



The 13th International Conference on Emerging Ubiquitous Systems and Pervasive Networks
(EUSPN 2022)
October 26-28, 2022, Leuven, Belgium

Towards Lifelong Federated Learning in Autonomous Mobile Robots with Continuous Sim-to-Real Transfer

Xianjia Yu^{a,*}, Jorge Peña Queraltá^a, Tomi Westerlund^a

^a*Turku Intelligent Embedded and Robotic Systems (TIERS) Lab, University of Turku, Finland*

Abstract

The role of deep learning (DL) in robotics has significantly deepened over the last decade. Intelligent robotic systems today are highly connected systems that rely on DL for a variety of perception, control and other tasks. At the same time, autonomous robots are being increasingly deployed as part of fleets, with collaboration among robots becoming a more relevant factor. From the perspective of collaborative learning, federated learning (FL) enables continuous training of models in a distributed, privacy-preserving way. This paper focuses on vision-based obstacle avoidance for mobile robot navigation. On this basis, we explore the potential of FL for distributed systems of mobile robots enabling continuous learning via the engagement of robots in both simulated and real-world scenarios. We extend previous works by studying the performance of different image classifiers for FL, compared to centralized, cloud-based learning with a priori aggregated data. We also introduce an approach to continuous learning from mobile robots with extended sensor suites able to provide automatically labelled data while they are completing other tasks. We show that higher accuracies can be achieved by training the models in both simulation and reality, enabling continuous updates to deployed models.

© 2022 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0>)

Peer-review under responsibility of the Conference Program Chairs

Keywords: Sim-to-real; Federated learning; Visual obstacle avoidance; Robot navigation; Continuous learning;

1. Introduction

Vision-based obstacle avoidance is an important enabler for widespread autonomous robot navigation and could also help visually handicapped citizens [1]. In general, vision sensors have been having an increasing role in intelligent robotic systems, as advances in deep learning (DL) have led to new state-of-the-art solutions mobile navigation, human-like walking, teaching through demonstration, and collaborative automation, among others in areas such as advanced perception and intelligent control [2, 3, 4].

* Corresponding author.

E-mail address: xianjia.yu@utu.fi

In this paper, we extend the previous work [5] on federated learning for vision-based obstacle avoidance in mobile robots towards analyzing different architectures and, most importantly, introducing an strategy for continuous learning from both simulated and real robots carrying out different missions requiring autonomous navigation. There are a number of works in the literature showing the potential of DL for vision-based obstacle avoidance [6], and works integrating other techniques such multi-view structure-from-motion for enhanced performance [7]. However, in addition to the use of DL in [5], the methods are also validated in the real world, a federated learning (FL) framework is introduced, and the effects of learning through heterogeneous environments are analyzed.

Federated learning enables sharing knowledge without transferring raw data. Therefore, FL allows for privacy-preserving distributed learning, which can then be enhanced by other technologies, such as distributed ledger technologies [8], has been utilized in multiple domains in robotics and autonomous system [9]. To the best of our knowledge, this is the first work to introduce an strategy for continuous learning together with federated learning and validating it in the real world. In addition to real-world robots, we also utilize simulations to obtain enough data for training the DL models. Due to the proliferation of photorealistic simulators in robotics fields, studies are increasingly relying on these simulators to supplement data collection in situations where we cannot reach or collect sufficient data. In addition, through these types of simulators, the deployment of robotic and autonomous systems can be more accessible. In this way, we believe it may be beneficial for real-world vision-based obstacle avoidance in the real world. Throughout their lives, humans and animals can acquire, fine-tune, and transfer knowledge and skills. It is instrumental and interesting for robots to have the same type of learning capabilities. To continually learn the model for obstacle avoidance like humans, we involved multiple robots performing other tasks in different places, including the simulator, into the FL-based lifelong learning system for data collection, model training, and model sharing. We compare the performance of both AlexNet and ResNet18 with centralized and federated learning approaches to model training. We evaluate the model on different synthetic and real-world datasets, and train the models on combinations of different subsets to better study the effect of environment heterogeneity during the training phase. New data is also acquired with a Clearpath Husky mobile robot while it is operating other tasks and autonomously navigating a new indoors environment. The new data is automatically labelled with additional sensor suited. This opens the door to wider usage of heterogeneous robot fleets where robots with additional sensor capabilities can generate labelled data to train models able to reproduce their autonomous behaviour with more limited sensors, mainly cameras.

