



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

Towards Collective Sentiment Analysis in IoT-Enabled Scenarios

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Savaglio C., Casadei R., Manzoni P., Viroli M., Fortino G. (2023). Towards Collective Sentiment Analysis in IoT-Enabled Scenarios. Institute of Electrical and Electronics Engineers Inc. [10.1109/DCOSS-IoT58021.2023.00118].

Availability:

This version is available at: <https://hdl.handle.net/11585/955659> since: 2024-02-05

Published:

DOI: <http://doi.org/10.1109/DCOSS-IoT58021.2023.00118>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

Towards Collective Sentiment Analysis in IoT-enabled Scenarios

1st Claudio Savaglio
DIMES Dept.
Università della Calabria
Rende, Italy
csavaglio@dimes.unical.it

2nd Roberto Casadei
DISI Dept.
Università di Bologna
Cesena, Italy
roby.casadei@unibo.it

3rd Pietro Manzoni
DISCA Dept.
Universidad Politécnica de Valencia
Valencia, Spain
pmanzoni@disca.upv.es

4th Mirko Viroli
DISI Dept.
Università di Bologna
Cesena, Italy
mirko.viroli@unibo.it

5th Giancarlo Fortino
DIMES Dept.
Università della Calabria
Rende, Italy
giancarlo.fortino@unical.it

Abstract—An interesting and innovative activity in Collective Intelligence systems is Sentiment Analysis (SA) which, starting from users’ feedback, aims to identify their opinion about a specific subject, for example in order to develop/improve/customize products and services. The feedback gathering, however, is complex, time-consuming, and often invasive, possibly resulting in decreased truthfulness and reliability for its outcome. Moreover, the subsequent feedback processing may suffer from scalability, cost, and privacy issues when the sample size is large or the data to be processed is sensitive. Internet of Things (IoT) and Edge Intelligence (EI) can greatly help in both aspects by providing, respectively, a pervasive and transparent way to collect a huge amount of heterogeneous data from users (e.g., audio, images, video, etc.) and an efficient, low-cost, and privacy-preserving solution to locally analyze them without resorting to Cloud computing-based platforms. Therefore, in this paper we outline an innovative collective SA system which leverages on IoT and EI (specifically, TinyML techniques and the EdgeImpulse platform) to gather and immediately process audio in the proximity of entities-of-interest in order to determine whether audience’ opinions are positive, negative, or neutral. The architecture of the proposed system, exemplified in a museum use case, is presented, and a preliminary, yet very promising, implementation is shown, revealing interesting insights towards its full development.

Index Terms—Internet of Things, Edge Intelligence, TinyML, Sentiment Analysis, Collective Intelligence

I. INTRODUCTION

Analyzing the level of satisfaction of a product or service offered is fundamental for product quality enhancement [1], i.e., to promote its general appreciation from its target audience. Public and private bodies in today’s society are hence motivated to collect, among other data, users’ feedback aiming to improve their offer, to be more competitive on the market, to raise their social reputation, or, more in general, to make informed decisions. The availability of feedback of good

quality as well as quantity, i.e., numerous, significant and reliable, is therefore a strategic goal and their processing paves the way towards an effective and actionable form of Collective Intelligence (CI) [2], [3].

User feedback can assume different forms (e.g., free-text reviews, star-based scores, phone interviews, etc.). Traditionally, user feedback was mainly gathered by means of explicit interactions between the users and the service/product providers, and, at the end, it was manually analyzed and evaluated [4]. Consider, for example, the request to insert a vote after having been in a restaurant or through interaction with a push-button panel at the exit of a museum. More recently, crowd-sourcing platforms [5] have been introduced to automate the collection of feedback from a large number of users and to stimulate discussions, brainstorming, etc. However, the aforementioned approaches for feedback gathering very often lack spontaneity (e.g., a user might be hesitant to provide negative feedback), are prone to misinterpretations (e.g., a user might not have all the relevant information or context to judge), and might be biased by collateral emotional state (e.g., a phone interview might bother a user during the work-time or resulting invasive). Conversely, it has been proved that *implicit* feedback, i.e., information left unintentionally and contextually by the users, grants a higher reliability and truthfulness, since the users are focused on the subject and they express opinions without filters [6].

