

Past, Present, Future: A Comprehensive Exploration of AI Use Cases in the UMBRELLA IoT Testbed

Peizheng Li*, Ioannis Mavromatis[†], Aftab Khan*

* Bristol Research and Innovation Laboratory, Toshiba Europe Ltd., U.K.

[†]Digital Catapult, London, U.K.

Email: {Peizheng.Li, Aftab.Khan}@toshiba-bril.com, Ioannis.Mavromatis@digicatapult.org.uk

Abstract—UMBRELLA is a large-scale, open-access Internet of Things (IoT) ecosystem incorporating over 200 multi-sensor multi-wireless nodes, 20 collaborative robots, and edge-intelligence-enabled devices. This paper provides a guide to the implemented and prospective artificial intelligence (AI) capabilities of UMBRELLA in real-world IoT systems. Four existing UMBRELLA applications are presented in detail: 1) An automated streetlight monitoring for detecting issues and triggering maintenance alerts; 2) A Digital twin of building environments providing enhanced air quality sensing with reduced cost; 3) A large-scale Federated Learning framework for reducing communication overhead; and 4) An intrusion detection for containerised applications identifying malicious activities. Additionally, the potential of UMBRELLA is outlined for future smart city and multi-robot crowdsensing applications enhanced by semantic communications and multi-agent planning. Finally, to realise the above use-cases we discuss the need for a tailored MLOps platform to automate UMBRELLA’s model pipelines and establish trust.

Index Terms—IoT, wireless communication, robotics, edge intelligence, use cases, MLOps

I. INTRODUCTION

The Internet of Things (IoT) stands as a crucial research field, revolutionising industries, addressing societal challenges, and fostering innovation in connectivity, security, energy efficiency, and societal impact. Nowadays, the advancements in radio technologies, network protocols, and artificial intelligence (AI)/Machine learning (ML) bring more choices for designing an efficient IoT system with appropriate device selection, network architecture, and data processing capability. Nevertheless, the flip side of this coin introduces more critical challenges in building an IoT system as people expect this IoT network to have enough flexibility for pre-simulation, reconfiguration, scalability control, and, most importantly, intelligence and automation capability [1].

For the existing IoT testbeds, considering their distinctive features, supported use cases, and major focuses, they can be classified into three main categories: wireless experimentation, robotics research, and AI-related activities¹. Here, we list some of the operational and accessible platforms of each category and briefly introduce their objectives.

- **Wireless Experimentation:** In this class, examples like FIT IoT-Lab [2], FlockLab 2 Testbed [3], and w-ILab [4]

¹It should be noted this classification has no clear boundaries, as there is some overlap among the functionality of testbeds.

provide public access, enabling Low-power WAN (LPWAN) experimentation. They are used primarily to test wireless networking protocols and algorithms.

- **Robotics Research:** Robotarium [5], IRIS [6], and LivingLab [7] are notable for robotics and swarm activities. Robotarium features miniature custom-designed robots, enhancing battery autonomy and reducing the testbed’s footprint but limiting hardware capabilities. IRIS robots have fully programmable wireless IEEE 802.15.4 [8] radio interfaces, yet position accuracy degrades in long experiments due to accumulated drift. FIT IoT-Lab and w-ILab are still under development, resulting in limited functionality.
- **AI/ML related:** For AI/ML, platforms like Gradient [9] and OpenAI [10] offer subscription-based access for evaluating AI/ML algorithms. Google Colab provides similar functionality for free. Despite powerful hardware, these platforms lack access to real-time generated data.

