As Jere et al. [17] presented, privacy attacks, such as inference and poisoning attacks, may compromise the model integrity and data, limiting Federated Learning use. To deal with it, Choudhury et al. [20] proposed implementing anonymisation methods to enhance the data utility and the model efficacy model. Integrating Adversarial Autoencoders, Federated Learning, and anonymisation techniques, coupled with caching strategies, offers a robust privacy-preserving approach for predicting and delivering popular content in vehicular networks.

Using anonymous data and Federated Learning can improve privacy. This approach aims to train the model to acquire an underlying representation of input data, including features like user demographics, content metadata, and network conditions. This latent representation is then passed through a discriminator network that differentiates between the true latent representation and randomly generated data. The developed model learns a representation that is informative for content prediction while maintaining robustness against adversarial perturbations, making it appropriate for implementation in a privacy-preserving and secure vehicular network context [15].

III. Proposal Architecture

Our approach consists of an Adversarial Autoencoder model employing Federated Learning and anonymised data to predict popular content within vehicular networks. The raw data is on vehicles, and before training, it goes through an anonymisation process. This step ensures that the shared model updates do not contain personally identifiable information (PII), providing an extra layer of privacy protection. Therefore, the vehicles train the model and return the top K content and weights from the model trained. At the RSU, the Federated Learning process aggregates the client's weights and computes a new model's version for the next training iteration.

The model architecture comprehends the encoder and the decoder, forming a Variational Autoencoder (VAE). To build the adversarial part, a discriminator is added to create the GAN (Generative Adversarial Network). The combination of them can achieve artificial user-item interactions that simulate real interactions. This approach can help reduce data sparsity issues and address the cold start problem [21]. The model illustration in Figure 1 depicts the overall framework. In this scheme, the input layer receives anonymised data and passes it through the encoder to generate a latent representation to feed the decoder part. The training process involves a min-max game between the encoder-decoder and the discriminator to learn a latent representation that effectively captures key content features while remaining resilient to adversarial attacks. The process is further detailed in Algorithm 1 where a subset of clients are selected to participate in each communication round. The clients download the current global model, preprocess and anonymise their local data, and perform several steps of local gradient descent updates on the model. The updated local models are sent back to the central server,

which performs a federated averaging step to aggregate the updates, generate a new global model, rank predicted content, and cache the top K contents.

Algorithm 1 explains how the model works, using a Federated Learning approach to predict popular content in vehicle networks with an Adversarial Autoencoder design. This prioritises protecting user privacy while optimising content caching for efficient delivery.

Algorithm 1 Model Execution Algorithm

```
Initialise ω_0
for each round r = 1, 2, ... do
    r_c: A set of selected clients in the r^th round
    for each client c ∈ C in parallel do
        Download current global model ω_r
        Preprocess and anonymise local data
        for epoch e = 1, 2, ... do
            Compute parameters with gradient descent
            ω_{c+1} = ω_c − η∇ι(ω_c; b)
        Predict popular contents pop_c
        Store pop_c into pop_r
        Compute federated averaging
    Rank predicted content
    Cache Top K contents
```

IV. Mehodology

The model was implemented and evaluated using Keras as the deep learning framework and Tensorflow as the backend on an M2 notebook with 24GB of memory and ten cores of graphics processing units (GPUs). The input dataset used for training the model is MovieLens 1M [22]. This dataset version counts over 1 million ratings from 6040 anonymised users, who rated 3883 movies.

A. Autoencoder Model

Before the training phase, the data was preprocessed for use. It includes applying an arbitrary threshold filter for ratings to ensure a more reliable statistic, excluding movies and users with low ratings. It improves the model's quality and efficiency and brings more accurate and relevant recommendations; checking the movie identifiers in the ratings and the movies datasets to ensure they are identical; after, a binary matrix is built with all rated movies marked with 1 (one), representing interesting content. For the uninterested and the unknown content, a random sample mechanism was applied to mark them as 1 (one) based on the probability of the user's preference for content and the rest as 0 (zero). This approach is presented in Yu et al. work [10]. To effectively contribute to the model, the system encapsulates the personal information from the users' dataset using one-hot encoding and then combines it with the previous binary matrix to conclude the preprocessing phase. The prediction results should be enhanced by incorporating user information as side information [11]. Our proposed approach examines the impact of anonymising that information on prediction accuracy.