In such direction, the Internet of Things (IoT) [7] provides an exceptional support for the implementation of pervasive infrastructures of heterogeneous sensors able to transparently capture the most natural kinds of human expressions, namely verbal and body language. Indeed, by means of cheap IoT devices (e.g., single-board computers augmented with cameras and microphones), it is easy to record audio as well as to catch images and video, and automatically treat them for further (semi-)automatic analysis. For instance, audio can be recorded,

This work was partially supported by the Italian MIUR, PRIN 2017 Project “Fluidware” (CUP H24I17000070001) and by the Italian MUR, PRIN 2022 Project “COMMON-WEARS” (CUP 2020HCWWLP).

automatically transcript and processed for feeding Sentiment Analysis (SA) systems, also referred to as Opinion Mining, to eventually determine the overall sentiment or emotional tone of the content [8]. By using Natural Language Processing (NLP) and Machine Learning (ML) techniques, it is possible to identify and extract subjective information from text data, thus inferring opinion on a particular topic, brand, or product—even without an explicit interaction. Such an innovative approach to user feedback gathering and analysis not only grants reliability, but also scalability, since a huge amount of data can be processed with minimal human intervention. Moreover, if this process takes place “in-situ” through techniques of Edge Intelligence (EI) [9], [10], namely without storing and analyzing raw data on remote servers, it results into efficient (by locality) and privacy-preserving data mining [11]. Precisely, this last aspect is extremely relevant as privacy fosters a safe and comfortable environment for users to share their honest opinions while protecting sensitive information and ensuring legal and ethical compliance.

Based on these considerations, in this work we design an innovative system for collective SA which leverages on IoT and EI (specifically, TinyML techniques [12] and the EdgeImpulse tool [13]) to catch and immediately process audio in the proximity of entities-of-interest, aiming to transparently gather users’ feedback. In particular, we focus on a museum setting where IoT devices are located in strategic points and they extrapolate target keywords (e.g., typical appreciation/critical words or expression) recorded by their microphones to determine whether users’ opinion are positive, negative, or neutral. Obviously, our solution is general-purpose and it can be easily extended/re-used in any context where CI can raise (e.g., national parks, fairs, exhibitions) or integrated with other forms of user-generated contents (e.g., social media posts, news articles, customer reviews, etc.). To the best of our knowledge, whereas the building blocks of our solutions are well-established paradigms (SA, IoT, EI), this is the first attempt towards the development of a such designed system.

The rest of the manuscript is organized as follows. Section II provides background and related work on Collective/Social EI and SA. Section III provides a brief description of our reference scenario (i.e., the Smart Museum), while the architecture of the proposed system and details of its early implementation with preliminary results are reported respectively in Section IV and Section V. Section VI closes the paper with concluding remarks and future works.

II. RELATED WORK

In this section, we cover related work, first on collective/social EI (Section II-A) and then on SA (Section II-B).

A. Collective/Social Edge Intelligence

Collective/Social intelligence refers to the ability of a group or community to collectively solve problems, make decisions, and achieve goals that are beyond the capabilities of individual members [2], [3]. Typical platforms and services for Collective/Social intelligence exploit the Cloud Computing

technology with remote servers, provided with relevant storage and processing power, able to offer robust, scalable, and distributed data management features. However, this work is motivated by efficiency and privacy goals for which such solutions are not appropriate.

Indeed, more recently, pivoted around the Edge Intelligence (EI) [9], [10] concepts (namely, the deployment of Artificial Intelligence (AI) and ML algorithms on devices located close to the source of data/event generation), the idea of letting CI emerge at the edge has arisen. Hence, novel approaches of *social edge intelligence* [14] (also involving a human-centric design) and *collective edge intelligence* [15] (where groups of software agents are involved) have gained traction, enabled especially by the rapid advancements of *TinyML techniques* [12]. TinyML promotes the provisioning of ML services in resource-constrained device without recurring to powerful servers and high-latency cloud computing premises, thus offering the potential to enable a wide range of applications (predictive maintenance, remote sensing, environmental monitoring, health monitoring, smart agriculture, and more) where small footprint, low power consumption, low-latency processing, robustness, privacy preservation and reliability are key requirements. Although the pair EI-Collective/Social Intelligence has great potential, research in this field is still at an initial stage. In this work, we consider these techniques for the design a Sentiment Analysis system.

