

Designing an ML-Friendly Wireless Physical Layer for Low-Power IoT

Akshay Gadre

*Electrical and Computer Engineering
Carnegie Mellon University
agadre@andrew.cmu.edu*

Swarun Kumar

*Electrical and Computer Engineering
Carnegie Mellon University
swarun@cmu.edu*

Abstract—With the advent of low-power Internet of Things (IoT), there is an increase in interest for designing inference systems in the cloud that aggregate and perform machine learning tasks from the low-power sensor data. Yet, unlike traditional mobile devices, low-power clients are too battery constrained to transmit large amounts of data within short time spans, as needed for many complex inference models.

In this paper, we present a vision for bridging the gap between the power-starved low-power clients and the data-starved inference engines in the cloud. We present a mechanism that takes into account the battery life of these clients and how it affects the traditional inference models.

Index Terms—LP-WANs, Machine Learning, data aggregation

I. INTRODUCTION

Low-Power Wide-Area Networks (LP-WANs) are empowering the urban and rural environments by enabling various applications such as metering, asset tracking and smart irrigation on farms. These networks connect ten-year battery powered devices to base stations several miles away, with each device transmitting a few kbps at maximum. These low-power sensors have created a new opportunity for leveraging compute in the cloud to develop novel inference models for various tasks. For e.g. "Which areas in the city are flooded?", "Is the air pollution sufficiently low outside for a patient with breathing difficulties?", "Should I irrigate my farm today?" are some questions that can be answered by deploying multiple inexpensive sensors across a city (or a farm) and using the data to make recommendations.

However, as the number of LP-WAN sensors scale, it becomes increasingly challenging to get frequent information from each and every sensor. Indeed, our recent work [1] has shown that querying 10,000 clients using traditional approaches may even take several hours. Despite recent advances [2] in developing novel MAC layers for LP-WAN technologies, the high latency in data retrieval at scale remains an issue.

Another important issue when dealing with low-power clients is the energy burden of communicating the information. Typically, these clients can provide a 10 year battery life on an AA battery albeit transmitting small amount of data every

15 minutes. This becomes worse as the range of the clients increases, lowering their data rate and in turn worsening their battery life. Thus, there is an important need to rethink the physical layer of low-power clients with a focus specifically on rapid machine learning inference.

In this paper, we present a vision for developing inference solutions in the era of low-power IoT clients and address the above issues in developing machine learning models in the cloud. We build on our recent work, QuAiL [1], on developing a PHY-layer for aggregating weighted linear combinations from low-power clients. This work develops a PHY-layer mechanism for low-power clients at scale to respond to various aggregate spatial, statistical and inference queries within the duration of a single packet. While our prior work [1] is primarily focused on data aggregation for statistics, this paper specifically expands on its impact on the performance and battery-drain when performing Machine Learning inference. While our approach is broadly applicable across various Machine Learning algorithms (e.g. SVM, naive Bayes, etc.), we focus on neural networks – one of the most general of such approaches.

The key idea behind our approach is to replace one of the layers of neural networks with a process that occurs naturally in wireless networks. To better understand our approach, recall that neural networks first construct linear combinations of the inputs before deriving useful inferences. Our solution replaces this linear combination phase of the first layer by a natural wireless phenomenon - how signals combine in the air. We force LP-WAN clients to transmit their signals concurrently following a query from the base station. The base station receives a linear combination of the overlapping signals weighted by wireless channels. Since transmissions from a large number of clients transmit concurrently, their signal powers add up at the base station. This allows clients to save precious battery-life by addressing key bottlenecks – lower air time and transmit power [3], [4]. The received signals are then passed along through the remaining layers of the neural network for performing the required inference task. The rest of this paper describes the various considerations in making such a design practical.

The first of those considerations is the ability of machine learning models such as neural networks (NNs) to be trained to maximize the robustness at the first layer. Indeed, while

This work is funded by NSF grants 1837607, 1942902, IoT@CyLab and Kavcic Moura Endowment Fund.




Fig. 1: Our approach co-optimizes the PHY layer and machine learning for inference on sensed data in LP-WAN to conserve client battery.

measuring these weighted linear combinations in the presence of noise, the received measurements suffer from additional errors reducing the accuracy of the inference models trained assuming ideal data. We present how these errors can be modeled and optimized for high-fidelity inference on sensor data despite the error overheads of quick aggregation schemes.