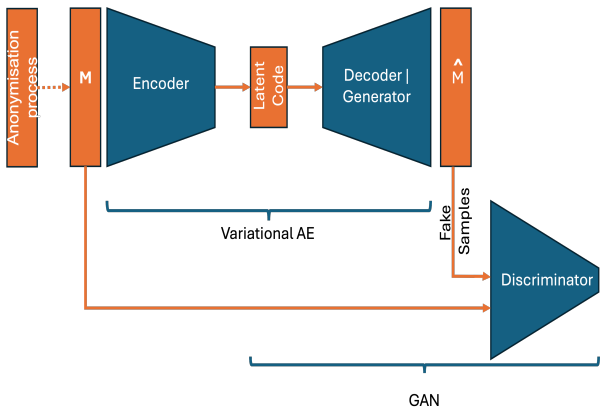

Fig. 1. AAE Model Illustration

After the preprocessing phase, the dataset was split into 80/10/10 rates for training, testing, and validation, respectively, which guaranteed a reliable evaluation and maintained statistical significance during the evaluation process, thus reaching a balance between the training, test and validation processes. The basis parameters for training are shown in Table I.

TABLE I
Model hyperparameters for training simulation

General Parameters			Learning Parameters	
Hyperparameters	**Value**		**Hyperparameters**	**Value**
Number of rounds	3		Dropout rate	0.8
Number of clients	50		VAE learning rate	0.002
Number of epochs	20		Discriminator learning rate	0.002
Batch size	256		Regularisation factor	0.008
Top K	400			

The initial training combines the binary matrix aggregated with raw users' features (gender, age, occupation, and zip code), generating the input matrix $\mathbf{M}$. After that, the model was trained with each user's features individually to examine the impact of that information on the model's performance. The training output can be seen in Figure 2. This figure shows the cache hit ratio (CHR) as a function of the cache size (CS) for the baseline dataset complemented with four user features.

User data anonymisation may reduce prediction accuracy due to losing valuable information. Therefore, to determine the significance of user features for the anonymisation process, we take the CHR for each CS and compute the overall $\text{CHR}_{\text{advantage}}$ regarding the baseline and each user feature. The $\text{CHR}_{\text{advantage}}$ of specific features is defined in Equation (1).

$$\text{CHR}_{\text{advantage}} = \frac{\sum \left(\frac{\text{CHR}_{\text{feature}} - \text{CHR}_{\text{base}}}{\text{CHR}_{\text{base}}} \times 100 \right)}{n} \quad (1)$$

where n is the total number of cache sizes, $\text{CHR}_{\text{feature}}$ is the cache hit ratio for the specific feature, and CHR_{base} is the cache hit ratio for the baseline model. The features that enable

a higher and $CHR_{advantage}$ are thus more important for popular content prediction tasks.

For better visualisation, we plotted the relationship between user features and the baseline individually. The plots in Figure 2 and Tables III and IV display more performance measures and respective CHR of each feature as well. We point out the last column of Table III, which shows the advantage of CHR for each case. The results indicate that the age feature has a higher $CHR_{advantage}$ for content prediction tasks, followed by gender. The graph highlights the relative importance of these user features in improving the model's performance compared to the baseline. Analysing the impact of each feature provides insights into which user information is most valuable for enhancing content predictions. This can help inform strategies to balance privacy and model accuracy through optimal data anonymisation.

B. Feature Anonymisation

For the anonymisation step, ARX Data Anonymization Tool[3] was used. ARX is a tool that implements various privacy models (e.g., k-anonymity, l-diversity, differential privacy) and enables their configuration to achieve the best privacy-utility trade-offs. While ARX offers a range of models, this work focuses on k-anonymity. For our anonymisation process, we primarily utilised ARX's generalisation, suppression, and attribute weights capabilities to create the most effective anonymisation configuration, balancing privacy and preserving predictive power. The key parameters we adjusted were the kappa value, suppression limit, attribute weights, and user feature generalisation hierarchy, the latter two considering the $CHR_{advantage}$ results.

