

A Federated Learning-enabled Smart Street Light Monitoring Application: Benefits and Future Challenges

Diya Anand^{1,2}, Ioannis Mavromatis², Pietro Carnelli², and Aftab Khan²

¹University of Bristol, Bristol, UK

²Bristol Research and Innovation Laboratory, Toshiba Europe Limited, Bristol, UK
{Ioannis.Mavromatis,Pietro.Carnelli,Aftab.Khan}@toshiba-bril.com

ABSTRACT

Data-enabled cities are recently accelerated and enhanced with automated learning for improved Smart Cities applications. In the context of an Internet of Things (IoT) ecosystem, the data communication is frequently costly, inefficient, not scalable and lacks security. Federated Learning (FL) plays a pivotal role in providing privacy-preserving and communication efficient Machine Learning (ML) frameworks. In this paper we evaluate the feasibility of FL in the context of a Smart Cities Street Light Monitoring application. FL is evaluated against benchmarks of centralised and (fully) personalised machine learning techniques for the classification task of the lampposts operation. Incorporating FL in such a scenario shows minimal performance reduction in terms of the classification task, but huge improvements in the communication cost and the privacy preserving. These outcomes strengthen FL's viability and potential for IoT applications.

CCS CONCEPTS

• **Computer systems organization** → **Client-server architectures; Sensor networks**; • **Computing methodologies** → **Supervised learning by classification; Neural networks**; • **Information systems** → *Data analytics*.

KEYWORDS

Smart Cities, IoT, Infrastructure, Monitoring, Lamppost, Neural Networks, Federated Learning

ACM Reference Format:

Diya Anand^{1,2}, Ioannis Mavromatis², Pietro Carnelli², and Aftab Khan². 2022. A Federated Learning-enabled Smart Street Light Monitoring Application: Benefits and Future Challenges. In *1st ACM Workshop on AI Empowered Mobile and Wireless Sensing (MORSE '22)*, October 21, 2022, Sydney, NSW, Australia. ACM, Linz, AT, 6 pages. <https://doi.org/10.1145/3556558.3558580>

1 INTRODUCTION

A Smart City is described in [12] as an urban medium using Information and Communication Technologies (ICT) to promote more efficient ordinary city operations and improve the Quality of Services (QoS) received by the citizens. The objective of a Smart City

is to enhance the Quality of Life (QoL) of the citizens and improve sustainability. This is achieved by promoting the digitisation of services, automation, and the use of data for intelligent responses and decisions, while autonomously adapting to different needs [18].

The realm of Smart Cities covers multiple areas and applications, such as Smart Transportation, Smart Urban Management, Smart Tourism, Green Cities, Smart Healthcare, etc. [19]. The technologies required for these applications are numerous but can be commonly grouped under three categories, i.e., sensing and data collection, intelligent decision support, and exchange of data and decisions between different system entities [21]. All these technologies are part of an Internet of Things (IoT) framework [24] that provides the underlying infrastructure and services for their operation.

Our work focuses on improving resource utilisation within an IoT ecosystem by moving intelligent decision-making closer to the sensors and the "edge". More specifically, we evaluate the feasibility of using Federated Learning (FL) within a Smart Street Light Monitoring application context. This application is a good representation of Smart Cities as it generates a large amount of data, requires increased communication bandwidth for their exchange, and is resource-intensive when classifying whether a lamppost is operational or not.

FL [26] is a branch of Machine Learning (ML) that relies on multiple "clients", e.g., edge devices, to collect and process local data for training an ML model. In turn, the client's ML model parameters are shared with a central FL server for global model aggregation. Once enough updates from clients (sample fraction per round of FL) within the network have been received, a new global model is generated and broadcast to all clients. Such methods of training local models have certain benefits important to IoT and "smart city"-centric networks [26].