Aided by AI, especially DL, each robot can have a model representing the environment based on their situated awareness. Different robots may have the limitation of detecting the environment due to their limited resources. It is of great necessity to have a collective model to share their knowledge about the environment. Instead of an individual robot, in a multi-robot system, multiple robots, situated in various whereabouts, collaboratively performing particular tasks is more efficient and of high success rate in heterogeneous environments including unknown ones [10]. However, because of the limits of scenarios where vision-based obstacle avoidance is applied, it is still challenging for DL to collect enough visual inputs for collaborative learning and share them with other agents for privacy and security reasons. To address these concerns, we proposed an approach that included FL, Sim-to-Real (Sim2Real) via a photorealistic simulator, and Lifelong Learning.

In summary, we investigated the possibility of federated and continuous learning within hybrid teams of simulated and real robotic agents in this work. We then evaluate the performance benefits of such an approach over offline learning or learning from more limited data sources. First, we evaluate two different deep obstacle avoidance neural networks with both synthetic and real-world data. For both of the architectures a FL-based knowledge sharing method (where the locally trained models are fused) is compared to a centralized training approach (where raw data needs to be aggregated before the training starts). Both models are validated with data from the photorealistic simulator and real-world environment separately. Second, we analyzed the sim-to-real performance of two different deep obstacle avoidance models generated by FL methods, with our results showing that FL outperforms the centralized data aggregation methods. Third, we integrate lidar-based navigation for automated labelled data gathering. We implemented an online FL-based visual obstacle avoidance system both in a simulator and real-world environment. With such a system, we can continuously collect data from obstacles and free paths and train the model while the robots operate other tasks.

The rest of this paper is organized as follows. In Section II we review related works in FL for autonomous robots, sim-to-real transfer and lifelong robot learning. The methodology followed for obtaining the results reported in the

manuscript is then introduced in Section III, with Section IV delving into the actual experimental results. Finally, section V concludes the paper and outlines future work directions.

2. Related Work

Multi-robot collaboration and DL for robotics have both played an increasingly important role in multiple robotic applications [4]. However, most of the work to date in learning from real-world experiences, sim-to-real transfer, and continuous learning, has been dedicated to reinforcement learning [3, 11] and robotic manipulation [12]. Within the possibilities to achieve collaborative learning, one of the most straightforward approaches is cloud-based centralized learning [13], with a server where data is aggregated and training occurs at once or in batches, but in an offline manner. Federated learning, in contrast, offers a distributed solution to collaborative learning where, in turn, the process is privacy-preserving (raw data can be processed locally). It also allows for management of networking resources (choosing when and how to transfer the models) through distributed computation at the edge [14].

In the area of lifelong learning and continuous learning for robotics, the literature has been in general dedicated to simulation environments only. While AI advances the intelligence of robots in various subfields of robotics, numerous challenges remain with real-world application, particularly in an unknown environment. Among these limitations are a lack of open-ended learning about item categories and scenes, the absence of a global navigation map, and the occurrence of object collisions, particularly in dynamic environments. With the assistance of Lifelong learning, the robot should be able to learn new object categories continuously and comprehend affordances, especially for robots manipulating tasks from a small number of on-site training examples [15]. Lifelong learning for navigation is a good application of lifelong learning. Compared with the traditional planner in navigation, this learning approach can improve the navigation performance based on its own experience and retain the ability after learning new ones [16].

In terms of deep learning vision-based obstacle avoidance, this has been a researched topic partly owing to the higher degree of maturity of DL methods for vision data in comparison to other modalities such as lidar data [17, 18]. Cameras are also more widely available and are relatively inexpensive sensors, while the methods can be generalized to various environments owing to the lack of the need for specific geometric or topological features. For example, in [19], autonomous aerial robots are deployed with minimal knowledge of the environment and obstacle avoidance is achieved through a single monocular camera. With models trained through deep reinforcement learning (DRL), noise to data in the training phase can bring increased resilience to changes in the environments [20], a method that is typical in other DRL approaches [3, 21]. In this paper, we also exploit heterogeneity of environments to show how a wider variety of the training data brings performance improvements to the DL models, especially when they are trained with a FL approach.