In general, the existing testbeds exhibit limited designed capabilities for handling only specific tasks, with short-supported use cases and a notable absence of access to real-time, real-world data, especially in AI/ML-enabled platforms. In this context, this paper introduces the unique work of UMBRELLA² - Urban Multi Wireless Broadband and IoT Testing for Local Authority and Industrial Applications [11]. Functioning as a dynamic testing environment, UMBRELLA allows users to emulate realistic scenarios, prototype solutions, and assess their IoT innovations. It provides foundational capabilities, tools, and infrastructure for testing physical and digital solutions across various use cases, including Smart City initiatives, IoT/Industrial IoT (IIoT), and AI/ML applications. This paper delves explicitly into the AI/ML use cases developed within the UMBRELLA ecosystem, emphasising their validation and deployment in real-world scenarios. All use cases discussed in this paper are supported by a unified cloud-native platform, to enhance model reliability, thereby enabling their scalability, extensibility and robustness. Simultaneously, the paper envisions the potential AI/ML use cases for UMBRELLA in complex real-world scenarios for construing trustable crowdsensing systems, and the importance of embracing ML Operations (MLOps), outlining them as future developments. The rest of the paper is organized as

²UMBRELLA - A living lab: <https://www.umbrellaiot.com/>

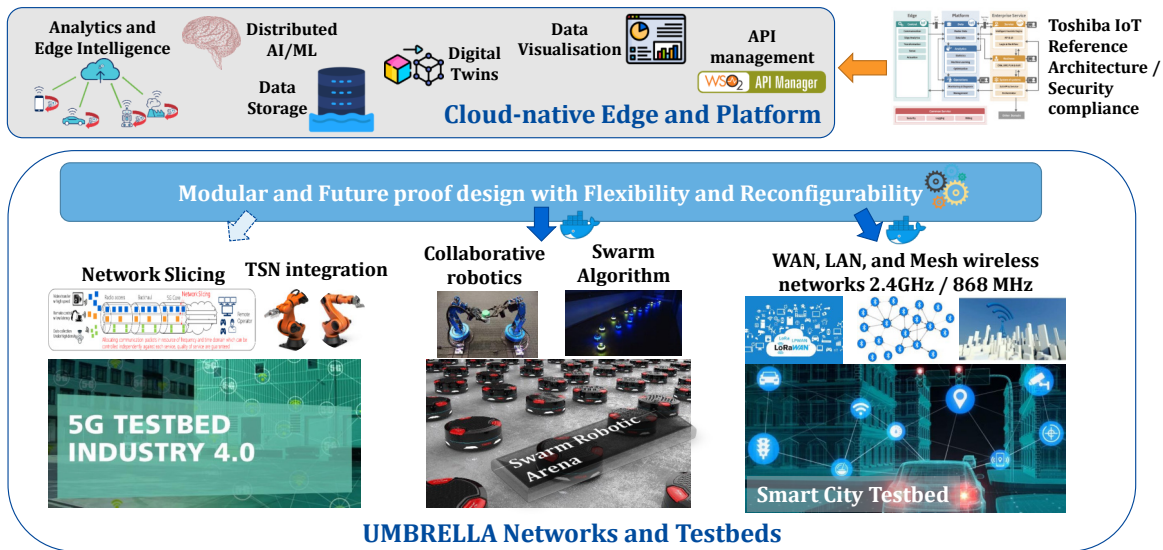


Figure 1. UMBRELLA SoS architecture overview with support of multiple sub-system testbeds.

follows: Sec. II introduces the UMBRELLA ecosystem and provides details about its features, such as system components, hardware design, and the implementation method for in-UMBRELLA applications. In Sec. III, various AI/ML use cases demonstrated within UMBRELLA are presented. This section includes both technical details and the results obtained. Sec. IV explores potential use cases scheduled to be implemented on the UMBRELLA platform. Sec. V emphasises the importance of adopting full-scale MLOps for AI/ML use cases in UMBRELLA. Finally, Sec. VI concludes this paper.

II. UMBRELLA TESTBED AND APPLICATIONS

A. What is UMBRELLA?

UMBRELLA [11] is a groundbreaking open-access IoT ecosystem deployed in the South Gloucestershire, U.K. region. Its primary goal is to stimulate innovation across diverse technological domains. Providing access to multiple testbed facilities, data, and computing resources, UMBRELLA links specialised test environments and offers a System-of-Systems (SoS) approach to address real-world technological challenges. Comprising 200+ multi-sensor, multi-wireless interface nodes installed on public streetlights, a robotics arena with 20 mobile robots, and a 5G network-in-a-box solution, all controlled via a unified backend platform, UMBRELLA ensures efficient management, control, and secure user access. UMBRELLA’s multi-domain architecture positions it as an ideal playground for IoT and IIoT innovation, embodying an open, sustainable testbed that actively guides future initiatives. With its unique blend of openness, heterogeneity, realism, and advanced tools, UMBRELLA is dedicated to advancing technology research, development, and translation into tangible real-world progress. The UMBRELLA SoS architecture overview supporting multiple sub-systems is illustrated in Fig. 1.