V. Results and Analysis

We generated multiple anonymisation configurations and evaluated their utility by comparing the risk level, number of records suppressed, and overall privacy-utility trade-offs. Those configurations, along with the risk metrics, are shown in Table II.

TABLE II
ANONYMISATION CONFIGURATION AND RISK METRICS

Kappa	ARX transformation	Records supressed	Highest risk	CHR mean at CS 100
7	0-0-1-5	66 (1.09%)	14.28%	28.56
7	0-0-2-4	48 (0.8%)	14.28%	28.59
7	0-1-1-5	66 (1.09%)	14.28%	28.72
14	0-0-1-5	203 (3.36%)	7.14%	28.83
14	0-0-2-4	266 (4.40%)	7.14%	28.64
14	0-2-1-5	100 (1.66%)	7.14%	28.70
21	0-0-1-5	396 (6.56%)	4.76%	28.72
21	0-0-24	543 (8.99%)	4.76%	28.72
21	0-2-1-5	235 (3.89%)	4.16%	28.75

The data indicates that varying ARX configurations impact suppression rates and risk outputs. Various anonymisation models and configurations were considered for the user data appended to the baseline dataset. An optimal balance is observed at kappa = 14 with transformation 0-2-1-5, which suppresses only 1.66% of records and achieves a relatively low risk of 7.14%. Nevertheless, the final row of the table, with kappa = 21 and transformation 0-2-1-5, suggests a configuration with slightly higher suppression yet lower risk. This is the preferred solution. The transformation indices indicate the generalisation levels for gender (0-1), age (0-4), occupation (0-2), and zip code (0-5) features, respectively. Gender was not generalised (level 0), with original values kept. Age had a moderate level (2) of abstraction. A simple level of abstraction (level 1) was used for occupation, and zip code had a higher degree of abstraction (level 5). This configuration ensures that each user remains indistinguishable within a group of at least 21 individuals (the Kappa value) by generalising attributes such as age, occupation, and zip code while preserving the original gender values.

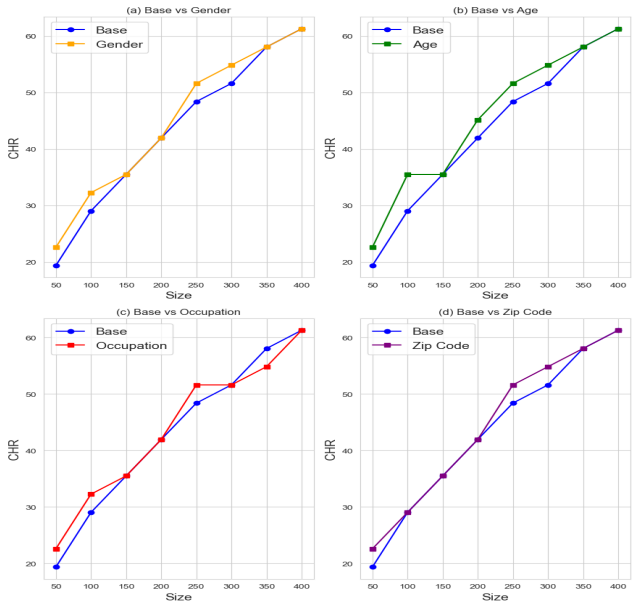

Fig. 2. Base x Feature Selection

A. Prediction with Feature Anonymisation

The CHR is the central metric observed to evaluate the model's performance. It is contrasted with defined parameters, like vehicle density, cache size, training time, and communication rounds. Aside from the cache hit ratio, the model's performance was evaluated using the training loss and the global accuracy computed as a mean square error or MSE.

