

Heterogeneity: An Open Challenge for Federated On-board Machine Learning

Maria Hartmann^{*1}, Grégoire Danoy^{1,2}, Pascal Bouvry²

¹*SnT, University of Luxembourg, Luxembourg*

²*FSTM, University of Luxembourg, Luxembourg*

The design of satellite missions is currently undergoing a paradigm shift from the historical approach of individualised monolithic satellites towards distributed mission configurations, consisting of multiple small satellites. With a rapidly growing number of such satellites now deployed in orbit, each collecting large amounts of data, interest in on-board orbital edge computing is rising. Federated Learning is a promising distributed computing approach in this context, allowing multiple satellites to collaborate efficiently in training on-board machine learning models. Though recent works on the use of Federated Learning in orbital edge computing have focused largely on homogeneous satellite constellations, Federated Learning could also be employed to allow heterogeneous satellites to form ad-hoc collaborations, e.g. in the case of communications satellites operated by different providers. Such an application presents additional challenges to the Federated Learning paradigm, arising largely from the heterogeneity of such a system. In this position paper, we offer a systematic review of these challenges in the context of the cross-provider use case, giving a brief overview of the state-of-the-art for each, and providing an entry point for deeper exploration of each issue.

1 Introduction

With advances in hardware and software capabilities, distributed satellite mission configurations are progressively replacing the classical paradigm of using a single monolithic spacecraft. With nanosatellites able to generate and store increasingly large amounts of data through various on-board sensors, downlink capacity is becoming a major bottleneck in processing the gathered information. To manage this problem, there is an ongoing drive towards shifting data processing onto satellites[1] – this strategy is referred to as Orbital Edge Computing (OEC)[2]. The overarching idea of OEC is to leverage on-board computing capabilities of each satellite to process locally gathered data, reducing the size and amount of required trans-

missions and speeding up evaluation. A promising variant of OEC proposes deploying Federated Learning [3] (FL) on satellites, allowing the joint training of on-board machine learning models across the data gathered by multiple satellites with a limited communication budget [4]. Under a FL scheme, each satellite performs on-board machine learning on the data it collects, training a local model – see Figure 1. These models are shared periodically among participants, allowing them to be aggregated into a more accurate global model on which to continue training. Aggregation can take place with the aid of a parameter server on the ground or in orbit, or in a fully decentralised manner between satellites. Fundamental advantages of this approach include a vastly reduced communication cost compared to the transmission of raw data, and the inherent privacy advantages of compartmentalising data on satellites.

Current literature on the use of Federated Learning in Orbital Edge Computing is focused primarily on a single use case: using Federated Learning in a single, dedicated constellation of satellites. However, another frequently occurring scenario appears largely unstudied: the potential for satellites from different missions and providers to form (ad-hoc) cross-provider collaborations. In this position paper, we offer an initial exploration of the conceptual challenges associated with this use case: we identify the characteristics of the problem, present a brief survey of the state of the art for each, connecting existing research from the application domain and the theoretical field, and discuss how existing approaches might fare in this scenario. We conclude with a gap analysis.

2 Survey results

In this section, we analyse the current state of the art in research relating to our use case of cross-provider FL. As this use case has not yet been addressed explicitly, we divide our analysis into different thematic

^{*}Corresponding author. E-Mail: {firstname.lastname}@uni.lu

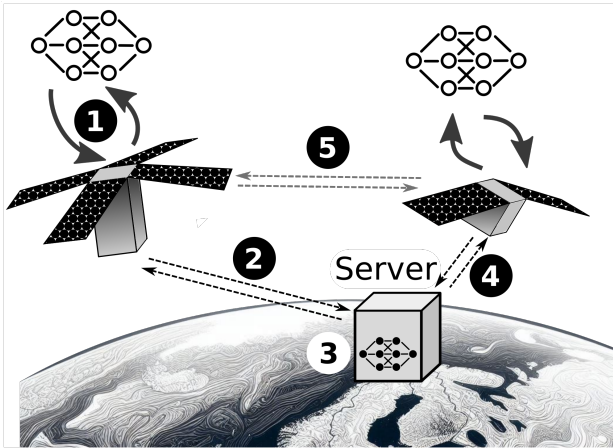


Figure 1: In Federated Learning, each satellite performs on-board machine learning to train a local model (1). Only these models are transmitted via satellite link to a server (2), here based on the ground, where multiple local models are aggregated into a single global model (3). This global model is transmitted back to the satellites (4) to continue the learning process. If necessary, satellites can act as relays for one another (5).