The guiding principle behind UMBRELLA is that each existing system can operate independently, but the interoperation

of various Constituent Systems (CSs) enables more complex scenario demonstrations. UMBRELLA has contributed to multiple use cases since its deployment in Sep. 2021. It also facilitates multiple systems developed by Small and Medium-sized Enterprises (SMEs) and has been fully functional since then, collecting over six million sensing samples daily. Essentially, UMBRELLA aims to foster a collaborative innovation ecosystem between users, industry, government, and academia to create economic and societal values.

B. UMBRELLA Hardware Design and the AI Capability

The hardware design of UMBRELLA places scalability and realism at the central stage; that is, the hardware should operate in realistic conditions and constraints, matching the requirements for the envisioned use cases. The hardware composition of UMBRELLA sensor node and collaborative Robotics testbed is illustrated as follows.

1) *UMBRELLA Sensors Node*: Fundamentally, the minimum functional unit of UMBRELLA is the “node”, which encompasses the mothership pod, edge computing pod, ambient sensing pod, and endpoint boards, as demonstrated in Fig. 2. The functionality of each node component is introduced below:

- **Mothership Pod**: It accommodates the Raspberry Pi CM3+ module on a custom carrier board, serving as the operational core of the node. Powered by Raspbian GNU/Linux 10 (buster) 32-bit and a custom kernel (ver. 4.19.95-v7+), the mothership ensures network connectivity and inter-module communication. The carrier board also integrates radio interfaces for network connectivity and wireless testbed capabilities.
- **Edge Computing Pod**: This pod contains an NVIDIA Jetson Nano connected via USB. The Jetson Nano operates on NVIDIA’s official Linux4Tegra 64-bit OS (derived from Ubuntu 18.04) with slight modifications. This module enables edge computing and supports GPU-intensive tasks.




Figure 2. (a) shows an UMBRELLA node with the enclosure open, showing modules for 1) mothership, 2) edge, 3) ambient sensing, 4) vacant ambient sensing expansion pod, and 5) the RPi Camera; (b) indicates the connection method of 1) mothership pod, 2) edge computing pod and 3) ambient sensing pod within a node.

- **Ambient Sensing Pod:** This pod includes an array of environmental sensors. A microcontroller board transmits sensor values to the mothership. The module, housed in a vented enclosure, is exposed to the atmosphere, ensuring accurate readings of the external environment.
- **Endpoint Boards:** The endpoint boards consist of customized Printed Circuit Board Assemblies (PCBAs) centred around a microcontroller, creating an abstraction layer between modular sensors or radios and the node’s remaining components. These boards facilitate lower-level communication with sensors through protocols like I2C, SPI, and I2S, effectively offloading time-critical operations from the main processing unit. UMBRELLA nodes offer two Bluetooth and one LoRaWAN interface for end-user access, a Wi-Fi interface for backbone connectivity, and a cellular interface for fail-safe management tasks in the absence of a main LAN connection.

2) *UMBRELLA Collaborative Robotics Testbed:* The basic element of this collaborative robot (cobot) system is the robot node, which supports multiple sensors, such as infrared, video cameras, Inertial Measurement Units (IMU), temperature, pressure, humidity sensors, etc. Two actuators (Omni-wheels, lift) are mounted in a node. Meanwhile, a number of wireless communication technologies (e.g., Wi-Fi, BLE, UWB, and 5G NR) are integrated into the node. The processor of each node is complemented with GPU (RockPi™ 4B) support, which permits embedded neural network model deployments for on-device processing. The software architecture is based on Docker containers and ROS2 DDS middleware for flexible and extensible evolution to support future sensors or network technologies. The node configuration and telemetry data exchange can be finished through user networks, service networks, access networks, proximity networks, and the meshed network. This testbed can be used for collaborative robotics experimentation, swarm algorithm development and wireless network technology evaluation³.