Another key consideration is that of client battery life – specifically the trade-off between system performance and battery-drain of the low-power clients. We analyze this trade-off and demonstrate how by reducing the transmit power can enable clients to even enable $7\times$ their typical battery life¹.

The rest of the paper is organized as follows: Sec. II presents related work on machine learning for low-power IoT, Sec. III gives a brief outline of our system, Sec. IV presents how machine learning can be trained differently to optimize for our approach, and Sec. V highlights the energy-accuracy tradeoff for machine learning applications leveraging the above approach. Sec. VI presents experimental results.

II. RELATED WORK

Related work falls broadly in three categories:

LP-WANs: LP-WAN deployments in both licensed (NB-IoT [5]) and unlicensed (LoRaWAN [6]) frequencies have witnessed rapid deployment globally [7], [8]. Towards developing these technologies, there has been much work done on developing novel approaches for synchronization [9], [10], association [11], [12], power optimization [3], [13], communication technology [2], [14], and power adaptation [15]. Despite these advances, it has been experimentally observed that the 10 year battery life is only achievable at extremely low reporting rates [3], [4], [13] with the synchronization overhead as the key bottleneck for smaller data reporting intervals. There has also been much work done on developing other LP-WAN technologies such as LoRaWAN [6] and SigFox [16], which provide significantly larger battery lives due to lack of requirement of backward compatibility with cellular infrastructure.

Our work aims to reduce the power burden of LP-WAN clients without sacrificing it altogether, when performing machine learning inference. We achieve this via a novel approach

to send aggregate information across clients in the network, each transmitting at a reduced power.

Machine Learning with Sensors: There has been much work on machine learning inference based on sensed data [17]–[21]. The emergence of smart cities and LP-WANs have enabled city scale sensing applications such as environmental monitoring [22], precision irrigation for parks [23], parking management [24], and other participatory sensing applications [25]. Many of these applications use of data from widely deployed sensors to make predictions about the state of the system. We specifically focus on tasks that can be posed as a probabilistic classification problem [26] over sensed data. This framework can answer both statistical questions, e.g. “what fraction of homes are cold today?” as well as more intricate inference questions, e.g. “is this home occupied based on its electricity and water usage?”. Our objective is to complement existing work on statistical inference on wide-area IoT data with a new approach to significantly reduce the time and battery-drain needed to collect this data.

Wireless coding in the air: Coding in the air using wireless signals has been deeply studied by wireless researchers over multiple decades [27] including leveraging MIMO [28], both in the WiFi [29] and LP-WAN contexts [13]. There has also been extensive work done on developing novel data accumulation techniques by clever encoding mechanisms [30] and optimal forwarding to minimize energy expenditure [31]. However, prior work in this space aims to retrieve individually sensed data at the base station, while our work aims to perform an inference task over aggregated sensed data.

There has also been much work done on developing compressive sensing in wireless sensor networks [32] where clients perform pre-processing to reduce communicated information. Other researchers have used various other lossy coding schemes [33] to optimize for energy by cutting redundancy. While our approach builds upon these solutions, it specifically optimizes neural-network based probabilistic classification on a low-power wide-area network.

III. PHY LAYER DESIGN OVERVIEW

This section provides a brief overview of how our architecture (see Fig. 1) can perform probabilistic neural-network based classification and regression on sensed values from LP-WAN clients while lowering their power consumption. Most LP-WAN clients are battery-starved, low data-rate transmitters deployed over a wide-area motivating the need for a time and energy efficient solution.

Our architecture (Fig. 2) has three steps: (1) The base station queries the network, addressing the clients that must transmit messages concurrently – either their sensed information (e.g. temperature, humidity) or diagnostics (e.g. queue size). The base station provides parameters dictating the timing and power scaling of the user’s transmission; (2) Our approach then designs a transmission mechanism that allows clients to transmit their sensed information rapidly and at a fraction of their typical battery overhead. We do so while remain-

¹Assuming inference is the typical operation of these clients.