IoT networks often build upon Low-power Wide-Area Network (LPWAN) wireless protocols such as LoRaWAN, Zigbee and Bluetooth [4]. For example, a Smart Street Light Monitoring application can easily generate hundreds of gigabytes of data if central training and processing are required [23]. However, this amount of data is almost impossible to be exchanged via the low-power, low data rate IoT wireless links. In such a scenario, FL can play a pivotal role by replacing the exchange of data with the exchange of the more lightweight prediction models.

Furthermore, FL provides a variety of privacy advantages required for real-world Smart Cities applications [11]. Data minimisation is achieved as the raw data stay on the edge device, and data leakage can be avoided in transit. For external users interacting with the system, having access to only aggregated models and data can enhance privacy. Traditional end-to-end encryption mechanisms can secure the models exchanged in transit. Finally, even

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

MORSE '22, October 21, 2022, Sydney, NSW, Australia

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9522-9/22/10...\$15.00

<https://doi.org/10.1145/3556558.3558580>

when an edge device is tampered with, and the model is altered, the nature of FL and the model aggregation on the server limits individual malicious models' influence on the global output. Of course, extensions and algorithms can provide more formal guarantees such as differential privacy [16], or algorithms for concept drift detection can identify such drifts [9].

The rest of the paper is structured as follows; Section 2 summarises the related work to optimising FL clients to their local datasets/environments. Section 3 discusses our methodology, experiments and introduces our evaluation dataset. Our experimental evaluation and results are discussed in Section 4, supported by a detailed discussion on the lessons learned from this activity. Finally, Section 5 summarises our findings, and provides suggestions for future research activities.

2 RELATED WORK

ML is currently being used in various intelligent fault diagnosis methods. For example, [10, 20] present two ML-based fault detection mechanisms for street light applications. Sending the collected data to a central server, the real-time illuminance of the lamppost is evaluated, and faults are reported to the maintenance staff. The results show up to 90% of fault detection accuracy for such scenarios. Such an approach increased the communication cost in the IoT system while introducing many challenges in securing the data in transit.

Personalisation in ML can enhance detection by targeting particular entities and optimising the trained models accordingly. Authors in [17] present three ways of personalisation utilising the MNIST dataset. Their results show increased performance compared to traditional ML strategies. However, personalised models trained on a server require again increased communication overhead. Personalised models trained at the edge, even though they minimise the communication overhead, require a large dataset available for each model trained. The intermittent nature of an IoT system, where data may be scarce, can create obstacles to collecting such vast amounts of data.

FL can bridge the gap between the personalised and centralised approach. Training models locally decrease communication costs while preserving data privacy. Moreover, for edge nodes with an abundance of data, aggregating the existing models on the server side and sharing them with all edge nodes can ensure that a highly accurate model is always available for inference. However, FL can suffer from highly skewed, non-Independent, Identically Distributed (IID) data. Knowing the data types and ways of clustering them can enhance FL's accuracy. Work carried out by [25] shows an improvement of circa 30% on the CIFAR-10 dataset compared to classical FedAvg using their proposed data-sharing strategy between participating clients. The authors show that sharing 5% of a separate global dataset across clients and initialising a model at the server on the dataset mentioned above leads to a classification performance increase. Federated Meta-Learning or FedMeta framework [3] uses parameterised algorithms such as MAML and Meta-SGD to train on the client's local data and communicate the updates to the server instead of updated models in traditional FL. In FedAMP [22], copies of local models are kept on the cloud server with attentive message passing between clients and server leading

to aggregated client models of messages passed. Both methods can again enhance FL's performance.

Whilst such techniques show general performance improvement, the practical challenges are numerous. For example, creating a separate dataset of similar distribution would require prior knowledge of client data and sharing this data with the server. This method is not aligned with FL principles of maintaining local datasets for the participating clients. Furthermore, such data-sharing techniques might not be feasible amongst thousands of sensors/edge devices in an IoT network. Similarly, sharing and storing multiple copies of clients' models requires more robust, scalable and resource-rich infrastructure and storage capabilities. This approach does not scale well in networks of hundreds of thousands of participating devices.