The proposed AAE model exhibited strong performance in predicting popular content on vehicular networks, achieving an accuracy of 94.74% on the validation set. The training process was conducted across 30 distinct random seed initialisations, with each one executed over 3 rounds, 20 epochs, and 50 clients. On average, this training process took 467.13 seconds and converged to a global loss of 4.711×10^{-2} and a validation loss of 2.1985×10^{-1}. A 95% confidence interval

[3]https://arx.deidentifier.org

TABLE III
DETAILED FEATURE SELECTION TRAINING: CACHE HIT RATIO

Input	CS 50	CS 100	CS 150	CS 200	CS 250	CS 300	CS 350	CS 400	$CHR_{advantage}$
Base	19.35	29.03	35.48	41.94	48.39	51.61	58.06	61.29	-
Base plus Gender	22.58	32.26	35.48	41.94	51.61	54.84	58.06	61.29	5.09
Base plus Age	22.58	35.48	35.48	45.16	51.61	54.84	58.06	61.29	7.44
Base plus Occupation	22.58	32.26	35.48	41.94	51.61	51.61	54.84	61.29	3.62
Base plus Zip Code	22.58	29.03	35.48	41.94	51.61	54.84	58.06	61.29	3.70

TABLE IV
DETAILED FEATURE SELECTION TRAINING: PERFORMANCE METRICS

Input	Global Accuracy (MSE)	Validation Loss	Validation Accuracy
Base	0.0473	0.1831	0.9481
Base plus Gender	0.0477	0.1926	0.9482
Base plus Age	0.0478	0.1924	0.9482
Base plus Occupation	0.0476	0.1904	0.9482
Base plus Zip Code	0.0475	0.1923	0.9482

was employed to quantify variability in the results. The model was evaluated on the test set to predict the top K most popular contents accurately.

Figure 3 presents the CHR against cache size for the baseline dataset (with raw users' feature) and the dataset with anonymised data following our proposed methodology. The figure indicates that the Federated Learning-Adversarial Autoencoder (FL-AAE) is capable of effectively learning patterns with anonymised data and generating precise predictions of popular content while maintaining user privacy. Even considering that the anonymisation process may affect system performance, the Figure 3 outcomes indicate that the model utility performs slightly better than the baseline model, as visible in the close alignment between the CHR curves for anonymised and baseline datasets across different cache sizes. However, the error bars for both datasets overlap slightly, suggesting that while the difference in performance is statistically significant, the margin of improvement is not very large. The anonymised data approach improves the base model without sacrificing performance significantly due to privacy measures. Using anonymisation appears effective in maintaining or slightly improving the CHR, even with potential data modifications for privacy. In particular, for the range of cache sizes, the baseline exhibits an average CHR of 39.03% against and an average CHR of 41.36% for the anonymised data.

The results highlight the effectiveness of the AAE in predicting popular content while preserving user privacy through data anonymisation. The anonymity of the user data is crucial for the model's performance. By preserving user privacy, the model can focus on learning the relevant patterns without being affected by potentially sensitive personal information. Therefore, the anonymised data solution is a viable alternative to the base method, providing comparable or better CHR results across all cache sizes while preserving user privacy. The confidence intervals suggest reliable performance, making the anonymised approach suitable for practical deployment.

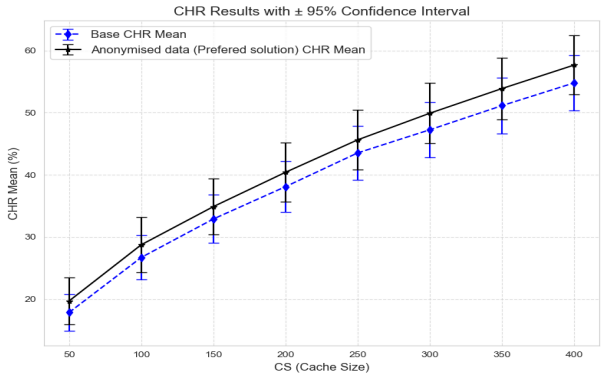

Fig. 3. Variation of the CHR with the cache size for the baseline (non-anonymised solution) in comparison with the anonymised solution (Preferred solution), as determined in section IV.