In the domain of sim-to-real to transfer for DL policies driving autonomous behaviour in robotic system, the most widely studied learning algorithms are DRL approaches [3]. Sim-to-real transfer with DRL has been particularly studied in depth in problems involving dexterous manipulation [21]. Even in this area, recent works in the literature also show the potential of DRL for vision-based obstacle avoidance [22, 23]. The authors of these works demonstrate the effectiveness of the methods even when the robots encounter in the real world new environments and objects.

3. Methodology

This section outlines the robot platform utilized in this experiment, two different deep neural networks to do vision-based obstacle avoidance for further validation of FL performance, and navigation settings for continual learning.

3.1. Data collection for FL

The details of the data gathering, including the usage of the photorealistic simulator Nvidia Isaac Sim, the environment settings, and data distributions, are detailed in [5], from where we reuse the base simulation and training datasets. The datasets from [5] include data from three distinct scenarios in NVIDIA Isaac Sim and from Jetbot robots deployed in three real-world rooms to train and validate the vision-based obstacle avoidance models. In this work, the datasets from the simulator are represented as $\mathcal{S}_{i, i \in \{0,1,2\}}$ where i indicates the simulated environments, including a hospital, office, and warehouse. The three real-world datasets are denoted as $\mathcal{R}_{i, i \in \{0,1,2\}}$ where i indicates office spaces,



Fig. 1: The customized Clearpath Husky platform.

hallways, and laboratory environments, respectively. Additional datasets acquired specifically for this work are introduced in the relevant sections, with Husky training data referred to as \mathcal{H}^S or \mathcal{H}^R for the simulated and real robots, respectively.

Regarding data training hardware, in this work, we utilized a Lambda Vector workstation equipped with two RTX 3080 GPU cards and a 24-core AMD Threadripper 3960X processor to train our models for vision-based obstacle avoidance.

3.2. Vision-based obstacle avoidance models

We trained two distinct types of DL models to ensure FL's performance in visual obstacle avoidance and assess how performance differs. These two deep convolutional neural networks (CNN) are vision-based obstacle classifiers for two classes that define whether the environment ahead is *blocked* or *free* for the robot to navigate. Owing to the relatively low level of complexity of the classification task and the size of potential datasets for such tasks, we have selected the AlexNet [24] and ResNet18 [25] architectures as appropriate for such binary classification task. AlexNet and ResNet are both commonly exploited backbones for conducting various tasks across multiple domains.

These two models are generic deep learning models designed to aid robots in discerning between various types of barriers in heterogeneous situations. This strategy contrasts with other possibilities, such as object detection or semantic segmentation (e.g., segmenting free floor from objects and walls). The chosen approach enables us to concentrate on examining the performance of a federated learning approach and its capacity for sim-to-real transfer rather than on developing a specific obstacle avoidance strategy, which is the study's primary purpose.

3.3. Proposed FL based lifelong learning obstacle avoidance

By incorporating lidar into the Husky navigation system, we developed a straightforward FL-based lifelong learning system for visual obstacle avoidance. We used lidar to differentiate between blocked and free space of visual data while performing specified navigation tasks. After collecting sufficient data, the robots will train the models independently and send them to the server for aggregation into a global model. Concerning the fusion of models in FL, federated averaging was applied in this work by taking the average of all model updates. The global model can then be deployed to another type of robot for obstacle avoidance performance evaluation.

Regarding the experimental environments, We implemented this with a Clearpath Husky Fig. 1 robot both in a simulator and real-world environments. Our customized Husky robot platform is equipped with one Ouster OS0-128 lidar for measuring distances to objects when performing autonomous navigation tasks. When the distances are lower to a certain threshold, an Intel Realsense L515 camera will be triggered to capture the images as obstacles included continuously. Otherwise, the images will be categorized to free labels. Once a sufficient number of images are collected, we will train local models based on these data, fuse them to be a global model by FL methods, and then apply the global model to a real-world robot obstacle avoidance operation.