³Reader could refer to [12] for a more detailed description of UMBRELLA cobot testbed.

3) *UMBRELLA Application Implementation:* All experiments and applications are deployed as containers across all devices, orchestrated through a cloud-native platform. A unified backend controls all applications, providing storage and visualisation capabilities and ensuring smooth operation across the entire ecosystem. UMBRELLA users can package their applications as containerised applications, upload them to the UMBRELLA portal, and deploy them on the nodes, where they are orchestrated within multiple Kubernetes clusters. Scheduled experiments operate as stateless instances on either Raspberry Pi or Jetson Nano.

III. EXISTING AI USECASES UNDER UMBRELLA

In this section, we present existing example use cases within UMBRELLA, encompassing, i.e., (1) monitoring of street lighting fixtures; (2) digital twin for smart building; (3) large-scale Federated Learning (FL) experiments on UMBRELLA nodes; and (4) intrusion detection for containers on UMBRELLA edge devices. A concise overview figure is illustrated in Fig. 3. We briefly describe each use case and discuss how the UMBRELLA ecosystem facilitated their development, delivery and demonstration.

A. Streetlight Monitoring

1) *Usecase Description:* This use case involves overseeing the functionality of streetlight fixtures to identify any irregularities or malfunctions promptly. Any unanticipated operations trigger alerts for the council’s streetlighting maintenance team. Most lights, excluding those following a customised schedule, are programmed to switch off 15 minutes after sunrise and illuminate 15 minutes before sunset. The municipal street lighting team conducts manual checks at regular intervals, approximately every four weeks, to ensure proper functionality. Typically, the team relies on public reports to identify issues with the light fixtures. When multiple fixtures display unexpected behaviour, such as being ON or OFF outside the expected times, the team addresses these issues in batch corrections, implemented over a section of a road for cost efficiency. Our implementation automates this process and provides a more cost-effective way of managing fixed street assets. This use case leverages the deployed UMBRELLA sensor nodes and platform to identify the state of traffic lighting, automates the reporting procedure and reduces the intervals on the scale of days.

2) *AI Technique and Results:* The UMBRELLA nodes are equipped with an RPi Camera Module ver.1, facing the streetlight. Taking periodic photos of the fixture, malfunctions can be identified and reported. Our decision-making is based on two methods. We initially employed a simplistic computer vision mechanism, where photos taken from each lamppost are highly exposed. Based on that, it is easy to identify the ON/OFF status. Secondly, an ML classification algorithm was developed using the normal images and providing a confidence interval as well. We have identified that the simple computer vision approach works well at night, while the ML-based approach operates accurately during the day. The ML model operates

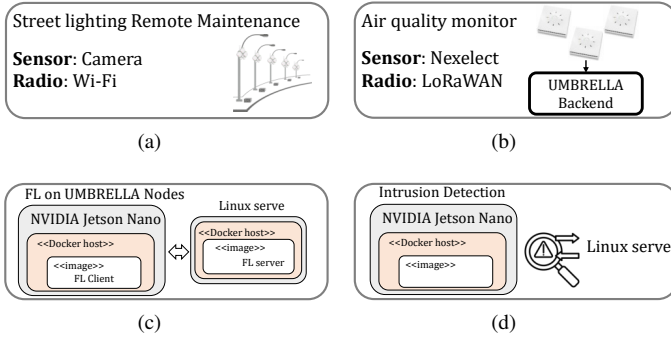


Figure 3. This figure illustrates an overview of the UMBRELLA AI use cases presented in this paper, where (a) refers to the street lighting remote monitoring; (b) is digital twin for building environment sensing; (c) depicts the FL trail on UMBRELLA node; (d) shows the intrusion detection on UMBRELLA edge devices.

with an accuracy of $\geq 90\%$ while the computer vision approach with almost 100%.

