

Federated Learning of Socially Appropriate Agent Behaviours in Simulated Home Environments

Saksham Checker
Delhi Technological University, India
checkersaksham01@gmail.com

Nikhil Churamani
University of Cambridge, UK
nikhil.churamani@cl.cam.ac.uk

Hatice Gunes
University of Cambridge, UK
hatice.gunes@cl.cam.ac.uk

ABSTRACT

As social robots become increasingly integrated into daily life, ensuring their behaviours align with social norms is crucial. For their widespread *open-world* application, it is important to explore Federated Learning (FL) settings where individual robots can learn about their unique environments while also learning from each others' experiences. In this paper, we present a novel FL benchmark that evaluates different strategies, using multi-label regression objectives, where each client individually learns to predict the social appropriateness of different robot actions while also sharing their learning with others. Furthermore, splitting the training data by different contexts such that each client *incrementally* learns across contexts, we present a novel Federated Continual Learning (FCL) benchmark that adapts FL-based methods to use state-of-the-art Continual Learning (CL) methods to continually learn socially appropriate agent behaviours under different contextual settings. Federated Averaging (FedAvg) of weights emerges as a robust FL strategy while rehearsal-based FCL enables *incrementally learning* the social appropriateness of robot actions, across contextual splits.

KEYWORDS

Federated Learning, Federated Continual Learning, Human-Robot Interaction, Distributed Learning, Social Robotics.

ACM Reference Format:

Saksham Checker, Nikhil Churamani, and Hatice Gunes. . Federated Learning of Socially Appropriate Agent Behaviours in Simulated Home Environments. In *HRI'24: Workshop on Lifelong Learning and Personalization in Long-Term Human-Robot Interaction (LEAP-HRI)*. ACM, New York, NY, USA, 5 pages.

1 INTRODUCTION

Social robots deployed in *open-world* human-centred environments are required to dynamically expand their knowledge by learning new tasks while preserving past knowledge [4, 5]. Such adaptation can enable them to support their users in day-to-day tasks by embedding themselves seamlessly within the social settings of their environments. Most robotic solutions for current day applications are designed as stand-alone implementations, tailored to specific tasks and/or environments [5]. As advances in Artificial Intelligence (AI) and Machine Learning (ML) gear robots towards a more ubiquitous presence, there is a need to explore adaptive learning paradigms that to not only facilitate a widespread and generalised application but also allow individual robots to personalise towards end-user requirements and preferences. This can be in the form of several robots deployed, in a distributed manner, across different contextual settings, interacting with multiple users at a time and learning different tasks [9]. Under such complex and diverse

application settings, there is a need to move beyond centralised platforms towards more distributed learning paradigms, enabling robots to learn *continually* while sharing their learning with others.

Traditional ML-based robotic applications, especially where multiple robots are deployed in parallel, usually follow a *centralised* learning (see Figure 1; left) approach where each robot collects data from its individual environment and communicates it to a central server. The data from each robot is then aggregated and a unified global model is trained to be used by each individual robot. Despite enabling robots to share experiences amongst each other, centralised learning approaches focus on developing a *one-size-fits-all* solution by training a unified model that can generalise across applications. Data privacy becomes a major concern as each robot shares the data collected by them with the central server which may not be acceptable in certain situations.

Federated Learning (FL) [22] (see Figure 1; middle), on the other hand, allows robots to learn independently from their own unique experiences, updating their learning models using only the data collected by them locally. Over time, these local updates for each agent can be aggregated across the centralised server, in the form of model updates that can inform training the unified centralised model. FL allows for a more privacy-preserving learning paradigm where local data is never shared with a centralised server. FL solutions have been used popularly in embedded or EdgeAI devices [13] that benefit from *distributed* learning settings [32] gathering and processing their own data in their unique application settings but also sharing their learning towards training a global aggregated model that allows devices to share knowledge between each other [18].