It's worth emphasizing that we used simulated data to help the continuous learning process, as obtaining real-world data is not always straightforward. In the following section, we evaluate the FL fused model's real-world deployment performance using either simulated or real-world data.




Fig. 2: Accuracy of models obtained through centralized learning with aggregated data or federated learning with fused local models based on ResNet18. These results are trained (t) and validated (v) with respective simulation datasets (S_i^t , S_i^v) and real datasets (\mathcal{R}_i^t , \mathcal{R}_i^v) independently.




Fig. 3: Accuracy of the different models obtained through centralized learning with aggregated data or federated learning with fused local models based on AlexNet. These results are trained (t) and validated (v) with respective simulation datasets (S_i^t , S_i^v) and real datasets (\mathcal{R}_i^t , \mathcal{R}_i^v) independently.

Table 1: Area under ROC curve (AUC) values for the aggregated centralized learning and federated learning approaches.

		Training datasets (AlexNet, ResNet18)								
		Centralized learning with aggregated data				Federated learning				
		$S_{0,1}^t$	$S_{0,2}^t$	$S_{1,2}^t$	$S_{1,2,3}^t$	$S_{0,1}^t$	$S_{0,2}^t$	$S_{1,2}^t$	$S_{1,2,3}^t$	
Validation datasets	Sim	$S_{0,1}^v$	(0.33, 0.53)	(0.50, 0.58)	(0.71, 0.56)	(0.52, 0.31)	(0.85, 0.74)	(0.85, 0.78)	(0.85, 0.79)	(0.85, 0.78)
		$S_{0,2}^v$	(0.33, 0.54)	(0.50, 0.40)	(0.75, 0.96)	(0.60, 0.21)	(0.94, 0.88)	(0.93, 0.89)	(0.93, 0.86)	(0.95, 0.89)
		$S_{1,2}^v$	(0.42, 0.64)	(0.50, 0.50)	(0.23, 0.26)	(0.62, 0.23)	(0.96, 0.84)	(0.95, 0.92)	(0.95, 0.93)	(0.95, 0.91)
Validation datasets	Real	$\mathcal{R}_{0,1}^v$	(0.75, 0.42)	(0.63, 0.50)	(0.08, 0.33)	(0.58, 1.00)	(0.88, 0.86)	(0.88, 0.75)	(0.88, 0.94)	(0.88, 1.00)
		$\mathcal{R}_{0,2}^v$	(0.35, 0.68)	(0.25, 0.69)	(0.50, 0.48)	(0.70, 0.71)	(0.83, 0.80)	(0.85, 0.86)	(0.85, 0.76)	(0.85, 0.82)
		$\mathcal{R}_{1,2}^v$	(0.46, 0.45)	(0.51, 0.42)	(0.48, 0.47)	(0.56, 0.67)	(0.90, 0.70)	(0.92, 0.87)	(0.90, 0.73)	(0.92, 0.82)

4. Experimental Results

This section presents experimental results obtained using data from both simulated and real robots. We first demonstrate the performance of various knowledge-sharing approaches for obstacle avoidance using AlexNet and ResNet18 models before delving into the possibility of sim-to-real knowledge transferability. Finally, we demonstrated the FL method's ability to facilitate lifelong learning by having Husky performing simulation and real-world navigation tasks.

4.1. Centralized training vs. federated learning

Our experiments begin by analyzing the performance improvements of federated learning over traditional centralized training with data aggregation. Both simulated and real-world data are evaluated separately in this part. Extending the previous works in [5], by including the AlexNet and ResNet18 in this study, we can compare how the performance will vary for centralized training and FL using different DL models. To accomplish this, we used the data we collected in the simulated hospital (S'_0), office (S'_1), and warehouse (S'_2) to train our model on each dataset and all possible combinations ($S'_{0,1}$, $S'_{0,2}$, $S'_{1,2}$, $S'_{0,1,2}$) of two or three of these datasets and validate the models on $S^v_{i,i \in \{0,1,2\}}$. Equivalently for the federated learning approach, we run different training rounds in which we simulate that a different subset of robots is collaboratively learning without sharing any actual raw data. Only the models are fused in this approach, and a global model is updated iteratively. In the case of real-world data, we repeat the procedure above with R^t_0 , R^t_1 , R^t_2 representing the training datasets from rooms, office and laboratory and $R^t_{0,1}$, $R^t_{0,2}$, $R^t_{1,2}$, $R^t_{0,1,2}$ representing the combinations of two or three of the previous sets while $R^v_{i,i \in \{0,1,2\}}$ denotes the corresponding validation data. Figure 3 and Fig. 2 report the accuracy of the different models (AlexNet and ResNet18) for centralized learning and FL, respectively. From the results, we found that FL based approach is robust both in a sim and real separately, and its accuracy is competitive with traditional centralized data aggregation training methods for performing vision-based obstacle avoidance tasks.