3 METHODOLOGY AND EXPERIMENTS

In this paper we investigate the feasibility of FL for a lamppost fault detection use-case and compare it against a centralised and a "fully" personalised approach. An FL benchmark method, using a typical averaging technique to establish a global model (inspired by [15]), was compared to a centralised method (i.e., classical ML) and an extreme version of or a "hyper-personalised" FL method, whereby each lamppost uses a model trained on **only** its dataset and was never aggregated centrally. Such "extreme" methods provide a good overview of their potential for accurate detection using very different training data splits.

3.1 Convolutional Neural Network Model Selection

For most FL IoT applications, the model should be lightweight enough to train on the edge device. Moreover, FL involves broadcasting the model between clients and the server. Hence a bigger model would not meet the bandwidth limitations of FL and increase communication costs significantly.

A Convolutional Neural Network (CNN) based on the Residual Networks (ResNet) architecture was considered "lightweight" enough for training and inference on IoT devices and capable of classifying incoming images. In particular, ResNets [7] were designed to combat deep neural networks which suffer from vanishing gradients (a common problem encountered in neural network training procedures). Furthermore, such a model allowed for fair comparison amongst the three different ML strategies being investigated.

3.2 The Dataset

Our investigation is based on a large volume dataset of street light images. Since our focus is the evaluation of FL under Smart Street Light Monitoring environments, we selected the "Dataset of Images of Public Streetlights" [13], generated as part of the UMBRELLA project [5]. The dataset is publicly available from Zenodo Open Repository [14]. The dataset consists of over 350,000 images of streetlights, collected hourly and over a period of six months. The images come from 140 UMBRELLA IoT nodes deployed across multiple locations in the South Gloucestershire region of the UK. The UMBRELLA nodes are currently installed at a public stretch of ~7.2 km road (about ~ 80% of the nodes) and around the University of the West of England (UWE) Frenchay Campus (about ~ 20% of the nodes). Since each lamppost had between 1,000 – 4,000




Figure 1: (a) Node 0: Ideal Lamppost. The light is clearly visible and this becomes easy detection of whether the light is on or off (b) Node 1: Normal Lamppost. The light is not directly visible; for detection (c) Node 2: Edge Case Lamppost. the light is completely covered by vegetation or the camera has slipped making for complicated detection.

images, personalised models per lamppost were possible to train and optimise for our comparison, achieving very high accuracy on each model independently.

The images in the dataset were used to determine whether the lamppost is operational or not, i.e., whether the lamppost is switched ON or OFF. The lamppost functionality is monitored during different times (once per hour), with the light expected to be ON at night and OFF during the day, as part of a partnership with the local government to ensure road safety. As “night” is considered the time period from “15-minutes before sunset” until “15-minutes after sunrise” and calculated independently for each day.

The images are in JPEG format with a resolution of 1024×768 pixels. The entries in the dataset are already pre-labelled. What is more, the dataset spans a large geographical area and various different lamppost designs, heights and operational modes. Several streetlights are partially obstructed by vegetation or are outside the Field of View (FoV) of the camera. Finally, the cameras facing the sky are susceptible to weather conditions (e.g., rain, snow, direct sunlight, etc.) that can partially or entirely alter the quality of the images taken. All the above generate “interesting” and unique edge-cases when evaluating FL within the context of a Smart Street Light Monitoring application.

3.3 Node Categorisation

Figure 1 shows some example images of the dataset used. Our evaluation and discussion are based on further grouping the nodes in the three categories seen in Figure 1. As discussed in Section 4, we grouped the nodes into three categories with respect to the Line-of-Sight (LoS) to the lamppost (being inside the camera’s FoV or not) and whether there is any obstruction by vegetation.