B. Sentiment Analysis

Sentiment Analysis (SA), or *opinion mining* [8], revolves around the study of the sentiment or opinion of people towards entities or objects of various kinds. A computer-based system can support SA by providing functionalities for the gathering and processing of user feedback. In this paper, we propose a SA system that is, crucially, transparent, privacy-preserving, and scalable, by leveraging IoT and EI. In the literature, other solutions have been proposed, based on different assumptions, goals, or solution components.

For instance, there are works that use public channels for gathering user-provided feedback. *SACI (Sentiment Analysis by Collective Inspection)* [16] is an unsupervised approach for extracting the collective sentiment through analysis of social media textual content. However, it leverages social media content, which means that the user feedback is explicitly provided and also there is no privacy in general. Approaches for (both offline and online) *privacy-preserving processing* of social media content are also possible, leveraging *anonymisation algorithms* [17], such as in [18], where the authors provide a proof-of-concept for SA-aimed anonymisation of edge data-streams. However, though the approach is privacy-preserving, it leverages Twitter content, i.e., text manually provided by users, whereas our approach provides implicit or transparent acquisition of user feedback.

There also exist work focussing on *implicit aspect extraction* in SA [19], [20], i.e., for extracting opinions that are not mentioned explicitly by the user in its text but rather implied implicitly. In our work, the “implicitness” is taken a step

further, where we also mean that the user does not provide explicit feedback, but rather just acts (potentially unaware of its contributions), providing feedback with no conscious effort.

Finally, that are also contributions devising effective and efficient machine-learning algorithms and models for sentiment analysis, e.g., based on NLP of textual content [21]. Our focus, instead, is on an overall system architecture for supporting end-to-end SA at the edge.

III. REFERENCE SCENARIO: THE SMART MUSEUM

Cultural Heritage has recently become a relevant playground for the IoT technologies [22]: indeed, an increasing number of IoT-based museum services have been designed, mostly exploiting proximity sensors and wearable devices for enabling visitors' indoor localization or content provision/enhancement through a virtual reality [22], [23]. The hence considered "dumb" infrastructure (e.g., NFC tags, Bluetooth beacons, etc.) can be configured to interact with the mobile (smart) devices of visitors to enhance their experience. Differently, to showcase the potential of our approach mixing IoT, EI and SA, we outlined a truly **Smart Museum** in which the exhibited works represent **entities-of-interest** and therefore are individually monitored through proper devices. A **point-of-interest**, instead, is a physical environment (e.g., a room or a corridor) within whose borders some entities-of-interest are collected according to a given criterion (e.g., all the works of a given author or century). The Smart Museum, therefore, can be seen as an ensemble of points-of-interest and the goal of its **administrator** is to understand the opinions of the museum **visitors** about an individual exhibited work or a collection without any explicit interaction nor the exploitation of their own devices. As primary source for visitors' feedback, among the possible ones, we opt for visitors' verbal actions: we focus on audio since it is, with respect to video, lightweight, cheap, and easy to gather and process, while keeping similar truthfulness and reliability. Audio can also be integrated at a later stage with other implicit or explicit, conventional feedback/sources (e.g., text reviews or posts on social networks or crowd-sourcing platforms) as well as with other information (e.g., the number of the visitors of the museum or of a specific entity-of-interest—obtained, for instance, by counting people from video snapshots). The domain model of our use case is reported in Fig.1.

IV. PROPOSED SYSTEM

The proposed system for realizing the Smart Museum according to the aforementioned desiderata leverages on the following components:

- **Edge Devices**, namely IoT devices such as single-board computers associated to a given entity-of-interest and provided with embedded/external sensors (like microphones or cameras) for transparently capturing visitors' feedback (audio, gestures, etc.) and transmitting a report.
- **Edge Servers**, like mini-computers, notebooks, or other workstations, able to collect the reports from the Edge

Devices located in its monitored point-of-interest and to infer visitors' opinion or sentiment.