Our use case is still operational and helps the council’s team identify malfunctions and abnormalities in the fixtures’ operation. Moreover, complementary to this activity we collected and shared a dataset with the research community [13]. The dataset was automatically labelled based on the mechanisms mentioned above and comprised 350,000 images, taken over a six-month period using 140 UMBRELLA nodes.

B. Digital Twin for Buildings’ Environmental Sensing

1) *Usecase Description:* In the wireless front, UMBRELLA integrates various technologies to facilitate short- and long-range communications. Bluetooth and LoRaWAN particularly are accessible by end users and be used for experimentation. This use case focuses on temperature and air quality (AQ) monitoring, from sensor data collected from devices remotely connected to UMBRELLA. The data collected are later stored in UMBRELLA’s backend and processed in real-time. Some examples of real-time sensor data collected include temperature, CO₂, humidity, pressure, noise, etc.

2) *AI Technique and Results:* The ultimate goal of this use case is to construct a high-fidelity 3D data-driven digital twin that virtually “senses” a building environment. In other words, based on a small scale of measurements, can extrapolate the expected air quality in the future. In order to expedite the rendering of the virtual environment and minimise the cost of sensor anchor points (AP), an autoencoder-based ML model is utilised. This model is responsible for generating virtual APs and inferring their measurements from the real-time sensor data stream deployed within the same experimental setting while preserving fidelity.

For this use case, remote nodes connected wirelessly to the UMBRELLA infrastructure were used. These nodes are battery-powered small-factor sensor boards. The sensor APs are linked to UMBRELLA’s platform, transmitting measurements every 10 minutes. UMBRELLA facilitated all the computing resources required for model training and inference as well as

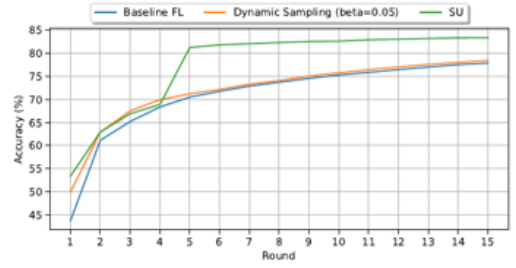


Figure 4. Preliminary results using four connected UMBRELLA IoT devices and containerized FL system.

the infrastructure for storing and visualising the data. The sensor boards were installed in care homes in the South Gloucestershire Council. Remarkably, the autoencoder model developed achieves an absolute percentage error of less than 1% for temperature and 10% for CO₂ measurements when compared to reference sensors while reducing the number of deployed sensors by one-fourth. The detailed description of this use case is available in [14].

C. Large-scale FL Experimentation

1) *Usecase Description:* Federated Learning (FL) is a decentralised machine learning framework designed to parallelise training across multiple interconnected devices simultaneously training a single model. In the context of UMBRELLA, UMBRELLA edge nodes, serve as network entry points and continuously collect raw data. In certain scenarios, IoT devices may face limitations in network quality, impacting their ability to share model parameters for FL. Transmitting substantial data volumes through low-power resource-constrained IoT networks (e.g. Bluetooth or LoRaWAN) becomes a significant bottleneck in FL. Some devices may also have constraints on energy usage, communication bandwidth, and other network resources, hindering their participation in model training. This use case aims to develop an FL system capable of training on UMBRELLA edge devices while minimising the transmission of model parameters. This approach offers the dual benefit of reducing network bandwidth requirements, improving quality of service, and lowering energy consumption during training. Additionally, the nature of FL ensures that the entire training dataset (raw sensor data) remains confidential, maintaining privacy during FL training.