More specifically, Figure 1a is labelled as “node type 0”, modelling an ideal lamppost image; it possesses a clear view of the lamppost light, making the binary classification task of whether the light is on or off relatively simple. Figure 1b is an intermediate case labelled as “node type 1”; due to the positioning of the UMBRELLA node, only the pole of the lamppost is visible. For this node type, the “turned-off” and “turned-on” lampposts can be classified by a human looking at the luminance of an image. However, depending on the camera’s position inside the node and the weather conditions,

the classification is not always easy with bare eyes. Finally, Figure 1c, referred to as “node type 2”, represents the most challenging type captured; this subclass consists of images with no view of the lamppost due to vegetation or the camera being mispositioned. The labelling of each node was done manually before the evaluation. For that, we considered the unique characteristics of each node. The labelling is later used during our evaluation process (fed as a CSV file in our algorithm)¹.

3.4 Data Pre-Processing

The original, ‘raw’ lamppost dataset consists of images of 1024×768 Red, Green, Blue (RGB) pixels which are far too big for most CNNs trained on edge devices (such as the Nvidia Jetson Nano). To reduce the computational and memory footprint during training and deployment we reduced the images to 32×32 pixels with three channels (RGB representation) as shown in Figure 2 by resizing, cropping and down-sampling the image using a bilinear interpolation method. Finally we normalised the reduced RGB images by subtracting the mean from each pixel and dividing it by the standard deviation (Figure 2).

4 EXPERIMENTAL EVALUATION

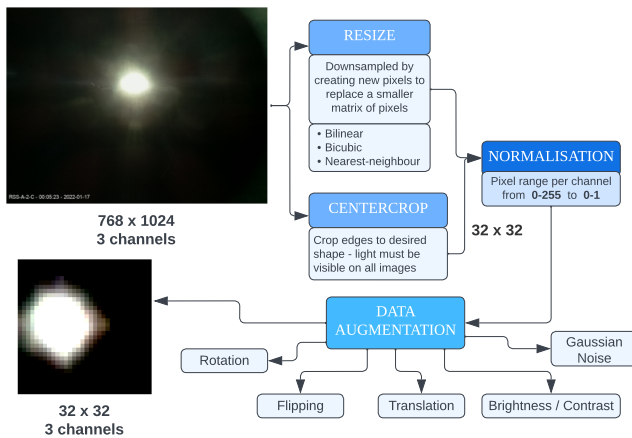
As discussed in Section 3 our evaluation compares a centralised, a “fully” personalised, and an FL approach. The dataset was split into a training and testing dataset; the test set consisted of 20% of images from every lamppost node/device, and the remaining 80% constituted the training dataset. This was consistent throughout all experiments. For our performance investigation, we also considered the type of nodes. We combined grouped the nodes of types 0 and 1 while we separately evaluated the nodes of type 2 to observe how the performance can degrade when facing such edge cases.

The developed centralised model and method generates a single model for the entire lamppost dataset for training. The “fully” personalised method is demonstrated by treating each lamppost as a separate entity. A model is trained solely on its own dataset (i.e., with no FL aggregation). Finally, the FL approach is based on 140

¹A copy of this file can be found in the following link:
<https://www.dropbox.com/s/ydxiooulet3gwf/nodetypes.csv?dl=0>

Table 1: Overview of Experiments

Method	Trained on	#Training Devices	Evaluated on	#Testing Devices	#Models
Personalised	Normal Nodes	133	Normal Nodes	133	133
	Edge-Case Nodes	7	Edge-Case Nodes	7	7
	All nodes or devices	140	All nodes or devices	140	140
Centralised	All nodes or devices	140	Normal Nodes	133	1
	All nodes or devices	140	Edge-Case Nodes	7	1
	All nodes or devices	140	All nodes or devices	140	1
FL benchmark	All nodes or devices	140	Normal Nodes	133	1
	All nodes or devices	140	Edge-Case Nodes	7	1
	All nodes or devices	140	All nodes or devices	140	1

**Figure 2: Flowchart of Data Pre-Processing Pipeline**

clients (i.e., one per lamppost), aggregating their models using the FedAvg algorithm. All experiments are summarised in Table 1.