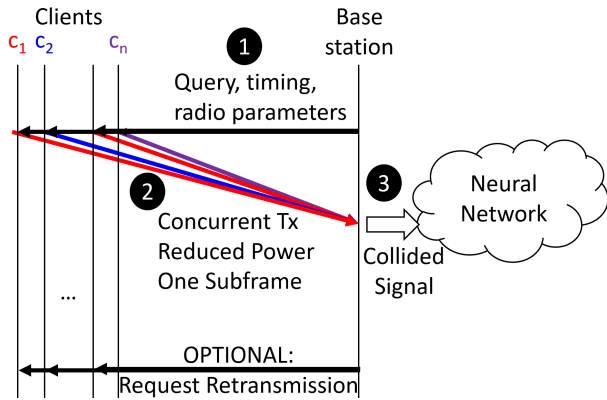


Fig. 2: System Architecture

ing compatible with the architecture of LoRa and NB-IoT, two common LP-WAN technologies. (3) Upon receiving the transmissions, our system runs the probabilistic classification algorithm and outputs the probability of different classes inferred by the neural network classifier.

Our key innovation is the co-optimization of neural-network based classification model and the LP-WAN client transmissions to minimize client battery-life and airtime. Traditionally, one would perform neural network inference by first collecting data individually from LP-WAN clients and later performing classification in the cloud. In contrast, our system replaces the very first layer of the neural network responding to the query by another natural linear process – collisions in the air. Specifically, our approach: (1) First requests LP-WAN clients to transmit synchronously within a single frame but at a reduced transmit power; (2) Next, in doing so, collisions from a large number of clients will add up in power over the air. This allows each client to save battery-life by reducing transmit power. Further, it delivers desired data in a time and spectrum efficient manner (within a few milliseconds of a frame). (3) Finally, we show how processing the resulting linear combination of signals at the base station can still solve a wide range of learning problems with minimal impact on accuracy. While the main challenges and solution to those challenges have been detailed in [1], some of the salient features, particularly of significance to machine learning are:

Compatibility with NB-IoT and LoRaWAN: While collisions are typical in LoRaWAN, cellular technologies such as NB-IoT are explicitly designed to avoid them. In light of these observations, our approach engineers collisions of NB-IoT clients in a unique part of the cellular PHY which is purely distributed and where collisions are common – the random access channel (RACH). Recall that to transmit any data, a client first needs to advertise its existence and request for time-frequency resources. Clients do this by sending a (randomly chosen) identifier on the random access channel, following which they are allotted resources by the base station. Each identifier is a unique frequency hopping code among a set of well-defined orthogonal codes [11]. Traditionally, two transmitters that transmit concurrently only collide when they

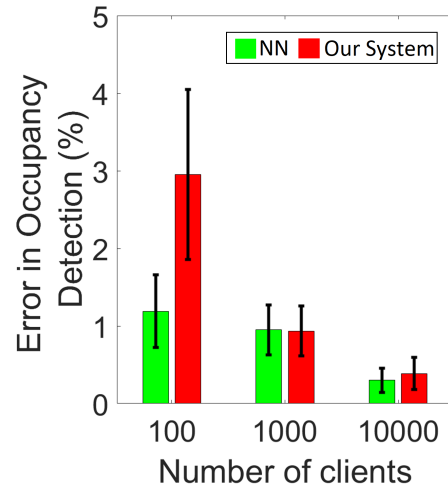


Fig. 3: Error in Occupancy detection reduces with larger number of clients

choose the same (random) identifier. Typically, this leads to packet loss and retransmission by both clients.

Our system chooses to re-use the random access channel to engineer collisions between NB-IoT clients. Specifically, each client modulates the power of the codes it transmits on the RACH, based on sensed data. Each client chooses the set of codes to transmit as well as any weights applied to them based on a predefined set of weights supplied by the base station. The base station therefore receives a linear combination of the clients’ signals weighted by the wireless channels.

Dealing with negative weights (typical in most inference models even at the first layer): Our system strives to make collisions of signals over the air emulate this very first layer of the neural network. To do so, it intelligently sets the weights applied to each client’s transmission per-code w_{ij} to emulate the precise linear combination that occurs in a neural network. In effect this makes the resulting output equivalent to the desired linear combination from a neural network’s first layer, modulo noise. A key advantage of this approach is that it does not require the neural network to be re-trained each time on the wireless channels.