These results show the feasibility of improving content prediction models in real-world vehicular networks while protecting personal data through privacy-preserving techniques.

VI. CONCLUSIONS AND FUTURE WORK

This research introduces a new method for forecasting popular content in vehicular networks. It uses an FL-AAE model combined with anonymisation techniques to enhance the data available for prediction while protecting user privacy. The key contributions are: it introduces a model that can accurately predict popular content while preserving user privacy through data anonymisation. It demonstrates that implementing data anonymisation enhances privacy without sacrificing prediction performance. Furthermore, it provides an assessment of the suggested methodology using real-world vehicular network data. It demonstrates that by employing anonymised user data and a Federated Learning approach, the proposed model can effectively forecast popular content in vehicular networks while maintaining high privacy protection.

Potential future research could explore integrating the content prediction model with caching and content delivery systems to attain thorough optimisation.

ACKNOWLEDGMENTS

This work is financed through national funds by FCT—Fundação para a Ciência e a Tecnologia, I.P., in the framework of the Project UIDB/00326/2025 and UIDP/00326/2025 and the Federal Institute of Education, Science, and Technology of Rio Grande do Norte—IFRN through the Institutional Development Plan.

References

[1] L. Liu, C. Chen, Q. Pei, S. Maharjan, and Y. Zhang, "Vehicular Edge Computing and Networking: A Survey," *Mobile Netw Appl*, Jul. 2020. DOI: 10.1007/s11036-020-01624-1.

[2] L. U. Khan, I. Yaqoob, N. H. Tran, S. M. A. Kazmi, T. N. Dang, and C. S. Hong, "Edge-Computing-Enabled Smart Cities: A Comprehensive Survey," *IEEE Internet of Things Journal*, vol. 7, no. 10, pp. 10 200–10 232, Apr. 2020. DOI: 10.1109/JIOT.2020.2987070.

[3] R. Meneguette, R. De Grande, J. Ueyama, G. P. R. Filho, and E. Madeira, "Vehicular Edge Computing: Architecture, Resource Management, Security, and Challenges," *ACM Comput. Surv.*, vol. 55, no. 1, 4:1–4:46, Nov. 2021. DOI: 10.1145/3485129.

[4] Z. Yu, J. Hu, G. Min, H. Xu, and J. Mills, "Proactive Content Caching for Internet-of-Vehicles based on Peer-to-Peer Federated Learning," in *2020 IEEE 26th International Conference on Parallel and Distributed Systems (ICPADS)*, Dec. 2020, pp. 601–608. DOI: 10.1109/ICPADS51040.2020.00083.

[5] B. C. Kim, J. U. Kim, H. Lee, and Y. M. Ro, "Revisiting Role of Autoencoders in Adversarial Settings," in *2020 IEEE International Conference on Image Processing (ICIP)*, Oct. 2020, pp. 1856–1860. DOI: 10.1109/ICIP40778.2020.9191259.

[6] V. Mothukuri, R. M. Parizi, S. Pouriyeh, Y. Huang, A. Dehghantanha, and G. Srivastava, "A survey on security and privacy of federated learning," *Future Generation Computer Systems*, vol. 115, pp. 619–640, Feb. 2021. DOI: 10.1016/j.future.2020.10.007.

[7] L. Lyu, H. Yu, J. Zhao, and Q. Yang, "Threats to Federated Learning," in *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Nov. 2020, pp. 3–16. DOI: 10.1007/978-3-030-63076-8_1.

[8] A. Mahmood, C. E. Casetti, C. F. Chiasserini, P. Giaccone, and J. Härri, "The RICH Prefetching in Edge Caches for In-Order Delivery to Connected Cars," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 1, pp. 4–18, Jan. 2019. DOI: 10.1109/TVT.2018.2879850.

[9] S. Raza, S. Wang, M. Ahmed, and M. R. Anwar, "A Survey on Vehicular Edge Computing: Architecture, Applications, Technical Issues, and Future Directions," *Wireless Communications and Mobile Computing*, vol. 2019, e3159762, Feb. 2019. DOI: 10.1155/2019/3159762.