Along with accuracy, we also calculate the area under the ROC curve (AUC) for each of the scenarios where training is conducted through either the centralized or federated approaches. The results are summarized in Table 1 where the AUC values for AlexNet-based and ResNet18-based models are enclosed with a bracket in order. This metric enables a better understanding of the models' reliability. In the context of robotic navigation, there is indeed a cost differential between false negatives over false positives in terms of the robots' integrity. However, from the point of view of performance, false positives can significantly degrade the navigation speed and time, while low-frequency collisions can be avoided with other sensors, considering as well that a false negative is not necessarily consistent over time and multiple observations of the same obstacle are processed before a collision may happen.

For both AlexNet and ResNet18, the AUC results together with the accuracy boxplots show that there is a performance boost with the federated learning approach in contrast to the centralized learning method where all data is first aggregated in a single training set.

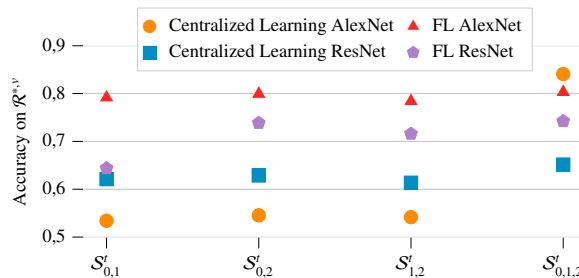


Fig. 4: Evaluation of sim-to-real capability of the trained models.

In terms of sim-to-real performance of the two architectures, we followed the steps in [5] and analyzed the performance of both the centralized and federated learning approaches to an independent real-world dataset. The results showing the potential for vision-based obstacle avoidance inference for both AlexNet and ResNet18 are shown in Fig. 4. The results indicate that both AlexNet and ResNet18-based obstacle avoidance models implemented with the FL approach can outperform the ones with the centralized data aggregation method in terms of sim-to-real performance. Additionally, AlexNet is more suitable for performing obstacle avoidance tasks in our dataset.

4.2. Federated continual learning for visual-based obstacle avoidance

In this section, we describe the results from experiments where the Husky robot has been utilized both in a simulated playpen (see Fig. 5a) and in a large-open indoor office environment (see Fig. 5c) to operate autonomous navigation tasks in their navigation maps (see Fig. 5b and in Fig. 5d, respectively).

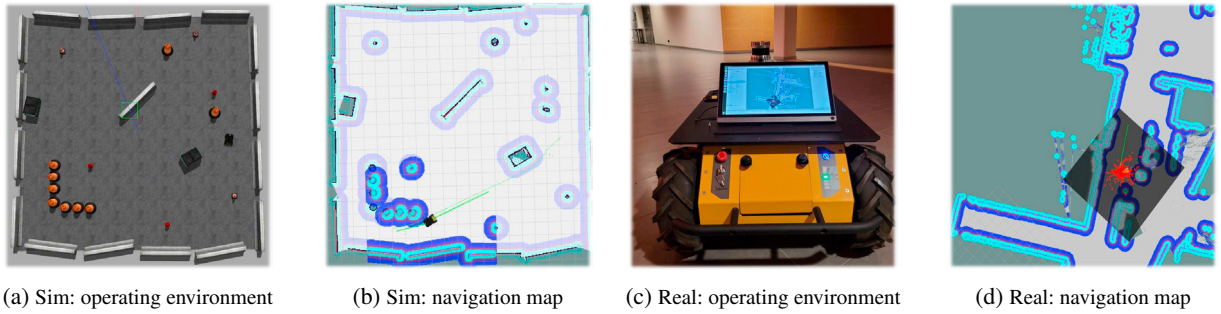


Fig. 5: Husky Navigation Environment and Maps




Fig. 6: Accuracy of FL fused models based on images collected while Husky was operating navigation tasks in simulated and real-world environments for continual learning purposes