As social robots interact with their environments gathering data, they need to efficiently discern novel knowledge from past experiences and adapt their learning models to accommodate new knowledge [25]. Under FL settings, this means that the data collected by each robot individually need not be *independent and identically distributed* (i.i.d), requiring the robot to learn with sequential streams of data in an incremental manner, personalising individual robots towards their environment and users. Continual Learning (CL) [10, 24] can help address this problem further by enabling robots to adapt their learning with continuous and sequential streams of data acquired from non-stationary or changing environments [5, 16]. This may be achieved by *regularising* model updates, *replaying* already seen information or dynamically *expanding* models to accommodate new information [5]. Combining FL and CL, Federated Continual Learning (FCL) [9, 30] (see Figure 1; right) allows for individual robots, learning with sequential streams of unique local data, to also benefit from other robots' learning. Each agent periodically sends their model parameters to the centralised server where the knowledge from all agents is aggregated into a unified model which is sent back.

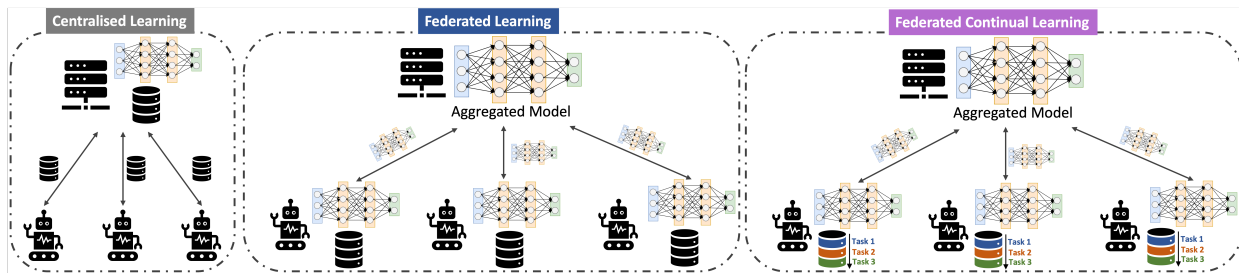


Figure 1: Centralised Learning (left) requires robots to share data with the server to train a shared model. **Federated Learning (FL) (middle)** allows local model weights or gradients to be aggregated on the server without sharing data. **Federated Continual Learning (FCL) (right)** further allows individual robots to incrementally learn tasks, sharing model updates with each other.




Figure 2: MANNERS-DB: A living room scenario with the Pepper robot. Adapted from [29].

Such distributed learning settings are particularly desirable for social robots operating in human-centred environments, to understand and learn socially appropriate behaviours, depending upon the context of the interaction, environmental factors as well as individual user preferences [5, 9]. Whether it is effectively navigating complex social environments [7, 21], learning approach and positioning behaviours [8, 23] or learning task-specific behaviours [29], it is essential for robots to consider the social-appropriateness of their behaviours in order to comply with social norms [2, 6].

In this paper, we explore simulated environments with humans and robotic agents to learn the social appropriateness of different high-level tasks as a use case for FL and FCL-based application of open-world learning. Depending upon the user, context or social norms [2], the agents need to learn what actions may be appropriate for them to perform and how they will be viewed by their users [23, 29]. Here, we explore the MANNERS-DB dataset [29] that provides social-appropriateness ratings for different agent actions in simulated home settings. We benchmark different FL and FCL methods to understand how such a learning of socially-appropriate agent behaviours can be realised in a distributed manner (FL) learning incrementally and sequentially (FCL), where individual agents effectively share their learning with each other.

2 METHODOLOGY

2.1 Learning Scenario: MANNERS-DB Dataset

Learning socially appropriate behaviours in complex home settings requires robots to be sensitive to its positioning with respect to other objects and users as well as individual user preferences. In this work, we explore a simulated living room scenario consisting of different actors where the agent is tasked upon learning the social appropriateness of different actions. For this, the MANNERS-DB

dataset [29] is used that consists of 3D scenes, created with Unity, of the Pepper robot co-inhabiting a living room space with other humans (adults and children) and animals under different social settings (see Figure 2). For each scene, the robot can perform 8 different tasks, that is, *vacuuming*, *mopping*, *carrying warm/cold food*, *carrying big/small objects*, *carrying drinks* or *cleaning/starting conversations*. These tasks can be performed by the robot either within a circle of influence or in the direction of operation with the only difference being cleaning within a circle and starting a conversation in the direction of the arrow. Crowd-sourced annotations are provided for the social appropriateness of each of these actions for every scene (out of a total of ≈ 1000 scenes), labelled on a 5-point Likert scale, ranging from very inappropriate to very appropriate.