However, an important problem remains in making the above design practical: while weights trained by a neural network can be positive or negative real numbers, weights used in our approach can only be positive. Recall that our solution relies on the linear combination of signal power, which is an inherently positive number. We address this challenge by moving the burden of subtraction from the air to the base station. Specifically, our system subdivides the codes of the RACH into two sets: positive codes and negative codes. It then asks each client to transmit using a positive code should the learned neural network weight be positive and the corresponding negative code otherwise. At the receiver, our approach simply subtracts the power received over each negative code from the positive code to retrieve the desired input for the neural network.

IV. IMPACT ON MACHINE LEARNING INFERENCE

In this section, we discuss how the design of a machine learning model can influence the accuracy of our distributed PHY-layer aggregation scheme.

A. Training Optimizations:

Traditional machine learning algorithms are trained to minimize classification/estimation error. While this works well where the input features are clean and without noise, doing so over low-power sensors may make the model highly volatile to variance in these cheap sensors. Further, many of these sensor behaviors change across environments, temperature and wear-and-tear.

There are two ways, training of machine learning model can be improved, to make the system robust to noisy aggregates:

(a) Prioritize gap over margin for classification: Traditional machine learning approaches attempt to maximize the classification margin for all training samples retrieved from the training dataset. For example, support vector machines are developed keeping in mind that the only points that govern the decision boundary between the two classes are those on the border between the two classes. This is then extended to non-linearly separable classes by using a higher dimension kernel to make these classes separable in that domain. However, there is an inherent assumption that accuracy is the only metric of interest.

Another critical metric is robustness to noise. Indeed, much of the information retrieved from these data retrieval approaches can highly vary rendering almost all training data highly noisy. While mapping it to a higher domain is the right approach for a hundred percent linear separability, it is more important to maximize the gap for the most points as well. Achieving this would require adding an additional term in the optimization function that penalizes the inverse of the gap for all training points classified correctly given a margin. This would provide an effective tradeoff parameter that can be controlled to estimate resilience to error in linear aggregates.

(b) Evolved Backpropagation in Neural Networks: Another important impact of our approach is upon multi-layer inference models such as neural nets or convolutional neural networks. The traditional assumption is that inside the network there is no additional noise being added. However, atypical approaches such as our own break that assumption. Indeed, if we use our solution to retrieve the linear combinations at the first layer of the neural network, there is an additional noise being added to the input of the next layer due to additive wireless noise. To address this additional challenge, the neural network optimization function will need to be designed to maximize robustness at the input of the second layer instead of the first layer.

B. Communicating models for Federated Learning:

An important question to address would be how does one learn models remotely with recent work on low-power machine learning enabled hardware and aggregate them at the base station. Indeed, federated machine learning is being

touted as a leading approach for many tasks involving mobile nodes.

Fortunately, our approach can assist the server in aggregating these small models learned at the end devices by providing weighted linear aggregates as a service. Further, many complex aggregations can be mapped to linear addition by a simple transform (for example, multiplication by log addition) enabling even non-linear aggregation for the client models.

V. ENERGY VS. ACCURACY TRADE-OFF

In this section, we study how we can retrieve major power savings for these low-power IoT clients while maintaining a low-error threshold for many applications involving complex machine learning models, including neural networks.

Power Savings At this point, should multiple clients transmit codes, the received signals may add-up constructively or destructively in-phase. However, it is well-known that as the number of interfering clients increases the power of the received signals will asymptotically add up [34], while the noise values will not. As a result, the net signal-to-noise ratio at the receiver along any code j is:

$$\text{SNR}_j = \frac{\sum_i |h_i|^2 |s_i|^2 w_{ij}}{\sigma^2}$$

Where σ^2 denotes noise power. It is easy to see that on average, signal-to-noise ratio scales linearly in the number of clients. Conversely, increasing the number of clients by any factor f allows each client to also reduce its transmit power by the same factor f without any loss in signal quality, reducing battery drain. This means that clients can dynamically scale down their transmit power in proportion to the number of clients in the network, as indicated by the base station, leading to significant battery benefits (See Sec. VI). Additional improvements in power savings are incurred by avoiding synchronization overhead. This effectively ensures that client spends an extremely small duration of time in transmit mode, and as a result, achieving significant battery life gains.