At the same time, while the Husky robot was operating autonomous navigation tasks in the sim and real environments, we collected images continuously for training local models for FL models fusion. By validation on an independent dataset $\mathcal{R}^{s,v}$, we evaluated the global models fused with the local model from simulated husky robot \mathcal{H}^S only, the local model from real-world Husky robot \mathcal{H}^R only, and the local models of both them $\mathcal{H}^{S,R}$. The results are shown in Fig. 6. Regarding AlexNet, fusing the models either from simulation or real-world can improve the accuracy of the vision-based obstacle avoidance. However, for ResNet18, our results show that fusing the local model from the simulator can improve the performance of the global model more than fusing the one from the real-world Husky. By fusing both of them, the accuracy can be improved, which shows the potential benefit of collaborative learning from both simulated and real robots.

5. Discussion and Conclusion

This work presents an FL-based lifelong learning method for vision-based obstacle avoidance among various mobile robots involving both simulated and real-world environments. Rather than applying a single deep neural network, we analyzed the performance of the FL-based method compared with the centralized data aggregation method with two different deep neural networks, providing to better generalize results. More specifically, we found that the FL approach can bring competitive accuracy compared to centralized learning across the simulated and real worlds while also delivering inherent benefits in communication optimization and data privacy preservation, enabling collaboration across organizations or users. Additionally, we evaluated the FL method's sim-to-real vision-based obstacle avoidance performance. The result indicates that transferring the obstacle knowledge from simulation to reality using the FL method is more effective and stable. Within the FL-based lifelong learning system, one agent can improve its obstacle avoidance performance by aggregating models from local models in other agents situated either in simulated or real-world environments in our study.

In future work, we will concentrate on utilizing and adapting the FL-based lifelong learning system to perform other robotic navigation tasks with a view toward sim-to-real capabilities. We also find potential in dynamically adjusting the simulation environments based on real-world robot experiences, e.g., adding 3D models of new objects found.

Acknowledgements

This research work is supported in part by the Academy of Finland's AutoSOS project (Grant No. 328755) and RoboMesh project (Grant No. 336061).