2.2 Experimentation Settings

2.2.1 Input Features and Data Augmentation: For each scene, a 29-dimensional descriptor is provided consisting of *features* such as a flags for circle of influence or direction of operation, number of humans, children and animals, distance between the robot and the 3 closest humans, amongst others (see [29] for the complete list). We use the 29-d scene descriptors as the input to the model to predict the social-appropriateness of each of the 8 actions (within the circle or in the direction of the arrow). For both FL and FCL evaluations, the data is split into training and test splits in the ratio of 75% : 25%. The training data is further split amongst the different clients (2 or 10) with the shared test-set used for evaluation. For FCL evaluations, the training set of individual clients (or simulation nodes) is further split into two tasks, that is, samples depicting the robot operating with an circle (Task 1) and in the direction of the arrow (Task 2). Since the MANNERS-DB dataset is a relatively small dataset with approx. 1000 samples, we also benchmark the different FL and FCL methods using data augmentation as well. For this, a Gaussian noise ($\mu = 0, \sigma = 0.01$) is added to each feature.

2.2.2 Implementation Details: For each FL and FCL approach, a Multilayer Perceptron (MLP)-based model is implemented consisting of two Fully Connected (FC) layers of 16 units each with a *linear* activation. Each FC layer is followed by a BatchNormalisation layer. The output of the last FC layer is passed to the 8-unit output layer, predicting the social appropriateness for each of the 8 robot actions. The experiments are run for 2 – 10 clients. This relatively low number of clients is to compare these methods for a potential real-world evaluation to be conducted using physical robots. For

Table 1: Federated Learning results for the MANNERS-DB dataset for two (left) and ten (right) clients. Bold values denote best while [bracketed] denote second-best values.

Method	Two Clients			Ten Clients		
	Loss	RMSE	PCC	Loss	RMSE	PCC
	<i>W/O Augmentation</i>					
FedAvg	0.219	0.468	0.445	[0.225]	0.475	0.453
FedBN	0.222	0.471	[0.456]	0.226	0.475	[0.454]
FedProx	[0.220]	[0.469]	0.458	[0.225]	[0.474]	0.465
FedOpt	0.223	0.472	0.448	0.224	0.473	0.452
FedDistill	0.250	0.500	0.425	0.670	0.818	0.419
	<i>W/ Augmentation</i>					
FedAvg _{Aug}	0.212	0.460	0.421	0.213	0.462	0.419
FedBN _{Aug}	0.231	0.480	[0.402]	[0.224]	[0.473]	0.402
FedProx _{Aug}	[0.222]	[0.471]	0.401	0.227	0.476	[0.404]
FedOpt _{Aug}	0.231	0.481	0.397	0.227	0.477	0.401
FedDistill _{Aug}	0.251	0.501	0.383	0.231	0.481	0.399

brevity, results for only 2 and 10 clients are presented. All models are implemented using the PyTorch¹ and Flower² Python Libraries.

2.2.3 Evaluation Metrics: Since each client is learning the social appropriateness of each of the 8 possible actions, we use regression-based loss and evaluation metrics. The models are trained using the Mean Squared Error (MSE), computed as an average across the 8 actions. Furthermore, the average Root Mean Squared Error (RMSE) and Pearson’s Correlation Coefficient (PCC) [3] scores are also reported, computed as an average across the 8 actions.

3 THE FL BENCHMARK

We compare different state-of-the-art FL methods, both without and with data augmentation, presenting a novel benchmark for learning social appropriateness for different robot actions simultaneously.

3.1 Compared approaches

3.1.1 FedAvg: Federated Averaging (FedAvg) [18] is a straightforward approach for weight aggregation across clients in rounds where at each round, a centralised server gathers the model weights from m clients, aggregates them by computing the average across these m clients to form the *global model weights* and then updates the weights of each client with the global model weights.

3.1.2 FedBN: One of the main problems with FedAvg comes under *heterogeneous* data conditions where local data is non-i.i.d. FedBN [19] aims to address this problem by adapting FedAvg by keeping the parameters for all the BatchNormalisation layers ‘strictly local’, that is, all other model weights are aggregated across clients apart from the BatchNormalisation parameters.