sections. We begin by considering the research closest to application, considering orbital edge computing (OEC) and Federated learning (FL) schemes tailored to use on satellites. Following this, we discuss cross-provider FL, an OEC use case that has, to the best of our knowledge, not been addressed to date. We explore related research from the field of FL that could be applicable for this use case, particularly works addressing the handling of different types of heterogeneity, and discuss their applicability.

2.1 Orbital edge computing and federated learning

Federated Learning could offer a flexible framework for satellites to collaborate on on-board information processing[1][5] while limiting communication cost and preserving data privacy. Various works modify the FL paradigm for the use case of LEO constellations, mainly focusing on adapted communication schedules to handle the intermittent connectivity of satellites. These approaches can broadly be divided by their proposed placement of a parameter server.

Initial works focused on the use of a ground-based server, offering a relatively higher resource capacity of the server; the drawback is a communication bottleneck caused by the intermittent connectivity of satellites. The FedSpace algorithm [6] attempts to overcome this constraint by performing semi-asynchronous federated aggregation, exploiting knowledge about clients’ orbital periods to calculate an aggregation schedule that yields an optimal trade-off between satellite idleness and model staleness. The

FedGSM algorithm[7] similarly makes use of known connectivity intervals to extrapolate model updates.

Conversely, FedISL[8], a synchronous FL scheme for a dense LEO constellation, hinges on the strategic placement of a server in medium Earth orbit (MEO); with convergence speed further enhanced by the use of intra-plane inter-satellite links. This concept is extended in [9] with the grouping of satellites sharing the same orbit to speed up aggregation. Similarly, the synchronous FedHAP [10] algorithm is based on the deployment of multiple high-altitude aerial platforms to accelerate aggregation; AsyncFLEO [11] rests on the same premise, but is capable of asynchronous aggregation. The DSFL[12] algorithm side-steps the challenge of server placement by performing fully decentralised aggregation. Finally, in their work on semi-supervised FL, Östman et al. [13] compare a decentralised aggregation strategy with two scenarios where a relay satellite and a set of ground stations, respectively, act as the aggregation server. All three variants are shown to achieve comparable accuracy performances with similar total training time and power consumption.

Note that the works presented in this section do not consider heterogeneity challenges in great depth; several works, e.g. [6][9], claim that any suitable FL algorithm could be utilised as a drop-in component. In the following section, we assess the additional characteristics that might be required of an algorithm to mitigate heterogeneity in the cross-provider use case.

2.2 Heterogeneity challenges of cross-provider FL

With the proliferation of private and commercial missions, an ever-increasing number of satellites with different capabilities and overlapping interests are active in Earth orbit. Enabling an ad-hoc collaboration between satellites of providers with compatible interests could serve to enhance the performance of all sides at a comparatively low communication cost. A natural example is the application to satellite communication problems: machine learning has the potential to assist with various SatComm-related problems[14], and collaboration between satellites could help solve these problems with greater accuracy and reliability. Many such satellites by different service providers are in operation today, with different hardware, different orbits and different underlying purposes, but nevertheless carrying out related functions. Compared to performing FL on single-mission satellite constellations, this application scenario presents unique challenges induced by the heterogeneity of satellites.

This could involve many different types of heterogeneity, known to present a challenge to FL algo-

rithms[15]. These include *data heterogeneity*, *feature heterogeneity*, *device heterogeneity*, and *preference heterogeneity*. In this section, we discuss the state-of-the-art approaches for each of these types, highlighting how each has been addressed for the OEC use case and, where missing, how existing solutions might transfer to this use case.