- **Remote Servers**, provided with relevant hardware resources and Internet connectivity, aiming to store the information provided by all the Edge Servers monitoring the various points-of-interest of the Smart Museum and to make it available to the administrator for visualization and further complex processing (for example, by exploiting historical data).

From such a description, it clearly emerges that instances of these three groups can be functionally and geographically arranged in a two-layer IoT architecture as in Fig.2 where (i) Edge Devices and Edge Servers constitute the **Edge Layer**, being deployed in the proximity of the entities-of-interests and connected through local (wired or wireless) connections for performing their lightweight "local" tasks of user feedback capturing, gathering and SA; while (ii) Remote Servers constitute the **Far-Edge Layer**, since they can be deployed anywhere to perform heavyweight tasks such as data analytics and visualization, models updating, feedback integration resorting on external sources etc.

Hence, the operation of the system can be traced back to five different moments:

- 1) *Keywords selection and recording.* Using dictionaries is likely the simplest possible way to perform SA. In our case, we aim at selecting a set of candidate keywords we repute relevant for a feedback and which, therefore, should be certainly captured within the visitors' comment. In such directions, an adequate dataset of keywords (performed through the Edge Devices by speakers purposely selected) should be preliminary built for supporting the subsequent steps aimed at the audio classification. Different languages, in the case of museums of international interest, should be also considered. Moreover, audio should be recorded from multiple directions and at multiple distances from the Edge Device's microphones.
- 2) *Model creation and deployment.* A ML model for audio classification needs to be created atop the keywords dataset, being effective for correctly classifying terms despite different accents and tones of voice but, at the same time, being optimized and quantized for matching with the resources of the Edge Devices. Thanks to such a model, keyword are recognized from live audio of museum visitors and a report is sent to the Edge Server for SA. The creation of such a model is not a trivial task since museums, either indoor or outdoor, are usually crowded and so background noise or overlapping voices can greatly affect the quality of audio data.
- 3) *Model updating.* Remote Server can improve the performance of the local models for the audio classifications in a privacy-preserving way by means of federated learning approaches. Indeed, local models (without original data) can be shared by the Edge Servers, aggregated on the Remote Server and updated to create a global one, more




Fig. 1. Domain model of the Smart Museum showing main relationships among considered actors, entities and devices

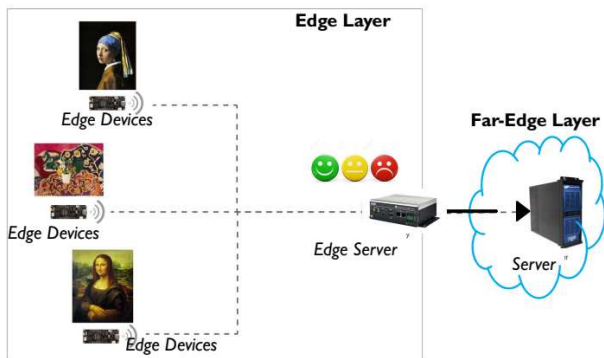


Fig. 2. Two-Layered System Architecture

effective and efficient than the original local ones.

- 4) *Sentiment analysis and feedback integration.* By having provided a score to each selected keyword, it is possible to assess the polarity of the visitors with respect to an individual entity-of-interest or a point-of-interest. This “local” SA can be strengthened through the Remote Servers, with other explicit feedback from outer sources (e.g., tags on the social networks, posts on crowdsourcing platforms), with historical or current data about the number of visitors, or with traditional feedback methods (e.g., push-button panels). The obtained information can be additionally exploited also to take future informed decision about the management of the museum itself (e.g., rearrangement of exhibited works or of the museum spaces).