VI. EVALUATION

We implement our approach using FSK radios on Semtech SX1276 radios to communicate with an Ettus USRP N210 emulating the base station. We collect more than 150,000 GPS-location and time stamped channel measurements across 3 km² area including geographical obstacles such as large buildings, hills, rivers in the City of Pittsburgh. We then emulate collisions using the collected channels to evaluate our system at scale.

We build a light-weight NB-IoT stack including the NP-RACH in C++/Gnuradio on the USRP base station. Our experiments first train weights (w_{ij}) of the machine learning model under consideration based on raw data. These weights are then provided to the clients for transmitting the appropriate energy across codes. Note that all the random matrices (M) applied to the weights are initialized using a seed provided during association. The base station receives the signals at




Fig. 4: Battery Life Improvement for various categories of clients

the base stations and measures the power for all the codes to decode them. This decoded vector is then passed on to the neural network for performing the desired inference task.

We measure three quantities of interest: (1) Quality of the solution against the ground-truth (pre-labeled data); (2) Battery-savings using the battery model in [13] and (3) Air-time. As a baseline, we consider a machine learning model that uses raw sensed data transmitted over NB-IoT.

We evaluate the battery benefits provided by our system as the number of clients increase. We evaluate the system for emulated collisions of 100, 1000 and 10000 clients with the occupancy dataset to estimate the average number of rooms occupied. This statistic is of interest to assess the impact in case of an emergency in a large urban space. Our evaluation is based on three publicly available datasets for sensor-based inference – the Gait dataset [35], Occupancy dataset [36] and Intel-Berkeley dataset [37].

Effect on Quality: As shown in Fig 3, when we increase the number of clients, the average quality of occupancy detection improves with an error of just 0.3% at 10,000 clients. This is due to the fact that as a greater number of clients participate, you have more information to estimate the number of occupied rooms and lesser room for error.

Effect on Battery Life: As shown in Fig 4, when we increase the number of clients, the average battery life improvement increases (albeit modestly) and is up to 3.13 \times for 10,000 clients. We note that the majority of the power savings of our approach stems from the absence of need for uplink synchronization and the short airtime. Any remaining benefits seen across number of clients is due to the lower transmit power per-client as more clients participate.

Battery Life Benefits across applications: We evaluate our system’s ability to perform weighted aggregation for three real world neural-network applications: (1) Estimating the distribution of activities being done by 100 users wearing

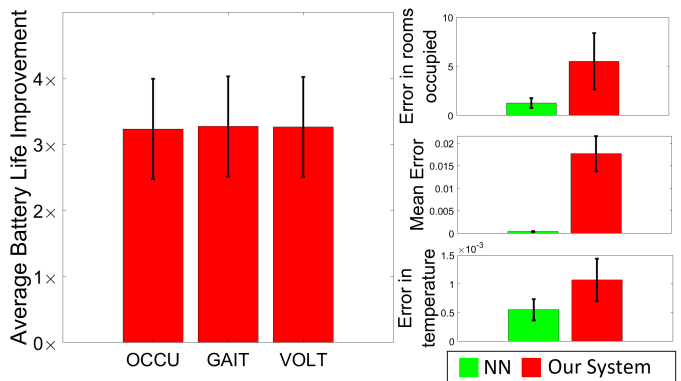


Fig. 5: Battery Life gains and effect on quality of estimation for various neural network applications

various sensors. (2) Estimating the occupancy of rooms using a client with four sensors in each room providing data. (3) Estimating the mean temperature of a room based on the voltage fluctuations in the sensors in that same room. Note that the above applications do not have simple models which can be applied and hence the models must be learned.

As shown in Fig 5, we achieve a substantial increase in the battery lives by up to 3.25 \times of the clients involved in all three tasks at modest performance cost. We see that for the occupancy task, we are able to assess the number of occupied rooms with a precision of five rooms instead of one room as the neural network does. In Gait identification of 100 clients, we are able to estimate the distribution of clients in each category to within 1%. Finally, our solution demonstrates about a 0.1% error in estimating the mean temperature of a room based on voltage fluctuations of 100 clients.