2) *AI technique and Results:* In this large-scale FL experiment, leveraging the flexible configuration capability of UMBRELLA, we implemented a Selective Updates (SU) FL method [15] to minimise communication volume in FL training while maintaining inference performance. UMBRELLA’s distributed edge processing nodes were used for scaling up our experiment and evaluating our scenario across a realistic system. The proposed method operates on a client-round basis, employing a sliding/adaptive observation window and threshold to effectively gather local client model performance before deciding whether to update our Parameter Server (PS) or conduct further local training. The model deployment and update procedure are illustrated in Fig. 3c. We utilised the

MNIST Fashion dataset [16] for testing, where each image is 28×28 pixels and greyscale, featuring ten classes for the FL model to detect and classify. The results, presented in Fig. 4, within UMBRELLA platform, the SU method has shown substantial advantages in terms of convergence speed and accuracy compared to baseline FL and dynamic sampling.

D. Intrusion Detection System at IoT Edge

1) *Usecase Description:* The open-access nature of UMBRELLA, where end-users can deploy their containerised experiments raises security concerns for the functionality of the system. Attacks like privilege escalation, can compromise both the host and the underlying IoT infrastructure. Inspired by that, we developed a semi-supervised FL solution, aiming to detect such issues at the edges, and ensured the uninterrupted operation of the UMBRELLA platform. The ultimate goal of this use case is to train and deploy a real-time anomaly detection system with updated models throughout the entire UMBRELLA IoT infrastructure.

2) *AI Technique and Results:* A novel systematic approach is proposed to perform intrusion detection on the UMBRELLA edge, to detect any malicious activity by a container. The approach is divided into two parts: data pipeline (ingesting the data from multiple hosts) and modelling (training model from ingested data). The data pipeline begins by processing raw audit logs at the edge, extracting field names and values. Then, event-based feature engineering is applied based on a methodology validated in [17]. This streamlined pipeline transforms raw logs into a structured and enriched dataset for efficient analysis and insights. In the modelling phase, three scenarios are considered for generating training/test datasets: S_A , where container escape involves launching a denial of service (DoS) attack using a host shell; S_B , where container escape involves writing to host permission files, granting passwordless sudo permission to the user; and S_C , a benign container with an Apache PHP webserver without any escape⁴. In each dataset, there are three classes of labels for intrusion detection, including the "Normal", "Anomaly", and "Uncertainty". The S_C dataset is split for the pre-training and validation of the autoencoder model, which is subsequently validated on the S_A and S_B datasets. After pre-training the autoencoder, the model undergoes FL updates. As in the previous use case, the UMBRELLA edge computing nodes and backend provide the basis and underlying infrastructure for the FL deployment. In this case, the FL training follows a classical manner. That is, in each FL round, the model is independently trained on individual clients, and the FL model is aggregated to the server to update the global model. This aggregated model is then broadcast to each client for the next FL round. The final FL model can reach more than 98% of identifying accuracy on S_A and S_B datasets.

IV. FUTURE AI USECASES TOWARDS CROWDSENSING

In the future, it is expected that UMBRELLA will serve as a distinctive platform for exploring and advancing IoT/IIoT

technologies. Due to UMBRELLA's nature and its real-world deployment, Crowdsensing applications and demonstrations can leverage the existing capabilities. UMBRELLA nodes installed across a very busy road in the northern part of Bristol, and inside the University of West of England campus can serve as the underlying infrastructure for real-world Smart City Crowdsensing applications. Additionally, through the utilisation of its cobot testbed, inventive mechanisms can be evaluated enabling robot-to-robot and robot-to-human interactions. This involves enhancing data exchange capabilities, improving collaborative abilities, and ultimately enhancing overall efficiency in co-working. These representative examples are illustrated as follows.

A. Smart City & UMBRELLA AI

Smart Cities are defined as an urban environment that leverages technology and data-driven solutions to improve efficiency, sustainability, and the quality of life for its residents. Undoubtedly, AI plays a crucial role in achieving these objectives. The UMBRELLA system is anticipated to contribute significantly to the future smart city landscape. UMBRELLA's comprehensive sensing, computing, and robust AI capabilities form the basis for developing reliable crowdsensing systems for smart cities. These systems utilise sensor data as a criterion for intelligent decision-making, encompassing tasks such as air pollution monitoring and traffic control. Moreover, the edge computing device on the UMBRELLA node can seamlessly function as a complimentary multi-access edge computing platform. This functionality facilitates the smooth offloading and balancing of cloud computing demands from pervasive computing, particularly in domains like environmental monitoring, video surveillance, and the evolution of intelligent transportation systems.