The outlined system is (i) *effective*, since it is suitable to carry out an implicit gathering of the heterogeneous feedback through pervasive yet non-invasive devices; (ii) *efficient*, since it has a low cost both for the hardware and the (local) com-

munication infrastructure, without resorting to remote servers or consuming bandwidth for the transmission of raw data; and (iii) *privacy-preserving*, because first the devices “hear” everything but “listen” only to the terms useful for the SA (that is, taken from a predefined vocabulary) without the possibility of linking them to a specific subject, and then they process the audio *in-situ*, thus avoiding data traveling on the network or passing through external servers.

The outlined system architecture as well as the listed design challenges have driven the early implementation phase focused on the first two points and reported in the following section.

V. EARLY IMPLEMENTATION AND PRELIMINARY RESULTS

Currently, we focused on the implementation of the system components and activities enabling the audio classification. Indeed, the goal of this early implementation is to preliminary assess if an Edge Device is able, simultaneously, to record audio and run an accurate model for its classification. We left the implementation of the other activities, e.g., the SA-related ones, carried out on devices without particular resource limitations (i.e., Edge Server and Remote Server), as future work.

First, the **Portenta H7** [24] developer board of the Arduino family has been designated as Edge Device: provided with two parallel cores and a graphics accelerator, it simultaneously runs high-level code along with real-time tasks while the onboard wireless module allows to simultaneously manage WiFi and Bluetooth connectivity. Moreover, augmented with the add-on board Arduino Portenta Vision Shield, the Portenta H7 can exploit also a camera module with an ultralow power CMOS Image Sensor (Himax HM-01B0), a LoRa connectivity module and an ultra-compact, low-power, omnidirectional, digital MEMS microphone (MP34DT05). The latter, in particular, is featured by a low-distortion with a 64 dB signal-to-noise ratio and -26 dBFS \pm 3dB sensitivity, which makes it suitable for human voice recording. Due to these features, its versatility, small dimensions (to maintain the spontaneity of data collection, edge devices must be as least visible and

invasive as possible), and low cost, these devices result very suitable for being deployed close to an entity-of-interest in order to record audio of excellent quality.

The audio collection as well as the model creation and deployment have been carried out mainly through **EdgeImpulse** [13]—see Fig.3. This is, indeed, an EI software platform which provides a set of tools and services for collecting and preprocessing sensor data, training, and optimizing TinyML models, and deploying them to microcontrollers or other Edge Devices. It includes a Graphical User Interface (GUI) for building and testing models, as well as command-line tools and Application Program Interfaces (APIs) for integrating with other development workflows, reason why it has been installed over a notebook acting as Edge Server. EdgeImpulse results particularly usable thanks to its integration with popular hardware platforms like Arduino, Raspberry Pi, and NVIDIA Jetson, and it provides a library of pre-built models for common use cases like anomaly detection and classification. For our purposes, we designed an “impulse”, namely a workflow made of multiple tasks, which: (i) sequentially resizes the captured audio data; (ii) processes it to extract its features (the Mel Frequency Cepstral Coefficients MFCCs have been used due to their suitability for human voice analysis); and (iii) classify it through a Convolutional Neural Network (CNN) made with Keras. This has been initially implemented using the EdgeImpulse GUI and, for matching with the Portenta resources, exported to **TensorFlowLite** [25] for being properly configured (from the convolution layer to the full connected layer) and optimized through the ADaptive Moment estimation (ADAM) method. Indeed, TensorFlow Lite is a subset of the TensorFlow library specifically purposed to run inference on microcontrollers and other resource-constrained (i.e., few few kilobytes of memory) IoT devices, without the support of operating system, any standard C or C++ libraries, or dynamic memory allocation.

A. Preliminary results

Aiming to preliminary assess the suitability of the considered hardware/software setting, we performed an initial evaluation of the system by focusing on the most innovative as well as sensitive activity, namely the audio classification through Edge Devices. As a first operation, we considered a list of twelve keywords in the Italian and Spanish languages and we asked 8 peoples of different countries, ages, and genders to pronounce each keyword ten times, being located in each of the 9 different positions with respect to monitored work item and the Edge Device’s microphone. Having built our draft dataset through EdgeImpulse, we went through the model creation by reaching, after a set of experiments, the following draft configuration for our CNN: 1-dimension network with both pool size and stride parameter equal to 2, dropout value 0.2, batch normalization for every stage, learning rate of 0.005, number of filter 32 and kernel size 3. Even in this early implementation, the model has demonstrated excellent behavior in terms of computational efficiency and quality of results, thus resulting suitable for being exploited over the

considered Edge Device. Indeed, an accuracy of 93% and a leak parameter of 0.43 were reported on the Test Set. The model performs the inference in an estimated time between 500 and 600 milliseconds; the memory occupied largely respects the constraints of the Portenta card. By simulating the real environment, with the use of different voices and various types of noise, an accuracy of 83.75% was reported in the various tests, showing an effective response to real use cases.