3.1.3 FedProx: Similar to FedBN, FedProx [17] also proposes improvements over FedAvg by allowing for only partial aggregation of weights by adding a proximal term to FedAvg. The objective for each client is thus modified to minimise $F_k(\omega) + \frac{\mu}{2} \|\omega - \omega^t\|^2$ where F_k is the loss, ω are the local model weights to optimise and ω^t are the global parameters at time-step t . FedAvg can be considered to be a special case of FedProx with $\mu = 0$.

3.1.4 FedOpt: Another challenge faced by FedAvg is that of *adaptivity*. To address this, FedOpt [26] is proposed as a ‘general optimisation framework’ where each client uses a *client optimiser* to

optimise on local data while the server updates apply a gradient-based *server optimiser* to the aggregated model weights. We use the Adam optimiser for both client and server optimisation. FedAvg can be considered to be a special case of FedOpt where both client and server optimisers use StochasticGradientDescent (SGD) with server learning rate set to 1.

3.1.5 FedDistill: The FedDistill [14] approach also aims to improve the ability of the clients to deal with *heterogeneous* data conditions by using *knowledge distillation* [11]. Each client maintains two models: (i) a local copy of the global model and (ii) a personalised model that acts as a teacher to the student global model. The updated student model is then aggregated across clients.

For the above-mentioned approaches, in our experiments, each client undergoes 10 aggregation rounds and test-metrics are computed at the end of each round using the aggregated global model.

3.2 Results and Discussion

Table 1 presents the FL benchmark results. For 2 clients, the train-set is split into two equal parts while for 10 clients, it is split into ten equal parts. Thus, when evaluating models without augmentation, there is relatively more data per client for two clients compared to ten clients. We see that FedAvg performs the best under such settings, especially as each sample enabling learning across all 8 actions, with FedProx being a close second. For 10 clients however, adaptive optimisers under FedOpt are able to work well with low amount of per-client data with FedProx performing the second best. Similar trends are witnessed in evaluations with augmentation where a relatively large amount of data is available for all clients under the 2 and 10-client splits. FedAvg performs the best here as well while FedProx and FedBN are the next best approaches. Our evaluations presents a multi-label regression problem which is different from classification where FedAvg offers a relatively simple and robust learning methodology to predict the social appropriateness of agent behaviours in the MANNERS-DB dataset. Even in situations when FedAvg is not the best performing approach the difference in model performances is marginal.

4 THE FCL BENCHMARK

As can be seen in the FL results (see Table 1), FedAvg emerges as a simple and robust approach in our multi-label regression set-up. For incremental learning across tasks under non-i.i.d settings, we adapt FedAvg for FCL using different state-of-the-art CL-based objectives proposing FCL variants for FedAvg. Each client incrementally learns the social appropriateness for different robot actions under different learning contexts using the following CL methods, followed by a weight aggregation round where model weights are averaged across clients. We focus primarily on regularisation-based CL as these methods do not require additional computational resources. Naive Rehearsal (NR) [12] is included as a baseline for rehearsal-based methods. All methods are compared, both without and with data augmentation, presenting a novel FCL benchmark.

4.1 Compared Approaches

4.1.1 FedAvg_{EWC}: The Elastic Weight Consolidation (EWC) [15] approach introduces quadratic penalties on weight updates between old and new tasks. For each parameter, an importance value is computed using that task’s training data, approximated as a *Gaussian*

¹<https://pytorch.org>

²<https://flower.dev>

Table 2: Federated Continual Learning results for the MANNERS-DB dataset for two (left) and ten (right) clients. Data is split into two tasks: Circle (Task 1) and Arrow (Task 2). Bold values denote best while [bracketed] denote second-best values.