Data heterogeneity. This type of heterogeneity, where data is imbalanced across participants, is discussed extensively in the literature[15], as it occurs naturally in most settings. In our use case, heterogeneous distributions of data across satellites are quite likely, with the extent dependent in part on the precise setting. For example, satellites gathering Earth observation images might collect significantly different samples based on their orbital planes, while for SatComm data might differ based on the role of the satellite or the associated service provider. The general issue of data heterogeneity is discussed in most FL variants proposed for the OEC use case, e.g. [12], [7]; however, their effectiveness is seldom demonstrated beyond preliminary benchmarking experiments. Therefore, it appears worthwhile to also consider the state of the art in the general field of FL. A taxonomy of variants of data heterogeneity is presented in [16], along with a comprehensive survey of current mitigation approaches. According to [16], these can be broadly divided into data-level, model-level and server-level interventions. Data-level approaches involve modifying the underlying training data to balance heterogeneity, e.g. by preprocessing [17][18], generating supplemental data using Generative Adversarial Networks [19], or transmitting information about data between clients [20]. However, these strategies often place a significant additional computing or communication burden on the clients, rendering them unattractive use on satellites. Selected model- and server-level strategies appear more promising, as they either require little additional computation cost, or can be carried out on the server-side. Notably, these include model regularisation [21], knowledge distillation [22], and personalised federated learning (PFL) approaches [23] such as client clustering[24][25], parameter decoupling [26] and model interpolation [27].

It is difficult to single out an optimal approach for the general version of our use case, where the data distribution pattern is unknown. The most promising approach would likely be an adaptive solution that modifies the aggregation approach during runtime based on observed metrics, e.g. clustering participants by similarity [24][25] or assigning importance weights for aggregation [27]. For mission architectures involving a powerful ground-based parameter server, more complex knowledge distillation-based approaches may also be an option.

Finally, we note that if the nature of the data distribution is known, such as in EO imaging missions, this could be exploited to the advantage of the algorithm, e.g. by grouping participants known to collect similar data, or conversely by exchanging small sets of selected samples to balance highly different datasets.

Device heterogeneity. Aside from the communication challenges induced by orbital trajectories, satellites in the cross-provider setting would also have hardware differences, leading to different levels of noise and training datasets of varying quality, and impacting computational speeds and capabilities. This is a common problem in the general field of FL [16]; standard approaches include adaptively assigning different weight contributions [28] or model architectures [29] to participants; possibly also reducing the consideration of lower-quality participants in selecting clients for aggregation [30][31]. It appears likely that such strategies would transfer well to the present use case, without a need for major modifications.

Feature heterogeneity. This setting corresponds to the collaboration of satellites with different types of sensors, collecting different types of data and potentially requiring different model architectures for on-board processing. Effectively integrating different features and models into a coherent federated model training process presents a difficult problem; to the best of our knowledge, it has not yet been tackled in research on OEC. Indeed, the general problem of performing FL in such a setting, known as Vertical Federated Learning (VFL) [32], remains largely unsolved beyond highly constrained artificial scenarios or costly compensation approaches [33]. This present lack of solutions renders the application of FL to the real-world satellite use case infeasible. The only approach that appears viable at present would involve feature distillation on a ground-station parameter server – a computationally expensive solution that does not yield workable models for the federated participants, and so would benefit only the server-side [33]. Beyond this niche variant, the currently most feasible approach would likely be to separate participants by features, eliminating this type of heterogeneity.

Preference heterogeneity. This is a novel type of heterogeneity that, as of now, has seen little recognition in the field of federated learning; yet it appears highly relevant to the present use case. Preference heterogeneity arises when participants are solving problems with multiple objectives, e.g. minimising communication cost while also minimising connection latency in a satellite communication network. In such multi-objective problems, different optimal solutions are generally possible, representing different trade-offs between the individual objectives. In practice, some trade-offs may be more desirable than others,

e.g. conserving energy by limiting communication for severely resource-constrained satellites, or minimising connection latency to boost service to certain geographical areas. This can be controlled by assigning importance weights to each objective when solving the learning problem. We call participants with different importance weights preference-heterogeneous.

At present, little research has been devoted to performing Federated Learning under preference heterogeneity; the closest related works consider federated multi-task learning [34], where participants have fully separate objectives, and federated multi-objective learning [35][36] without allowance for different preferences.