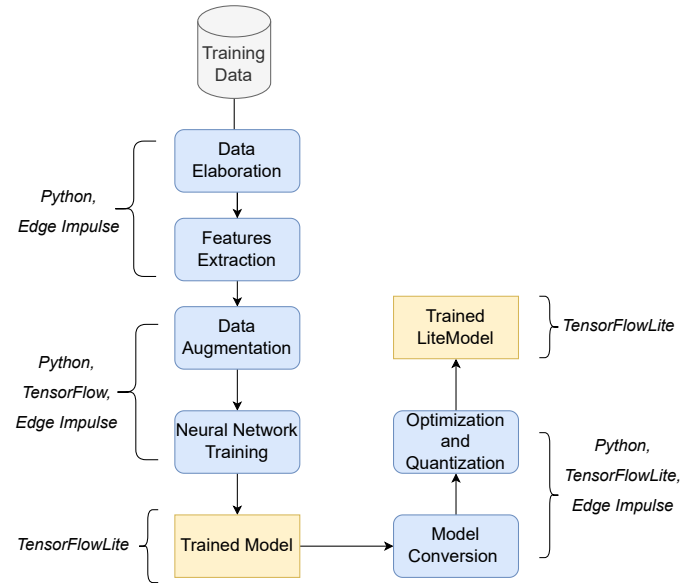


Fig. 3. Model Development for Audio Classification

VI. CONCLUSION AND FUTURE WORK

Collective intelligence can be applied to user feedback analysis by leveraging the insights and perspectives of a diverse group of users to inform product/service improvements.

IoT technology and EI enable implicit, yet privacy-preserving, feedback gathering and processing: this type of innovative approach makes it possible to overcome the limits imposed by traditional assessment tools which depend on explicit, and therefore not always spontaneous, interactions with the customer. Exactly in the direction of a Collective Edge Intelligence, in this paper we presented an early implementation of a Smart Museum where IoT devices co-located with the exhibited work item capture the visitors’ audio and locally perform SA. Preliminary tests verify the robustness of the system. The results obtained are significant: the model reached 93.00% of accuracy on the test set, while in the live tests on the device it reached values equal to 83.75%. In the latter, however, there is a greater level of complexity, which strictly depends on the algorithm by which data is acquired through the microphone.

As future work, we plan to address current limitations and complete the Smart Museum development, specifically by:

- enriching the keywords selection and its dataset by involving more speakers;

- enhancing the context-awareness of the system but still preserving users' privacy;
- providing a more detailed description of the domain, according, for example, to the high-level metamodels of [26], also for fostering the integration of the Smart Museum in Digital Libraries and other Web-based systems;
- identifying, testing and embedding suitable SA models in the system;
- estimating the deployment performance of smart collective services at the edge according to the methodology and simulation-based tool-chain presented in [27].

ACKNOWLEDGMENT

We acknowledge the support of the PNRR project FAIR - Future AI Research (PE00000013), Spoke 9 - Green-aware AI, under the NRRP MUR program funded by the NextGenerationEU. We also thank Paolo Fusca for his help in the early development activities.