Method	Two Clients						Ten Clients					
	After Task 1			After Task 2			After Task 1			After Task 2		
	Loss	RMSE	PCC	Loss	RMSE	PCC	Loss	RMSE	PCC	Loss	RMSE	PCC
	<i>W/O Augmentation</i>											
FedAvg _{EWC}	[0.248]	[0.492]	0.562	0.265	0.503	0.579	0.263	0.504	0.543	0.262	0.502	0.557
FedAvg _{EWCOnline}	0.242	0.488	0.621	0.261	0.501	0.581	0.249	0.491	0.582	0.268	0.507	0.550
FedAvg _{SI}	0.258	0.502	[0.577]	[0.240]	0.483	[0.585]	0.262	0.505	0.546	0.249	0.489	[0.574]
FedAvg _{MAS}	0.271	0.513	0.567	0.241	[0.481]	0.589	[0.251]	[0.494]	0.573	0.261	0.502	0.556
FedAvg _{NR}	0.262	0.503	0.531	0.235	0.480	0.564	0.252	[0.494]	[0.577]	[0.250]	[0.494]	0.583
	<i>W/ Data-augmentation</i>											
FedAvg _{EWC}	0.197	0.439	0.631	0.240	0.482	0.528	0.179	0.419	0.620	[0.226]	[0.470]	[0.541]
FedAvg _{EWCOnline}	0.188	0.429	0.619	0.251	0.491	0.530	[0.180]	[0.420]	0.634	[0.226]	[0.470]	0.542
FedAvg _{SI}	0.195	0.438	[0.622]	0.236	0.481	[0.538]	0.184	0.425	[0.637]	0.240	0.481	0.529
FedAvg _{MAS}	[0.192]	[0.434]	0.621	[0.232]	[0.475]	0.534	0.184	0.424	0.640	0.232	0.474	0.534
FedAvg _{NR}	0.205	0.447	0.613	0.211	0.456	0.550	0.182	0.422	0.631	0.221	0.466	0.540

Distribution with its mean as the task parameters and the importance determined by the diagonal of the Fischer Information Matrix.

4.1.2 FedAvg_{EWCOnline}: EWOnline [27] offers an improvement over EWC where, instead of maintaining individual quadratic penalty terms for each of the tasks, a *running sum* of the Fischer Information Matrices for the previous tasks is maintained.

4.1.3 FedAvg_{SI}: The Synaptic Intelligence (SI) [31] approach penalises changes to weight parameters or synapses such that new tasks can be learnt without forgetting the old. To avoid forgetting, importance for solving a learned task is computed for each parameter and changes in important parameters are discouraged.

4.1.4 FedAvg_{MAS}: The Memory Aware Synapses (MAS) [1] approach attempts to alleviate forgetting by calculating an importance value for each parameter by examining the sensitivity of the output function instead of the loss function. Higher the impact of changes to a parameter, higher is the importance assigned and higher is the penalty imposed. Yet, different from EWC and SI, parameter importance is calculated using only unlabelled data.

4.1.5 FedAvg_{NR}: For the Naive Rehearsal (NR) [12] approach, each client maintains a replay buffer where a fraction of previously seen data is stored. This *old data* is interleaved with the new data to create mixed mini-batches to train the model by simulating i.i.d data settings in an attempt to mitigate forgetting in the model. The above-mentioned CL-based adaptations to the FedAvg approach are applied locally for each client. The importance values for EWC, EWOnline, SI and MAS approaches are calculated before aggregation across clients. The computed importance values for each of the parameters are then used to penalise changes in local weight updates between tasks, mitigating *forgetting*.

4.2 Results and Discussion

Table 2 presents the FCL results. For each approach, average test-set metrics are calculated using the aggregated global model after all aggregation rounds are completed for each task. Task 1 results represent test-set results only for data pertaining to the ‘circle’ split while the entire test-set is used to evaluate the models after Task 2. Learning incrementally is seen to have a positive effect on model performance. This is evidenced from the average PCC values (*after task 2*) being better for FCL vs. FL evaluations, for both 2 and 10 clients. Without using data augmentation, rehearsal-based NR

approach performs better than regularisation-based approaches after witnessing both tasks as it maintains a memory buffer to store previously seen task 1 samples. With the relatively low number of samples in the MANNERS-DB dataset, almost all the samples from task 1 can be maintained in the memory buffer, resulting in the better performance scores for FedAvg_{NR}. For 10 clients, SI comes closer however NR still achieves the best PCC scores. A similar trend is seen with data augmentation as well where NR still is able to retain past knowledge the best. Yet, as a separate memory buffer needs to be maintained for NR, it may not be the most resource-efficient approach. This may become particularly challenging when dealing with high-dimensional data such as images or videos [28].