2.3 Other considerations

2.3.1 Fairness

In addition to the technical considerations of the previous section, this cross-provider use case also differs from the single-provider variant in the assumptions made about participants' intentions. Satellites designed to collaborate with each other as part of a single mission can generally be assumed to act altruistically in federation, i.e. towards the benefit of the larger system. In the cross-provider setting, however, no such assumption should be made, as has been noted e.g. in [9]. Instead, we assume that participants may act 'selfishly', valuing their own success over that of the federated system. This could occur e.g. if satellites contribute low-quality updates, leading to an overall degradation of the global model and benefiting from the collaboration at the cost of others. Similarly, a participant could limit the frequency of its contributions to conserve communication budget or maintain privacy, to the detriment of others in the federation, while still receiving global model updates. A successful cross-provider collaboration scheme should guard against such exploitation. These fairness considerations have yet to be addressed in works targeting the OEC use case; a full survey on the state of the art of general FL approaches is provided in [37]. The same work gives a detailed account of the different definitions of fairness and underlying assumptions; the choice of an appropriate mitigation approach for the present depends on these characteristics.

2.3.2 Guarding against malicious participants

This is a more extreme case of the challenges discussed in the previous section – here participants intentionally attempt to sabotage the performance of the federated system through their participation, e.g. through submitting intentionally false model up-

dates (known as model poisoning attacks). A thorough overview of possible attack vectors and approaches for guarding against such attacks is given in [15] and more recently [38], suggesting e.g. the assignment of confidence scores [39], filtering outliers [40], or normalising model updates before aggregation [41]. A number of these solutions appear to transfer well to the cross-provider OEC use case, given a trustworthy server, with the selection of approach dependent on the particular parameters of the system and the results of a risk analysis.

2.3.3 The role of standardisation

Along with the algorithmic approaches towards preventing misuse of the collaboration, standardisation likely has an important role to play in the deployment of FL to the present use case involving multiple providers. This could e.g. include a requirement for certification of machine learning pipelines in accordance with certain quality standards to obtain access to such federated exchange schemes, to decrease the likelihood of interference by malicious or poorly engineered participants. A review of standards relating to the trustworthiness of machine learning for space applications suggests that such standards largely have yet to be defined [42]; a recently published handbook provides a first glimpse of such considerations [43].

Finally, we note another crucial challenge of this particular use case: the need for a unified communication protocol between participants, capable of negotiating the parameters of the FL scheme between participants and able to transmit machine learning models unambiguously, without interpretation errors caused e.g. by differences in hardware or software. For ad-hoc collaboration, standardised communication formats are of critical importance, both to negotiate the parameters of the FL protocol during the initialisation phase, and to transmit model updates without error. There appears to be no existing solution, as most works in the literature tend to assume a group of satellites collaborating as part of a single unified mission.

3 Discussion

In our analysis of the state of the art, we have seen that current works investigating the use of Federated Learning in satellite on-board edge computing have largely focused on a single scenario: a large constellation of homogeneous satellites deployed as part of a single mission. We have introduced a different use case, relevant to the present or near future, where heterogeneous satellites of multiple different providers could collaborate under a FL scheme to enhance on-

board learning. We have elucidated the unique conceptual challenges of this use case, with a focus on the different types of heterogeneity and conflicts of interest that might arise. We have provided a broad perspective of the state of the art for each, discussing both the existing work close to the use case, and the general state of the art in the theoretical literature. Our brief survey shows that several aspects remain to be addressed to adequately solve this real-world use case. In particular, there is a need to further investigate (1) how state-of-the-art solutions can be combined in settings where multiple types of heterogeneity occur simultaneously; (2) which heterogeneity-mitigating algorithms could be selected to fit with the use case, and with existing OEC schemes; (3) how to perform Federated Learning under preference heterogeneity. Finally, we suggest that any engineering approach should be supported by additional measures reducing the complexity of the system by providing appropriate constraints, e.g. through the use of standardisation.

Acknowledgements

This work is partially funded by the joint research programme UL/SnT-ILNAS on Technical Standardisation for Trustworthy ICT, Aerospace, and Construction.