REFERENCES

- [1] G. Köksal, I. Batmaz, and M. C. Testik, "A review of data mining applications for quality improvement in manufacturing industry," *Expert Syst. Appl.*, vol. 38, no. 10, pp. 13448–13467, 2011. doi: 10.1016/j.eswa.2011.04.063. [Online]. Available: <https://doi.org/10.1016/j.eswa.2011.04.063>
- [2] S. Alag, *Collective intelligence in action*. Simon and Schuster, 2008.
- [3] R. Casadei, "Artificial collective intelligence engineering: a survey of concepts and perspectives," 2023. doi: 10.48550/arXiv.2304.05147
- [4] R. Santos, E. C. Groen, and K. Villela, "An overview of user feedback classification approaches." in *REFSQ Workshops*, vol. 3, 2019, pp. 357–369.
- [5] S. S. Bhatti, X. Gao, and G. Chen, "General framework, opportunities and challenges for crowdsourcing techniques: A comprehensive survey," *J. Syst. Softw.*, vol. 167, p. 110611, 2020. doi: 10.1016/j.jss.2020.110611. [Online]. Available: <https://doi.org/10.1016/j.jss.2020.110611>
- [6] G. Jawaheer, P. Weller, and P. Kostkova, "Modeling user preferences in recommender systems: A classification framework for explicit and implicit user feedback," *ACM Transactions on Interactive Intelligent Systems (TiiS)*, vol. 4, no. 2, pp. 1–26, 2014.
- [7] L. Atzori, A. Iera, and G. Morabito, "The internet of things: A survey," *Computer networks*, vol. 54, no. 15, pp. 2787–2805, 2010.
- [8] B. Liu and L. Zhang, "A survey of opinion mining and sentiment analysis," in *Mining text data*. Springer, 2012, pp. 415–463.
- [9] V. Barbuto, C. Savaglio, M. Chen, and G. Fortino, "Disclosing edge intelligence: A systematic meta-survey," *Big Data and Cognitive Computing*, vol. 7, no. 1, p. 44, 2023.
- [10] S. Deng, H. Zhao, W. Fang, J. Yin, S. Dustdar, and A. Y. Zomaya, "Edge intelligence: The confluence of edge computing and artificial intelligence," *IEEE Internet Things J.*, vol. 7, no. 8, pp. 7457–7469, 2020. doi: 10.1109/JIOT.2020.2984887. [Online]. Available: <https://doi.org/10.1109/JIOT.2020.2984887>
- [11] C. Savaglio, P. Gerace, G. Di Fatta, and G. Fortino, "Data mining at the iot edge," in *2019 28th International Conference on Computer Communication and Networks (ICCCN)*. IEEE, 2019, pp. 1–6.
- [12] L. Dutta and S. Bharali, "Tinyml meets iot: A comprehensive survey," *Internet Things*, vol. 16, p. 100461, 2021. doi: 10.1016/j.iot.2021.100461. [Online]. Available: <https://doi.org/10.1016/j.iot.2021.100461>
- [13] S. Hymel, C. R. Banbury, D. Situnayake, A. Elium, C. Ward, M. Kelcey, M. Baaijens, M. Majchrzycki, J. Plunkett, D. Tischler, A. Grande, L. Moreau, D. Maslov, A. Beavis, J. Jongboom, and V. J. Reddi, "Edge impulse: An ml ops platform for tiny machine learning," *CoRR*, vol. abs/2212.03332, 2022. doi: 10.48550/arXiv.2212.03332. [Online]. Available: <https://doi.org/10.48550/arXiv.2212.03332>
- [14] D. Wang, D. Zhang, Y. Zhang, M. T. Rashid, L. Shang, and N. Wei, "Social edge intelligence: Integrating human and artificial intelligence at the edge," in *2019 IEEE First International Conference on Cognitive Machine Intelligence (CogMI), Los Angeles, CA, USA, December 12-14, 2019*. IEEE, 2019. doi: 10.1109/CogMI48466.2019.00036 pp. 194–201. [Online]. Available: <https://doi.org/10.1109/CogMI48466.2019.00036>
- [15] R. Casadei, M. Viroli, G. Audrito, D. Pianini, and F. Damiani, "Engineering collective intelligence at the edge with aggregate processes," *Eng. Appl. Artif. Intell.*, vol. 97, p. 104081, 2021. doi: 10.1016/j.engappai.2020.104081. [Online]. Available: <https://doi.org/10.1016/j.engappai.2020.104081>