VII. CONCLUSION AND FUTURE WORK

In this paper, we presented a vision for developing inference solutions in the era of low-power IoT clients and address the above issues in developing machine learning models in the cloud. Our solution builds upon our recent work on developing a machine learning friendly PHY-layer for retrieving information from the clients. We then show how we can cleverly design machine learning models to optimize for lower errors over such schemes where the error likelihood is significantly larger at the first layer. We show how by designing our models correctly, we can optimize for energy trade-offs in the client demonstrating significant power benefits over naïve approaches for inference from LP-WAN clients. While this paper focuses on LP-WAN technologies, we believe that developing a general purpose low-power PHY layer co-optimized for machine learning in the IoT context remains an important open problem.

Acknowledgements: This work was supported by NSF grants 1942902, 1837607, 1646235, IoT@CyLab and the Kavcic-Moura Fund.

REFERENCES

- [1] A. Gadre, F. Yi, A. Rowe, B. Iannucci, and S. Kumar, "Quick (and Dirty) Aggregate Queries on Low-Power WANs," in *ACM/IEEE IPSN*, 2020.
- [2] R. Eletreby, D. Zhang, S. Kumar, and O. Yağın, "Empowering low-power wide area networks in urban settings," in *ACM SIGCOMM*, 2017, pp. 309–321.
- [3] P. Andres-Maldonado, P. Ameigeiras, J. Prados-Garzon, J. Navarro-Ortiz, and J. M. Lopez-Soler, "Narrowband iot data transmission procedures for massive machine-type communications," *IEEE Network*, vol. 31, no. 6, pp. 8–15, 2017.
- [4] 3GPP RP-151393, "NB LTE - Battery Lifetime Evaluation," <https://portal.3gpp.org/ngppapp/CreateTdoc.aspx?mode=view&contributionId=672720>, 2015, accessed: 2019-01-30.
- [5] 3GPP, "NarrowBand - Internet of Things," <https://www.gsma.com/iot/narrow-band-internet-of-things-nb-iot/>, 2018, accessed: 2019-01-30.
- [6] LoRa Alliance, "What is the LoRaWAN Specification?" <https://loralliance.org/about-lorawan>, 2016, accessed: 2019-01-30.
- [7] T-Mobile, "T-Mobile Launches America's First Nationwide Narrowband IoT Network," <https://www.t-mobile.com/news/americas-first-narrowband-iot-network>, 2018, accessed: 2019-01-30.
- [8] R. Grech, "Semtech and Comcast's machineQ Announce LoRaWAN Network Availability in 10 Cities," <https://www.semtech.com/company/press/semtech-and-comcasts-machineq-announce-lorawan-network-availability-in-10-cities>, 2018.
- [9] A. Ali and W. Hamouda, "On the cell search and initial synchronization for nb-iot lte systems," *IEEE Communications Letters*, vol. 21, no. 8, pp. 1843–1846, 2017.
- [10] C. G. Ramirez, A. Sergejev, A. Dyussenova, and B. Iannucci, "Long-shot: long-range synchronization of time," in *ACM/IEEE IPSN*, 2019, pp. 289–300.
- [11] X. Lin, A. Adhikary, and Y.-P. E. Wang, "Random access preamble design and detection for 3gpp narrowband iot systems," *IEEE Wireless Commun. Letters*, vol. 5, no. 6, pp. 640–643, 2016.
- [12] A. Gadre, R. Narayanan, A. Luong, S. Kumar, B. Iannucci, and A. Rowe, "Frequency Configuration for Low-Power Wide-Area Networks in a Heartbeat," in *USENIX NSDI*, 2020.
- [13] A. Dongare, R. Narayanan, A. Gadre, A. Luong, A. Balanuta, S. Kumar, B. Iannucci, and A. Rowe, "Charm: exploiting geographical diversity through coherent combining in low-power wide-area networks," in *IPSN*. IEEE Press, 2018, pp. 60–71.
- [14] R. Ratasuk, B. Vejlgaard, N. Mangalvedhe, and A. Ghosh, "Nb-iot system for m2m communication," in *Wireless Communications and Networking Conference (WCNC), 2016 IEEE*. IEEE, 2016, pp. 1–5.