4.1 Results for all Training/Testing Methods

Our results are summarised in Table 2. Considering the “normal” nodes, as expected, the fully personalised and centralised methods generated higher accuracy (98.57% and 98.41% respectively) and F1-scores (0.99 and 0.988 respectively) than the benchmark FL method (accuracy of 96.7% and F1-score of 0.967). Both methods were allowed to train until fully converged (with minimal overfitting).

The difference is less prominent when the “edge” nodes (node type 2) are considered. Again, personalised and centralised methods slightly outperform FL, but only by a couple of percentage points. The above results are due to having either access to the entire dataset with a single model to train (centralised) or having each client train on a single lamppost’s datasets (fully personalised method).

4.2 Discussion and Observations

Whilst the personalised method fractionally outperformed the centralised and FL training methods, it still achieved a lower accuracy and F1-score than required for immediate deployment. Considering a city-scale deployment, such a system will still produce false positives/negatives at a rate not easily monitored by local government officials. Even at the fairly limited coverage provided within our dataset (140 nodes and lampposts), an error rate of 1.75% will result in tens if not more daily alerts. Given the diversity of the node types and camera locations, more sophisticated classification algorithms are required for better accuracy. For example, taking into account multiple concurrent decisions can reduce the error rate as falsely classified results will be the minority of the reported values.

The “fully” personalised method relies on having a large enough amount of data stored locally on the lamppost edge device for individual training. This constitutes a significant problem when considering the resource-constrained nature of the current IoT devices or when “new” lampposts join the network (as there is no available global model).

Considering the centralised classification, as seen, it performs almost as well as the fully personalised method. However, it relies on having all the lamppost data transferred to a central server for processing and training. Currently, the entire compressed lamppost dataset is circa 150GB in volume. Admittedly, this is without any of our pre-processing and dimensionality reduction techniques applied to the images. However, this is still a good indication of the costly nature of exchanging such a large volume of data. Furthermore, this will also result in privacy and confidentiality issues that may arise with the data being in transit.

The benchmark FL performed the worst out of the three evaluated. This is likely due to the combination and equal weighting during aggregation of the edge case nodes (i.e., type 2). The FL method was several percentage points lower in accuracy and F1-score than the personalised and centralised methods. However, such a method provides not only increased data privacy but also a global model too, ready to be used for initialising training or inference on a new lamppost/edge device joining the network. Therefore, gaining some scalability advantage and drastically reducing the overall data communication volume during training.

Table 2: Experimental Results for all Three Implemented Methods of Training and Testing.

Method	#training lampposts	#test lampposts	#FL clients	#Models	#Training samples	#Test samples	Accuracy (%)	F1-Score
Personalised	133 [†]	133 [†]	–	133	281804	70549	98.57	0.990
	7 [‡]	7 [‡]	–	7	5891	1476	94.82	0.945
	140	140	–	140	287695	72025	98.25	0.984
Centralised	140	133 [†]	–	1	287695	70549	98.41	0.988
	140	7 [‡]	–	1	287695	1476	93.39	0.943
	140	140	–	1	287695	72025	98.01	0.983
FL benchmark	140	133 [†]	140	1	287695	70549	95.89	0.967
	140	7 [‡]	140	1	287695	1476	92.15	0.932
	μ	140	–	–	287695	72025	94.02	0.949