- [16] L. Rocha, F. Mourão, T. Silveira, R. Chaves, G. Sá, F. Teixeira, R. Vieira, and R. Ferreira, "SACI: sentiment analysis by collective inspection on social media content," *J. Web Semant.*, vol. 34, pp. 27–39, 2015. doi: 10.1016/j.websem.2015.05.006. [Online]. Available: <https://doi.org/10.1016/j.websem.2015.05.006>
- [17] M. Cunha, R. Mendes, and J. P. Vilela, "A survey of privacy-preserving mechanisms for heterogeneous data types," *Comput. Sci. Rev.*, vol. 41, p. 100403, 2021. doi: 10.1016/j.cosrev.2021.100403. [Online]. Available: <https://doi.org/10.1016/j.cosrev.2021.100403>
- [18] A. Pandya, P. Kostakos, H. Mehmood, M. Cortés, E. Gilman, M. Oussalah, and S. Pirttikangas, "Privacy preserving sentiment analysis on multiple edge data streams with apache nifi," in *European Intelligence and Security Informatics Conference, EISIC 2019, Oulu, Finland, November 26-27, 2019*. IEEE, 2019. doi: 10.1109/EISIC49498.2019.9108851 pp. 130–133. [Online]. Available: <https://doi.org/10.1109/EISIC49498.2019.9108851>
- [19] M. Tubishat, N. Idris, and M. A. M. Abushariah, "Implicit aspect extraction in sentiment analysis: Review, taxonomy, opportunities, and open challenges," *Inf. Process. Manag.*, vol. 54, no. 4, pp. 545–563, 2018. doi: 10.1016/j.ipm.2018.03.008. [Online]. Available: <https://doi.org/10.1016/j.ipm.2018.03.008>
- [20] J. Z. Maitama, N. Idris, A. Abdi, L. Shuib, and R. Fauzi, "A systematic review on implicit and explicit aspect extraction in sentiment analysis," *IEEE Access*, vol. 8, pp. 194166–194191, 2020. doi: 10.1109/ACCESS.2020.3031217. [Online]. Available: <https://doi.org/10.1109/ACCESS.2020.3031217>
- [21] H. Sankar, V. Subramaniaswamy, V. Vijayakumar, A. K. Sangaiah, L. Ravi, and A. Umamakeswari, "Intelligent sentiment analysis approach using edge computing-based deep learning technique," *Softw. Pract. Exp.*, vol. 50, no. 5, pp. 645–657, 2020. doi: 10.1002/spe.2687. [Online]. Available: <https://doi.org/10.1002/spe.2687>
- [22] A. Chianese and F. Piccialli, "Designing a smart museum: When cultural heritage joins iot," in *2014 Eighth International Conference on Next Generation Mobile Apps, Services and Technologies*, 2014. doi: 10.1109/NGMAST.2014.21 pp. 300–306.
- [23] S. Alletto, R. Cucchiara, G. Del Fiore, L. Mainetti, V. Mighali, L. Patrono, and G. Serra, "An indoor location-aware system for an iot-based smart museum," *IEEE Internet of Things Journal*, vol. 3, no. 2, pp. 244–253, 2015.
- [24] "Portenta h7," accessed on 2023-04-14. [Online]. Available: <https://docs.arduino.cc/hardware/portenta-h7>
- [25] "Tensorflow lite," accessed on 2023-04-14. [Online]. Available: <https://www.tensorflow.org/lite>
- [26] G. Fortino, A. Rovella, W. Russo, and C. Savaglio, "Including cyber-physical smart objects into digital libraries," in *Internet and Distributed Computing Systems: 7th International Conference, IDCS 2014, Calabria, Italy, September 22-24, 2014. Proceedings 7*. Springer, 2014, pp. 147–158.
- [27] R. Casadei, G. Fortino, D. Pianini, A. Placuzzi, C. Savaglio, and M. Viroli, "A methodology and simulation-based toolchain for estimating deployment performance of smart collective services at the edge," *IEEE Internet Things J.*, vol. 9, no. 20, pp. 20136–20148, 2022. doi: 10.1109/JIOT.2022.3172470. [Online]. Available: <https://doi.org/10.1109/JIOT.2022.3172470>