References

- Izzo, D., Meoni, G., Gómez, P., Dold, D. & Zochbauer, A. in *Artificial Intelligence for Space: AI4SPACE* 21–52 (CRC Press, 2022).
- Denby, B. & Lucia, B. *Orbital Edge Computing: Nanosatellite Constellations as a New Class of Computer System* in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems* (ACM, Mar. 2020). <http://dx.doi.org/10.1145/3373376.3378473>.
- McMahan, B., Moore, E., Ramage, D., Hampson, S. & y Arcas, B. A. *Communication-Efficient Learning of Deep Networks from Decentralized Data* in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA* (eds Singh, A. & Zhu, X.) **54** (PMLR, 2017), 1273–1282. <http://proceedings.mlr.press/v54/mcmahan17a.html>.
- Jabbarpour, M. R. *et al. Performance Analysis of Federated Learning in Orbital Edge Computing* in *Proceedings of the IEEE/ACM 16th International Conference on Utility and Cloud Computing* (ACM, Dec. 2023). <http://dx.doi.org/10.1145/3603166.3632140>.
- Chen, H., Xiao, M. & Pang, Z. *Satellite-Based Computing Networks with Federated Learning*. *IEEE Wireless Communications* **29**, 78–84. issn: 1558-0687. <http://dx.doi.org/10.1109/MWC.008.00353> (Feb. 2022).
- So, J. *et al. FedSpace: An Efficient Federated Learning Framework at Satellites and Ground Stations* 2022. <https://arxiv.org/abs/2202.01267>.
- Wu, L. & Zhang, J. *FedGSM: Efficient Federated Learning for LEO Constellations with Gradient Staleness Mitigation* in *2023 IEEE 24th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)* (IEEE, Sept. 2023). <http://dx.doi.org/10.1109/SPAWC53906.2023.10304545>.
- Razmi, N., Matthiesen, B., Dekorsy, A. & Popovski, P. *On-Board Federated Learning for Dense LEO Constellations* in *ICC 2022 - IEEE International Conference on Communications* (IEEE, May 2022). <http://dx.doi.org/10.1109/ICC45855.2022.9838619>.
- Razmi, N., Matthiesen, B., Dekorsy, A. & Popovski, P. *On-board Federated Learning for Satellite Clusters with Inter-Satellite Links*. *IEEE Transactions on Communications*, 1–1. issn: 1558-0857. <http://dx.doi.org/10.1109/TCOMM.2024.3356429> (2024).
- Elmahallawy, M. & Luo, T. *FedHAP: Fast Federated Learning for LEO Constellations using Collaborative HAPs* in *2022 14th International Conference on Wireless Communications and Signal Processing (WCSP)* (IEEE, Nov. 2022). <http://dx.doi.org/10.1109/WCSP55476.2022.10039157>.
- Elmahallawy, M. & Luo, T. *AsyncFLEO: Asynchronous Federated Learning for LEO Satellite Constellations with High-Altitude Platforms* in *2022 IEEE International Conference on Big Data (Big Data)* (IEEE, Dec. 2022). <http://dx.doi.org/10.1109/BigData55660.2022.10021101>.
- Wu, C., Zhu, Y. & Wang, F. *DSFL: Decentralized Satellite Federated Learning for Energy-Aware LEO Constellation Computing* in *2022 IEEE International Conference on Satellite Computing (Satellite)* (IEEE, Nov. 2022). <http://dx.doi.org/10.1109/Satellite55519.2022.00013>.
- Östman, J., Gomez, P., Shreenath, V. M. & Meoni, G. *Decentralised Semi-supervised Onboard Learning for Scene Classification in Low-Earth Orbit* 2023. <https://arxiv.org/abs/2305.04059>.
- Fourati, F. & Alouini, M.-S. *Artificial intelligence for satellite communication: A review*. *Intelligent and Converged Networks* **2**, 213–243. issn: 2708-6240. <http://dx.doi.org/10.23919/ICN.2021.0015> (Sept. 2021).
- Kairouz, P. *et al. Advances and Open Problems in Federated Learning*. *Foundations and Trends® in Machine Learning* **14**, 1–210. issn: 1935-8245. <http://dx.doi.org/10.1561/22000000083> (2021).
- Ye, M., Fang, X., Du, B., Yuen, P. C. & Tao, D. *Heterogeneous Federated Learning: State-of-the-art and Research Challenges*. *ACM Computing Surveys* **56**, 1–44. issn: 1557-7341. <http://dx.doi.org/10.1145/3625558> (Oct. 2023).
- Li, A. *et al. Sample-level Data Selection for Federated Learning* in *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications* (IEEE, May 2021). <http://dx.doi.org/10.1109/INFOCOM42981.2021.9488723>.
- Xu, X., Li, H., Li, Z. & Zhou, X. *Safe: Synergic Data Filtering for Federated Learning in Cloud-Edge Computing*. *IEEE Transactions on Industrial Informatics* **19**, 1655–1665. issn: 1941-0050. <http://dx.doi.org/10.1109/TII.2022.3195896> (Feb. 2023).
- Goodfellow, I. *et al. Generative adversarial networks*. *Communications of the ACM* **63**, 139–144. issn: 1557-7317. <http://dx.doi.org/10.1145/3422622> (Oct. 2020).
- Yoon, T., Shin, S., Hwang, S. J. & Yang, E. *FedMix: Approximation of Mixup under Mean Augmented Federated Learning*. *ArXiv abs/2107.00233*. <https://api.semanticscholar.org/CorpusID:235614375> (2021).