[†] actual normal case
[‡] actual edge case
 μ average of models

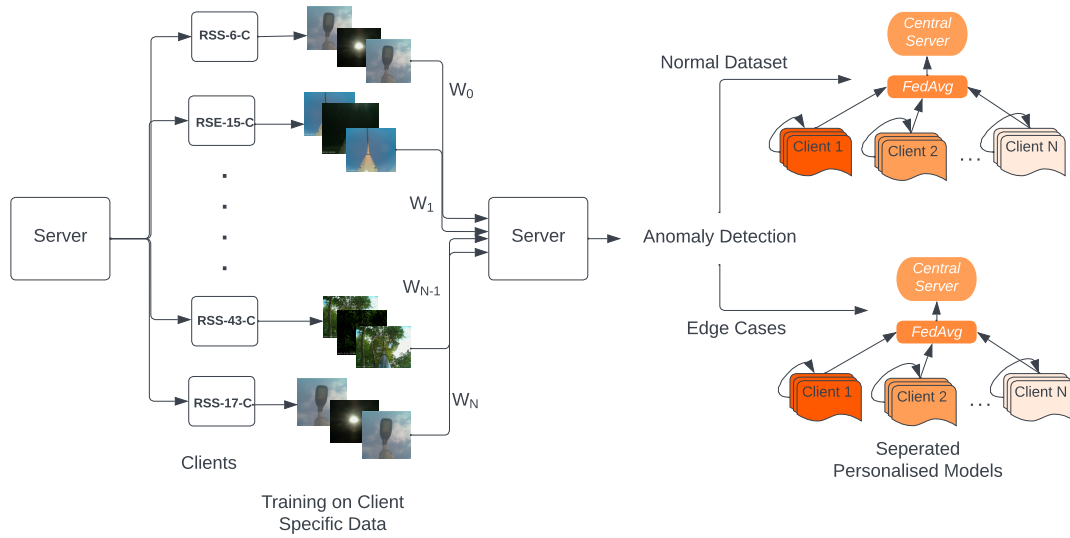


Figure 3: Flowchart of Proposed Approach for Personalised Federated Learning

5 CONCLUSIONS AND FUTURE RESEARCH

In this paper, we evaluated three different methods of detecting lamppost operability in a smart city environment. As seen, a “fully” personalised method provides a strong performance but does not scale well. On the other hand, a centralised approach is very demanding on the communication overhead introduced. FL can provide benefits of both worlds but still lacks in terms of accuracy. We suggest the following for future FL and personalised FL research challenges:

- (i) We are currently experimenting with a “tuned” personalised method, whereby certain model layers are trained and optimised solely on each device and dataset. In contrast, the

remaining model layers are used for aggregation when receiving updates from the global FL parameter server. Work is still ongoing, with promising results regarding accuracy observed and the reduction of communication costs.

- (ii) Clustered personalised FL based on client model parameters received at the FL parameter server (proposed method shown in Figure 3). Ideally, this would allow for early detection of the lamppost node types, which could then be separately aggregated into multiple (but concurrent) FL global models. Any new connected lamppost joining the network could receive a copy of each FL global model and run local tests to evaluate performance on its local dataset before conducting

local training/tuning optimisation. Early experiments suggest it is possible to detect and classify extreme edge case nodes (i.e., node type 2) fairly accurately, but we fail to detect the other node types. Whilst having only two global models for FL would work, ideally, we want to personalise the clusters to the extent that we achieve even higher accuracies, again to drastically reduce notifications/alerts received by the government officials monitoring the system and prioritising lampposts for maintenance.

- (iii) Given the hypothesised upper bound of performance (currently, the personalised method achieved an averaged F1-score of 0.98), it may suggest that further FL personalisation strategies may struggle to gain significant improvements. As such, in particular with ML, often the dataset might be limiting in terms of achieving such a high performance; for example details/nuances might be missed or averaged out when the images are drastically reduced in size during our pre-processing step (see Figure 2). Consequently, we have been experimenting with image metadata and other image statistics, particularly the mean and median green pixel values. Our intuition is that in extreme edge cases, where the lighting element is not visible to the monitoring sensor camera, small amounts of reflected light (or remnants of refracted light through vegetation) might be detectable. Our preliminary experiments using such metadata to improve detection accuracy have proven relatively successful but need careful calibration and integration into a robust and scalable personalised FL method.
- (iv) Communication overhead in FL can be further improved by introducing selective update strategies, such as dynamic sampling or selective masking on the models exchanged [1, 8]. Such methods can enhance the system's scalability, mainly when introducing thousands of clients. Furthermore, adaptive compression on the generated models [2, 6] can bring even more benefits by reducing the communication overhead and enabling the exchange of data over low data rate IoT technologies and longer distances.