[10] Z. Yu, J. Hu, G. Min, Z. Zhao, W. Miao, and M. S. Hossain, "Mobility-Aware Proactive Edge Caching for Connected Vehicles Using Federated Learning," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 8, pp. 5341–5351, Aug. 2021. DOI: 10.1109/TITS.2020.3017474.

[11] M. Yasir, S. K. uz Zaman, T. Maqsood, F. Rehman, and S. Mustafa, "CoPUP: Content popularity and user preferences aware content caching framework in mobile edge computing," *Cluster Comput*, vol. 26, no. 1, pp. 267–281, Feb. 2023. DOI: 10.1007/s10586-022-03624-0.

[12] Y. Mao, C. You, J. Zhang, K. Huang, and K. B. Letaief, "A survey on Mobile Edge Computing: The communication perspective," *IEEE Communications Surveys Tutorials*, vol. 19, no. 4, pp. 2322–2358, 2017. DOI: 10.1109/COMST.2017.2745201.

[13] C. Xie, Z. Cao, Y. Long, D. Yang, D. Zhao, and B. Li, *Privacy of Autonomous Vehicles: Risks, Protection Methods, and Future Directions*, Sep. 2022. DOI: 10.48550/arXiv.2209.04022. arXiv: 2209.04022 [cs].

[14] Y. He, D. Zhai, F. Huang, D. Wang, X. Tang, and R. Zhang, "Joint Task Offloading, Resource Allocation, and Security Assurance for Mobile Edge Computing-Enabled UAV-Assisted VANETs," *Remote Sensing*, vol. 13, no. 8, p. 1547, Jan. 2021. DOI: 10.3390/rs13081547.

[15] S. Gharib, M. Tran, D. Luong, K. Drossos, and T. Virtanen, "Adversarial Representation Learning for Robust Privacy Preservation in Audio," *IEEE Open Journal of Signal Processing*, vol. 5, pp. 294–302, 2024. DOI: 10.1109/OJSP.2023.3349113.

[16] G. Gondim-Ribeiro, P. Tabacof, and E. Valle, *Adversarial Attacks on Variational Autoencoders*, Jun. 2018. DOI: 10.48550/arXiv.1806.04646. arXiv: 1806.04646 [cs].

[17] M. S. Jere, T. Farnan, and F. Koushanfar, "A Taxonomy of Attacks on Federated Learning," *IEEE Security & Privacy*, vol. 19, no. 2, pp. 20–28, Mar. 2021. DOI: 10.1109/MSEC.2020.3039941.

[18] J. Zhang, B. Chen, X. Cheng, H. T. T. Binh, and S. Yu, "PoisonGAN: Generative Poisoning Attacks Against Federated Learning in Edge Computing Systems," *IEEE Internet of Things Journal*, vol. 8, no. 5, pp. 3310–3322, Mar. 2021. DOI: 10.1109/JIOT.2020.3023126.

[19] P. Kairouz, H. B. McMahan, B. Avent, *et al.*, "Advances and Open Problems in Federated Learning," *MAL*, vol. 14, no. 1-2, pp. 1–210, Jun. 2021. DOI: 10.1561/2200000083.

[20] O. Choudhury, A. Gkoulalas-Divanis, T. Salonidis, *et al.*, *Anonymizing Data for Privacy-Preserving Federated Learning*, Feb. 2020. arXiv: 2002.09096 [cs].

[21] F. Yuan, L. Yao, and B. Benatallah, "Exploring Missing Interactions: A Convolutional Generative Adversarial Network for Collaborative Filtering," in *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, Virtual Event Ireland: ACM, Oct. 2020, pp. 1773–1782. DOI: 10.1145/3340531.3411917.

[22] F. M. Harper and J. A. Konstan, "The MovieLens Datasets: History and Context," *ACM Trans. Interact. Intell. Syst.*, vol. 5, no. 4, pp. 1–19, Jan. 2016. DOI: 10.1145/2827872.