21. Kim, J., Kim, G. & Han, B. Multi-Level Branched Regularization for Federated Learning. *ArXiv* **abs/2207.06936**. <https://api.semanticscholar.org/CorpusID:250340586> (2022).
22. Zhu, Z., Hong, J. & Zhou, J. Data-Free Knowledge Distillation for Heterogeneous Federated Learning. *Proceedings of machine learning research* **139**, 12878–12889 (July 2021).
23. Tan, A. Z., Yu, H., Cui, L. & Yang, Q. Towards Personalized Federated Learning. *IEEE Transactions on Neural Networks and Learning Systems* **34**, 9587–9603. ISSN: 2162-2388. <http://dx.doi.org/10.1109/TNNLS.2022.3160699> (Dec. 2023).
24. Ghosh, A., Chung, J., Yin, D. & Ramchandran, K. An Efficient Framework for Clustered Federated Learning. *IEEE Transactions on Information Theory* **68**, 8076–8091. ISSN: 1557-9654. <http://dx.doi.org/10.1109/TIT.2022.3192506> (Dec. 2022).
25. Duan, M. et al. *FedGroup: Efficient Federated Learning via Decomposed Similarity-Based Clustering* in 2021 IEEE Intl Conf on Parallel & Distributed Processing with Applications, Big Data & Cloud Computing, Sustainable Computing & Communications, Social Computing & Networking (ISPA/BDCLOUD/SocialCom/SustainCom) (IEEE, Sept. 2021). <http://dx.doi.org/10.1109/ISPA-BDCLOUD-SocialCom-SustainCom52081.2021.00042>.
26. Arivazhagan, M. G., Aggarwal, V., Singh, A. K. & Choudhary, S. *Federated Learning with Personalization Layers* 2019. <https://arxiv.org/abs/1912.00818>.
27. Hanzely, F. & Richtárik, P. Federated Learning of a Mixture of Global and Local Models. *ArXiv* **abs/2002.05516**. <https://api.semanticscholar.org/CorpusID:211096943> (2020).
28. Ma, X., Zhang, J., Guo, S. & Xu, W. *Layer-wised Model Aggregation for Personalized Federated Learning* in 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (IEEE, June 2022). <http://dx.doi.org/10.1109/CVPR52688.2022.00985>.
29. Diao, E., Ding, J. & Tarokh, V. HeteroFL: Computation and Communication Efficient Federated Learning for Heterogeneous Clients. *ArXiv* **abs/2010.01264**. <https://api.semanticscholar.org/CorpusID:222133374> (2020).
30. Li, L. et al. *FedSAE: A Novel Self-Adaptive Federated Learning Framework in Heterogeneous Systems* in 2021 International Joint Conference on Neural Networks (IJCNN) (IEEE, July 2021). <http://dx.doi.org/10.1109/IJCNN52387.2021.9533876>.
31. Nishio, T. & Yonetani, R. *Client Selection for Federated Learning with Heterogeneous Resources in Mobile Edge* in ICC 2019 - 2019 IEEE International Conference on Communications (ICC) (IEEE, May 2019). <http://dx.doi.org/10.1109/ICC.2019.8761315>.
32. Yang, Q., Liu, Y., Chen, T. & Tong, Y. Federated Machine Learning: Concept and Applications. *ACM Transactions on Intelligent Systems and Technology* **10**, 1–19. ISSN: 2157-6912. <http://dx.doi.org/10.1145/3298981> (Jan. 2019).