ACKNOWLEDGMENTS

This work is funded in part by Toshiba Europe Ltd. UMBRELLA project is funded in conjunction with South Gloucestershire Council by the West of England Local Enterprise Partnership through the Local Growth Fund, administered by the West of England Combined Authority.

REFERENCES

- [1] Saif Anwar, Pietro E Carnelli, and Aftab Khan. 2022. Methods and Systems for Remote Training of a Machine Learning Model. US Patent App. 16/952,308.
- [2] Saif Anwar, Pietro E Carnelli, and Aftab Khan. 2022. System and Method for Adaptive Compression in Federated Learning. US Patent App. 16/952,705.
- [3] F. Chen, M. Luo, Z. Dong, Z. Li, and X. He. 2018. Federated meta-learning with fast convergence and efficient communication. *arXiv preprint* (2018).
- [4] Jie Ding, Mahyar Nemati, Chathurika Ranaweera, and Jinho Choi. 2020. IoT Connectivity Technologies and Applications: A Survey. *IEEE Access* 8 (2020), 67646–67673. <https://doi.org/10.1109/ACCESS.2020.2985932>
- [5] Tim Farnham, Simon Jones, Adnan Aijaz, Yichao Jin, Ioannis Mavromatis, Usman Raza, Anthony Portelli, Aleksandar Stanoev, and Mahesh Sooriyabandara. 2021. UMBRELLA Collaborative Robotics Testbed and IoT Platform. In *2021 IEEE 18th Annual Consumer Communications and Networking Conference (CCNC)*. IEEE, 1–7. <https://doi.org/10.1109/CCNC49032.2021.9369615>
- [6] C. Hardy, E. Le Merrer, and B. Sericola. 2017. Distributed Deep Learning on Edge-Devices: Feasibility via Adaptive Compression. In *2017 IEEE 16th International Symposium on Network Computing and Applications (NCA)*. IEEE Computer Society, Los Alamitos, CA, USA, 1–8. <https://doi.org/10.1109/NCA.2017.8171350>
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep Residual Learning for Image Recognition. (12 2015).
- [8] S. Ji, W. Jiang, A. Walid, and X. Li. 2022. Dynamic Sampling and Selective Masking for Communication-Efficient Federated Learning. *IEEE Intelligent Systems* 37, 02 (mar 2022), 27–34. <https://doi.org/10.1109/MIS.2021.3114610>
- [9] Aftab Khan, Pietro E Carnelli, Timothy David Farnham, Ioannis Mavromatis, and Anthony Portelli. 2022. System and Method for Detecting and Rectifying Concept Drift in Federated Learning. US Patent App. 17/018,923.
- [10] Yanming Lee, Hongyi Zhang, and Jimmi Rosa. 2019. Street Lamp Fault Diagnosis System Based on Extreme Learning Machine. *IOP Conference Series: Materials Science and Engineering* 490 (apr 2019), 042053. <https://doi.org/10.1088/1757-899x/490/4/042053>
- [11] Hui Lin, Wenxin Liu, and Xiaoding Wang. 2021. A Secure Federated Learning Mechanism for Data Privacy Protection. In *2021 20th International Conference on Ubiquitous Computing and Communications (IUCC/CIT/DSCI/SmartCNS)*. 25–31. <https://doi.org/10.1109/IUCC-CIT-DSCI-SmartCNS55181.2021.00019>
- [12] M-Lluïsa Marsal-Llacuna, J. Colomer-Llinàs, and J. Meléndez-Frigola. 2015. Lessons in urban monitoring taken from sustainable and livable cities to better address the Smart Cities initiative. *Technological Forecasting and Social Change* 90 (2015), 611–622.
- [13] Ioannis Mavromatis, Aleksandar Stanoev, Pietro Carnelli, Yichao Jin, Mahesh Sooriyabandara, and Aftab Khan. 2022. A Dataset of Images of Public Streetlights with Operational Monitoring using Computer Vision Techniques. <https://doi.org/10.48550/ARXIV.2203.16915>