References

- [1] M. Mendes, A. Coimbra, M. Crisostomo, Assis-cicerone robot with visual obstacle avoidance using a stack of odometric data, *IAENG Int. J. Comput. Sci* 45 (2018) 219–227.
- [2] H. A. Pierson, M. S. Gashler, Deep learning in robotics: a review of recent research, *Advanced Robotics* 31 (16) (2017) 821–835.
- [3] W. Zhao, J. Peña Queraltá, T. Westerlund, Sim-to-real transfer in deep reinforcement learning for robotics: a survey, in: *IEEE Symposium Series on Computational Intelligence*, IEEE, 2020.
- [4] J. Peña Queraltá, J. Taipalmaa, B. C. Pullinen, V. K. Sarker, T. N. Gia, H. Tenhunen, M. Gabbouj, J. Raitoharju, T. Westerlund, Collaborative multi-robot search and rescue: Planning, coordination, perception, and active vision, *IEEE Access* 8 (2020) 191617–191643.
- [5] X. Yu, J. P. Queraltá, T. Westerlund, Federated learning for vision-based obstacle avoidance in the internet of robotic things, *arXiv preprint arXiv:2204.06949* (2022).
- [6] J. O. Gaya, L. T. Gonçalves, A. C. Duarte, B. Zanchetta, P. Drews, S. S. Botelho, Vision-based obstacle avoidance using deep learning, in: *2016 XIII Latin American Robotics Symposium and IV Brazilian Robotics Symposium*, IEEE, 2016, pp. 7–12.
- [7] S. Yang, S. Konam, C. Ma, S. Rosenthal, M. Veloso, S. Scherer, Obstacle avoidance through deep networks based intermediate perception, *arXiv preprint arXiv:1704.08759* (2017).
- [8] Y. Xianjia, J. P. Queraltá, J. Heikkonen, T. Westerlund, Federated learning in robotic and autonomous systems, *arXiv preprint arXiv:2104.10141* (2021).
- [9] K. Kishor, Communication-efficient federated learning, in: *Federated Learning for IoT Applications*, Springer, 2022, pp. 135–156.
- [10] E. Olcay, F. Schuhmann, B. Lohmann, Collective navigation of a multi-robot system in an unknown environment, *Robotics and Autonomous Systems* 132 (2020) 103604.
- [11] B. Liu, L. Wang, M. Liu, Lifelong federated reinforcement learning: a learning architecture for navigation in cloud robotic systems, *IEEE Robotics and Automation Letters* 4 (4) (2019) 4555–4562.
- [12] W. Zhao, J. Peña Queraltá, L. Qingqing, T. Westerlund, Ubiquitous distributed deep reinforcement learning at the edge: Analyzing byzantine agents in discrete action spaces, *Procedia Computer Science* 177 (2020) 324–329.
- [13] B. Liu, L. Wang, X. Chen, L. Huang, D. Han, C.-Z. Xu, Peer-assisted robotic learning: a data-driven collaborative learning approach for cloud robotic systems, in: *2021 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2021, pp. 4062–4070.
- [14] A. Imteaj, U. Thakker, S. Wang, J. Li, M. H. Amini, A survey on federated learning for resource-constrained iot devices, *IEEE Internet of Things Journal* 9 (1) (2021) 1–24.
- [15] S. H. Kasaei, J. Melsen, F. van Beers, C. Steenkist, K. Voncina, The state of lifelong learning in service robots, *Journal of Intelligent & Robotic Systems* 103 (1) (2021) 1–31.
- [16] B. Liu, X. Xiao, P. Stone, A lifelong learning approach to mobile robot navigation, *IEEE Robotics and Automation Letters* 6 (2) (2021) 1090–1096.
- [17] Y. Li, L. Ma, Z. Zhong, F. Liu, M. A. Chapman, D. Cao, J. Li, Deep learning for lidar point clouds in autonomous driving: A review, *IEEE Transactions on Neural Networks and Learning Systems* 32 (8) (2020) 3412–3432.
- [18] Y. Xianjia, S. Salimpour, J. Peña Queraltá, T. Westerlund, Analyzing general-purpose deep-learning detection and segmentation models with images from a lidar as a camera sensor, in: *International Conference on Intelligent Systems Design and Engineering Applications*, Lecture Notes in Electrical Engineering (to appear), Springer, 2022.
- [19] A. Singla, S. Padakandla, S. Bhatnagar, Memory-based deep reinforcement learning for obstacle avoidance in uav with limited environment knowledge, *IEEE Transactions on Intelligent Transportation Systems* 22 (1) (2021) 107–118. doi:10.1109/TITS.2019.2954952.
- [20] L. Gao, J. Ding, W. Liu, H. Piao, Y. Wang, X. Yang, B. Yin, A vision-based irregular obstacle avoidance framework via deep reinforcement learning, in: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021. doi:10.1109/IR08S1168.2021.9636512.
- [21] W. Zhao, J. Peña Queraltá, L. Qingqing, T. Westerlund, Towards closing the sim-to-real gap in collaborative multi-robot deep reinforcement learning, in: *2020 5th International Conference on Robotics and Automation Engineering*, IEEE, 2020, pp. 7–12.
- [22] K. Zhang, J. Lin, L. Bi, T. Zhang, Sim2real learning of visionbased obstacle avoidance for robotic manipulators, in: *RSS Workshop*, 2020.
- [23] T. Zhang, K. Zhang, J. Lin, W.-Y. G. Louie, H. Huang, Sim2real learning of obstacle avoidance for robotic manipulators in uncertain environments, *IEEE Robotics and Automation Letters* 7 (1) (2021) 65–72.
- [24] A. Krizhevsky, I. Sutskever, G. E. Hinton, Imagenet classification with deep convolutional neural networks, *Advances in neural information processing systems* 25 (2012) 1097–1105.
- [25] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016.