5 CONCLUSIONS AND FUTURE WORK

This work presents a novel benchmark for learning socially appropriate robot behaviours in home settings comparing different FL and FCL approaches. For FL evaluations, FedAvg offers a relatively simple and robust learning methodology matching baseline evaluation scores from traditional ML-based methods [29]. This motivates the use of FedAvg to be adapted for FCL evaluations when incrementally learning different tasks. Our FCL evaluations show that rehearsal-based NR approach is best suited for such applications albeit being memory intensive. In this work, we explore pre-extracted 29-d scene descriptions to predict the social appropriateness of different robot actions. Our future work will focus on end-to-end learning directly using scene renders while exploring more resource-efficient generative feature replay methods [20, 28].

ACKNOWLEDGMENTS

Funding: S. Checker contributed to this work while undertaking a remote visiting studentship at the Department of Computer Science and Technology, University of Cambridge. N. Churamani and H. Gunes are supported by Google under the GIG Funding Scheme.

Open Access: For open access purposes, the authors have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

Data Access Statement: This study involves secondary analyses of the existing datasets, that are described and cited in the text.

Code Access: <https://github.com/nchuramani/FCL-MannersDB>.

REFERENCES

- [1] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. 2018. Memory Aware Synapses: Learning What (not) to Forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 139–154.
- [2] Ali Ayub and Alan R. Wagner. 2020. What Am I Allowed to Do Here?: Online Learning of Context-Specific Norms by Pepper. In *Social Robotics*, Alan R. Wagner, David Feil-Seifer, Kerstin S. Haring, Silvia Rossi, Thomas Williams, Hongsheng He, and Shuzhi Sam Ge (Eds.). Springer International Publishing, 220–231.
- [3] Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. 2009. Pearson correlation coefficient. In *Noise reduction in speech processing*. Springer, 37–40.
- [4] Zhiyuan Chen and Bing Liu. 2018. *Lifelong Machine Learning*. Springer International Publishing, Cham.
- [5] Nikhil Churamani, Sinan Kalkan, and Hatice Gunes. 2020. Continual Learning for Affective Robotics: Why, What and How?. In *29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 425–431.
- [6] Henriette Cramer, Jorrit Goddijn, Bob Wielinga, and Vanessa Evers. 2010. Effects of (in) accurate empathy and situational valence on attitudes towards robots. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 141–142.
- [7] Anthony Francis, Claudia Pérez-d'Arpino, Chengshu Li, Fei Xia, Alexandre Alahi, Rachid Alami, Aniket Bera, Abhijit Biswas, Joydeep Biswas, Rohan Chandra, et al. 2023. Principles and guidelines for evaluating social robot navigation algorithms. *arXiv preprint arXiv:2306.16740* (2023).
- [8] Yuan Gao, Fangkai Yang, Martin Frisk, Daniel Hernandez, Christopher Peters, and Ginevra Castellano. 2019. Learning socially appropriate robot approaching behavior toward groups using deep reinforcement learning. In *2019 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 1–8.
- [9] Luke Guerdan and Hatice Gunes. 2023. Federated Continual Learning for Socially Aware Robotics. In *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. 1522–1529.
- [10] Raia Hadsell, Dushyant Rao, Andrei A. Rusu, and Razvan Pascanu. 2020. Embracing Change: Continual Learning in Deep Neural Networks. *Trends in Cognitive Sciences* 24, 12 (Dec. 2020), 1028–1040.
- [11] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling the Knowledge in a Neural Network. In *NIPS Deep Learning and Representation Learning Workshop*.
- [12] Yen-Chang Hsu, Yen-Cheng Liu, Anita Ramasamy, and Zsolt Kira. 2018. Re-evaluating Continual Learning Scenarios: A Categorization and Case for Strong Baselines. In *NeurIPS Continual learning Workshop*.
- [13] Ahmed Imteaj, Urmish Thakker, Shiqiang Wang, Jian Li, and M. Hadi Amini. 2022. A Survey on Federated Learning for Resource-Constrained IoT Devices. *IEEE Internet of Things Journal* 9, 1 (2022), 1–24.