- [14] Ioannis Mavromatis, Aleksandar Stanoev, Pietro Carnelli, Yichao Jin, Mahesh Sooriyabandara, and Aftab Khan. 2022. Images of Public Streetlights with Operational Monitoring using Computer Vision Techniques. <https://doi.org/10.5281/zenodo.6046759>
- [15] H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2016. Communication-Efficient Learning of Deep Networks from Decentralized Data. (2 2016).
- [16] Ahmed El Oudrhiri and Ahmed Abdelhadi. 2022. Differential Privacy for Deep and Federated Learning: A Survey. *IEEE Access* 10 (2022), 22359–22380. <https://doi.org/10.1109/ACCESS.2022.3151670>
- [17] Johannes Schneider and Michalis Vlachos. 2021. Personalization of Deep Learning. In *Data Science – Analytics and Applications*, Peter Haber, Thomas Lampoltshammer, Manfred Mayr, and Kathrin Plankensteiner (Eds.). Springer Fachmedien Wiesbaden, Wiesbaden, 89–96.
- [18] Bhagya Nathali Silva, Murad Khan, and Kijun Han. 2018. Towards Sustainable Smart Cities: A Review of Trends, Architectures, Components, and Open Challenges in Smart Cities. *Sustainable Cities and Society* 38 (2018), 697–713. <https://doi.org/10.1016/j.scs.2018.01.053>
- [19] Kehua Su, Jie Li, and Hongbo Fu. 2011. Smart City and The Applications. In *2011 International Conference on Electronics, Communications and Control (ICECC)*. 1028–1031. <https://doi.org/10.1109/ICECC.2011.6066743>
- [20] Dongjun Tang, Fulai Ding, Boya Deng, Pingkang Zhang, Qingming Wang, and Hui Lv. 2021. An Intelligent Fault Diagnosis Method for Street Lamps. In *2021 International Conference on Internet, Education and Information Technology (IEIT)*. 300–303. <https://doi.org/10.1109/IEIT53597.2021.00073>
- [21] S. Umamaheswari, K.Hari Priya, and S.Arun Kumar. 2021. Technologies Used in Smart City Applications – An Overview. In *2021 International Conference on Advancements in Electrical, Electronics, Communication, Computing and Automation (ICAECA)*. 1–6. <https://doi.org/10.1109/ICAECA52838.2021.9675707>
- [22] Y. Huang, L. Chu, Z. Zhou, L. Wang, J. Liu, J. Pei, and Y. Zhang. 2021. Personalized cross-silo federated learning on non-iid data. *Association for the Advancement of Artificial Intelligence (AAAI)* (2021).
- [23] Yu-Sheng Yang, Shih-Hsiung Lee, Guan-Sheng Chen, Chu-Sing Yang, Yueh-Min Huang, and Ting-Wei Hou. 2020. An Implementation of High Efficient Smart Street Light Management System for Smart City. *IEEE Access* 8 (2020), 38568–38585. <https://doi.org/10.1109/ACCESS.2020.2975708>
- [24] Andrea Zanella, Nicola Bui, Angelo Castellani, Lorenzo Vangelista, and Michele Zorzi. 2014. Internet of Things for Smart Cities. *IEEE Internet of Things Journal* 1, 1 (2014), 22–32. <https://doi.org/10.1109/JIOT.2014.2306328>
- [25] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. 2018. Federated Learning with Non-IID Data. (6 2018).
- [26] Zhaohua Zheng, Yize Zhou, Yilong Sun, Zhang Wang, Boyi Liu, and Keqiu Li. 2022. Applications of Federated Learning in Smart Cities: Recent Advances, Taxonomy, and Open Challenges. *Connection Science* 34, 1 (2022), 1–28. <https://doi.org/10.1080/09540091.2021.1936455>