B. Robotics & UMBRELLA AI

The UMBRELLA cobot testbed is designed for experiments where multiple robots work together or jointly with humans for a shared goal. Cobot can naturally be seen as a crowdsensing system. The success of the cobot task heavily relies on effective communication (data exchange) among them to process task allocation and coordination, sensor data sharing, position and location sharing, status updates, error detection, and recovery. Reducing the latency of such data sharing and decision-making will be the focus of the next step of UMBRELLA cobot-based crowdsensing system development. From the techniques perspective, potential consideration will be given to the following points.

- **Semantic communication:** It provides a promising technique to compress such raw sensor data in a more meaningful and context-aware manner. Deploying the embedded semantic encoder/decoder model on the UMBRELLA robot node can reduce the communication overhead, and improve the efficiency of cooperative robotics systems [19].
- **Communication and control co-design:** This approach entails the collaborative design and optimisation of communication and control systems, with the former playing a crucial

⁴The dataset of container escape detection is available on [18].

role in transmitting control information and feedback signals among system components. This co-design methodology is anticipated to enhance the robustness, agility, and determinism of the cobot system. By customising communication protocols on UMBRELLA node to meet the specific requirements of control algorithms, we can minimise unnecessary data transmission and reduce processing overhead.

- **Automated Planning:** The completion of this process is expected to be achieved by harnessing the inherent AI capabilities of the UMBRELLA robot node and utilising connected cloud/edge intelligence. This will enable the automatic determination of optimal paths and collaborative methods for the involved robots, facilitating a streamlined operation for the entire cobot system.

V. THE NEED FOR FULL-SCALE MLOPS

MLOps is a method of controlling the lifecycle of an ML model, right from design to deployment. It encompasses a set of principles that merge ML, DevOps, and data engineering to ensure the dependable and efficient deployment and maintenance of ML models in a production environment. It essentially involves delivering ML applications through the lens of DevOps, with a particular emphasis on data and ML models. MLOps is all about streamlining and expediting processes, wherein the streamlining entails automating the ML pipeline, covering the journey from data processing to model development for continuous training, as well as automating the Continuous Integration/Continuous Deployment (CI/CD) processes for ML applications [20]. Expediting refers to the ability to enhance the speed of delivering these applications while upholding service quality standards. The practice of the MLOps system brings three benefits for the ML model development and deployment, including (1) automated and effective model deployment via a unified workflow; (2) easy for quality control, ensuring the reliability of ML models in the real-world; (3) speed up model management and updating.

So far, there is a missing native MLOps platform that is adopted or specifically designed for UMBRELLA. The implementation and operation of applications in UMBRELLA are predominantly handled by Kubernetes orchestration systems, functioning as a DevOps platform. Consequently, the capabilities of UMBRELLA in terms of ML-related automation functions, such as processing training data, conducting model training/validation and ML model lifecycle management, are relatively limited. This limitation could pose challenges to the establishment of trustworthy crowd-sensing systems. Consequently, we intend to delve into the creation of an in-UMBRELLA MLOps platform. It is anticipated that all the AI use cases outlined in this paper will be built upon the native MLOps infrastructure in the future.

VI. CONCLUSIONS

This paper provides a comprehensive overview of the AI/ML applications crafted within the UMBRELLA framework. Four use cases accompanied by in-depth technical details and the

accomplished outcomes, are presented. Simultaneously, this paper explores potential crowdsensing applications in smart cities and robotics that are earmarked for future development, underscoring the potential growth of UMBRELLA. Additionally, the advantages of embracing full-scale MLOps in the context of UMBRELLA are highlighted. The authors humbly hope that this paper can positively contribute to the advancements of both the theoretical understanding and practical implementation of intelligent IoT application development.

ACKNOWLEDGMENTS

This work was supported by Toshiba Europe Ltd. and Bristol Research and Innovation Laboratory (BRIL).