33. Liu, Y. et al. Vertical Federated Learning: Concepts, Advances, and Challenges. *IEEE Transactions on Knowledge and Data Engineering*, 1–20. ISSN: 2326-3865. <http://dx.doi.org/10.1109/TKDE.2024.3352628> (2024).
34. Smith, V., Chiang, C.-K., Sanjabi, M. & Talwalkar, A. S. *Federated Multi-Task Learning* in *Advances in Neural Information Processing Systems* (eds Guyon, I. et al.) **30** (Curran Associates, Inc., 2017). https://proceedings.neurips.cc/paper_files/paper/2017/file/6211080fa89981f66b1a0c9d55c61d0f-Paper.pdf.
35. Yang, H., Liu, Z., Liu, J., Dong, C. & Momma, M. *Federated Multi-Objective Learning* in *Advances in Neural Information Processing Systems* (eds Oh, A. et al.) **36** (Curran Associates, Inc., 2023), 39602–39625. https://proceedings.neurips.cc/paper_files/paper/2023/file/7cb2c2a8d35576c00078b6591ec26a7d-Paper-Conference.pdf.
36. Hartmann, M., Danoy, G., Alswaitti, M. & Bouvry, P. *MOFL/D: A Federated Multi-objective Learning Framework with Decomposition* in *International Workshop on Federated Learning in the Age of Foundation Models in Conjunction with NeurIPS 2023* (2023). <https://openreview.net/forum?id=Pj6BPHZy56>.
37. Shi, Y., Yu, H. & Leung, C. Towards Fairness-Aware Federated Learning. *IEEE Transactions on Neural Networks and Learning Systems*, 1–17. ISSN: 2162-2388. <http://dx.doi.org/10.1109/TNNLS.2023.3263594> (2024).
38. Rodríguez-Barroso, N., Jiménez-López, D., Luzón, M. V., Herrera, F. & Martínez-Cámara, E. Survey on federated learning threats: Concepts, taxonomy on attacks and defences, experimental study and challenges. *Information Fusion* **90**, 148–173. ISSN: 1566-2535. <http://dx.doi.org/10.1016/j.inffus.2022.09.011> (Feb. 2023).
39. Muñoz-González, L., Co, K. T. & Lupu, E. C. Byzantine-Robust Federated Machine Learning through Adaptive Model Averaging. *ArXiv* **abs/1909.05125**. <https://api.semanticscholar.org/CorpusID:202558449> (2019).
40. Yin, D., Chen, Y., Kannan, R. & Bartlett, P. *Byzantine-Robust Distributed Learning: Towards Optimal Statistical Rates* in *Proceedings of the 35th International Conference on Machine Learning* (eds Dy, J. & Krause, A.) **80** (PMLR, Oct. 2018), 5650–5659. <https://proceedings.mlr.press/v80/yin18a.html>.
41. Sun, Z., Kairouz, P., Suresh, A. T. & McMahan, H. B. *Can You Really Backdoor Federated Learning?* 2019. <https://arxiv.org/abs/1911.07963>.
42. Reiff, J.-M. et al. *Trustworthiness in ICT, Aerospace, and Construction applications* tech. rep. (Luxembourg Institute of Standardisation, Accreditation, Safety, Quality of Products, and Services (ILNAS), 2023).
43. ECSS. *ECSS-E-HB-40-02A: Machine Learning Qualification for Space Applications Handbook* tech. rep. (European Cooperation for Space Standardization (ECSS), 2023).