- [14] Donglin Jiang, Chen Shan, and Zhihui Zhang. 2020. Federated Learning Algorithm Based on Knowledge Distillation. *2020 International Conference on Artificial Intelligence and Computer Engineering (ICAICE)*, 163–167.
- [15] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences* 114, 13 (2017), 3521–3526.
- [16] Timothée Lesort, Vincenzo Lomonaco, Andrei Stoian, Davide Maltoni, David Filliat, and Natalia Díaz-Rodríguez. 2020. Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges. *Information Fusion* 58 (2020), 52–68.
- [17] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. 2020. Federated Optimization in Heterogeneous Networks. In *Proceedings of Machine Learning and Systems*, I. Dhillon, D. Papailiopoulos, and V. Sze (Eds.), Vol. 2. 429–450.
- [18] Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zhihua Zhang. 2020. On the Convergence of FedAvg on Non-IID Data. In *International Conference on Learning Representations*.
- [19] Xiaoxiao Li, Meirui JIANG, Xiaofei Zhang, Michael Kamp, and Qi Dou. 2021. FedBN: Federated Learning on Non-IID Features via Local Batch Normalization. In *International Conference on Learning Representations*.
- [20] Xialei Liu, Chenshen Wu, Mikel Menta, Luis Herranz, Bogdan Raducanu, Andrew D. Bagdanov, Shangling Jui, and Joost van de Weijer. 2020. Generative Feature Replay For Class-Incremental Learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE.
- [21] Luis J. Manso, Pedro Nuñez, Luis V. Calderita, Diego R. Faria, and Pilar Bachiller. 2020. SocNav1: A Dataset to Benchmark and Learn Social Navigation Conventions. *Data* 5, 1 (2020).
- [22] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 54)*, Aarti Singh and Jerry Zhu (Eds.). PMLR, 1273–1282.
- [23] Emily McQuillin, Nikhil Churamani, and Hatice Gunes. 2022. Learning Socially Appropriate Robo-Waiter Behaviours through Real-Time User Feedback. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (Sapporo, Hokkaido, Japan) (HRI '22). IEEE Press, 541–550.
- [24] German I. Parisi, Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter. 2019. Continual Lifelong Learning with Neural Networks: A review. *Neural Networks* 113 (2019), 54–71.
- [25] Jitendra Parmar, Satyendra Chouhan, Vaskar Raychoudhury, and Santosh Rathore. 2023. Open-World Machine Learning: Applications, Challenges, and Opportunities. *ACM Comput. Surv.* 55, 10, Article 205 (feb 2023), 37 pages.
- [26] Sashank J. Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and Hugh Brendan McMahan. 2021. Adaptive Federated Optimization. In *International Conference on Learning Representations*.
- [27] Jonathan Schwarz, Wojciech Czarniecki, Jelena Luketina, Agnieszka Grabska-Barwinska, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. 2018. Progress & Compress: A scalable framework for continual learning. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, Stockholm, Sweden, 4528–4537.
- [28] Samuil Stoychev, Nikhil Churamani, and Hatice Gunes. 2023. Latent Generative Replay for Resource-Efficient Continual Learning of Facial Expressions. In *17th IEEE International Conference on Automatic Face and Gesture Recognition (FG)*. 1–8.
- [29] Jonas Tjomsland, Sinan Kalkan, and Hatice Gunes. 2022. Mind Your Manners! A Dataset and a Continual Learning Approach for Assessing Social Appropriateness of Robot Actions. *Frontiers in Robotics and AI* 9 (2022).
- [30] Jaehong Yoon, Wonyong Jeong, Giwoong Lee, Eunho Yang, and Sung Ju Hwang. 2021. Federated Continual Learning with Weighted Inter-client Transfer. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 12073–12086.
- [31] Friedemann Zenke, Ben Poole, and Surya Ganguli. 2017. Continual Learning Through Synaptic Intelligence. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70 (Sydney, NSW, Australia). Proceedings of machine learning research* 70, 3987–3995.
- [32] Chen Zhang, Yu Xie, Hang Bai, Bin Yu, Weihong Li, and Yuan Gao. 2021. A survey on federated learning. *Knowledge-Based Systems* 216 (2021), 106775.