REFERENCES

- [1] K. Shafique *et al.*, "Internet of things (IoT) for next-generation smart systems: A review of current challenges, future trends and prospects for emerging 5G-IoT scenarios," *IEEE Access*, pp. 23 022–23 040, 2020.
- [2] C. Adjih, E. Baccelli *et al.*, "FIT IoT-LAB: A Large Scale Open Experimental IoT Testbed," in *Proc. of IEEE WF-IoT*, Dec. 2015, pp. 459–464.
- [3] R. Trüb, R. Da Forno *et al.*, "FlockLab 2: Multi-Modal Testing and Validation for Wireless IoT," in *Proc. of Int. Conf. Work. CPS-IoTBenCh*, Sep. 2020.
- [4] S. Bouckaert, W. Vandenberghe *et al.*, "The w-iLab.t Testbed," in *Testbeds and Research Infrastructures. Development of Networks and Communities*, T. Magedanz, A. Gavras *et al.*, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 145–154.
- [5] D. Pickem, P. Glotfelter *et al.*, "The Robotarium: A Remotely Accessible Swarm Robotics Research Testbed," in *Proc. of IEEE ICRA*, Sep. 2017, pp. 1699–1706.
- [6] J. A. Tran, P. Ghosh *et al.*, "Intelligent Robotic IoT System (IRIS) Testbed," in *Proc. of IEEE/RSJ IROS*, Oct. 2018, pp. 1–9.
- [7] "Fraunhofer IML Robotic Testbed," https://www.iml.fraunhofer.de/en/news/_archiv/research-infrastructure---testbeds-.html, Accessed: 30-11-2023.
- [8] IEEE, "IEEE Standard for Low-Rate Wireless Networks," *IEEE Std 802.15.4-2020 (Revision of IEEE Std 802.15.4-2015)*, pp. 1–800, 2020.
- [9] "Gradient: A Platform for Building and Scaling Real-world Machine Learning Applications," <https://gradient.run/>, Accessed: 30-11-2023.
- [10] "OpenAI: Build next-gen apps with OpenAI's powerful models," <https://openai.com/>, Accessed: 01-11-2023.
- [11] I. Mavromatis, Y. Jin *et al.*, "UMBRELLA: A One-stop Shop Bridging the Gap from Lab to Real-World IoT Experimentation," *arXiv: Networking and Internet Architecture*, vol. abs/2401.14829 [cs], Jan. 2024.
- [12] T. Farnham, S. Jones *et al.*, "UMBRELLA Collaborative Robotics Testbed and IoT Platform," in *Proc. of IEEE CCNC 2021*, Jan. 2021, pp. 1–7.
- [13] I. Mavromatis, A. Stanoev *et al.*, "A Dataset of Images of Public Streetlights with Operational Monitoring using Computer Vision Techniques," *Data in Brief*, vol. 45, p. 108658, Oct. 2022.
- [14] Z. Liang *et al.*, "Demo: BuildTwin: Towards Real-time High-fidelity Digital Twin for Smart Building Management," in *Proc. of IEEE ICNP*, 2023.
- [15] S. Anwar, P. E. Carnelli, and A. Khan, "Methods and Systems for Remote Training of a Machine Learning Model," May 2022, US Patent App. 16/952,308.
- [16] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms," *arXiv preprint arXiv:1708.07747*, 2017.
- [17] C.-W. Tien *et al.*, "Kubanomaly: Anomaly detection for the docker orchestration platform with neural network approaches," *Engineering reports*, 2019.
- [18] J. Pope *et al.*, "Container escape detection for edge devices," in *Proc. of the 19th ACM SenSys*, 2021, p. 532–536.
- [19] P. Li and A. Aijaz, "Open RAN meets Semantic Communications: A Synergistic Match for Open, Intelligent, and Knowledge-Driven 6G," in *Proc. of IEEE CSCN Conf.*, 2023.
- [20] P. Li *et al.*, "RLOps: Development Life-Cycle of Reinforcement Learning Aided Open RAN," *IEEE Access*, vol. 10, pp. 113 808–113 826, 2022.