- [15] C. Yu, L. Yu, Y. Wu, Y. He, and Q. Lu, "Uplink scheduling and link adaptation for narrowband internet of things systems," *IEEE Access*, vol. 5, pp. 1724–1734, 2017.
- [16] "Sigfox," Accessed: 2019-01-30. [Online]. Available: <https://www.sigfox.com/>
- [17] Intercomp, "Triumph in Dubai with tests for the development of NB - IoT technologies applied to parking sensors," <https://www.intercomp.it/nb-iot-technologies-applied-to-parking-sensors/?lang=en>, 2017, accessed: 2019-01-30.
- [23] L. Sanchez, L. Muñoz, J. A. Galache, P. Sotres, J. R. Santana, V. Gutierrez, R. Ramdhany, A. Gluhak, S. Krco, E. Theodoridis *et al.*, "Smart-santander: Iot experimentation over a smart city testbed," *Computer Networks*, vol. 61, pp. 217–238, 2014.
- [18] S. online, "Waste-management system rides narrowband iot network," <https://goo.gl/oHCgFK>, 2018, accessed: 2019-01-30.
- [19] GSMA, "Smart agriculture proves a natural use of NB-IoT," <https://www.gsma.com/iot/news/smart-agriculture-nb-iot/>, 2018, accessed: 2019-01-30.
- [20] Ericsson, "The world's first cellular IoT-based smart factory," <https://www.ericsson.com/assets/local/networks/offerings/cellular-iot/connected-factory-case-study.pdf>, 2018, accessed: 2019-01-30.
- [21] IoT for all, "IoT in the Workplace: Smart Office Applications for Better Productivity," <https://www.iotforall.com/iot-smart-office-applications/>, 2018, accessed: 2019-01-30.
- [22] J. M. Talavera, L. E. Tobón, J. A. Gómez, M. A. Culman, J. M. Aranda, D. T. Parra, L. A. Quiroz, A. Hoyos, and L. E. Garreta, "Review of iot applications in agro-industrial and environmental fields," *Computers and Electronics in Agriculture*, vol. 142, pp. 283–297, 2017.
- [24] Smart Parking Systems, "Great success of NB-IoT technology tests in Madrid in Huawei laboratories," <http://smartparkingsystems.com/en/nb-iot-sensors-applied-to-parking-system/>, 2018, accessed: 2019-01-30.
- [25] Huawei, "Smart Shared Bicycle Lock," <https://www.huawei.com/minisite/iot/en/smart-bike-sharing.html>, 2017, accessed: 2019-01-30.
- [26] J. Friedman, T. Hastie, and R. Tibshirani, *The elements of statistical learning*. Springer series in statistics New York, NY, USA., 2001, vol. 1, no. 10.
- [27] T. S. Rappaport *et al.*, *Wireless communications: principles and practice*. prentice hall PTR New Jersey, 1996, vol. 2.
- [28] H. Lipfert, "Mimo ofdm space time coding–spatial multiplexing, increasing performance and spectral efficiency in wireless systems, part i technical basis (technical report)," *Institut für Rundfunktechnik*, 2007.
- [29] L. Hanzo, Y. Akhtman, J. Akhtman, L. Wang, and M. Jiang, *MIMO-OFDM for LTE, WiFi and WiMAX: Coherent versus non-coherent and cooperative turbo transceivers*. John Wiley & Sons, 2010.
- [30] R. Rajagopalan and P. K. Varshney, "Data aggregation techniques in sensor networks: A survey," 2006.
- [31] H. Ö. Tan and I. Körpeoğlu, "Power efficient data gathering and aggregation in wireless sensor networks," *ACM Sigmod Record*, vol. 32, no. 4, pp. 66–71, 2003.
- [32] W. Bajwa, J. Haupt, A. Sayeed, and R. Nowak, "Compressive wireless sensing," in *IPSN*. ACM, 2006, pp. 134–142.
- [33] X. Zhang and B. Li, "Optimized multipath network coding in lossy wireless networks," *IEEE journal on selected areas in communications*, vol. 27, no. 5, pp. 622–634, 2009.
- [34] H. Rahul, H. Hassanieh, and D. Katabi, "Sourcesync: a distributed wireless architecture for exploiting sender diversity," *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4, pp. 171–182, 2011.
- [35] R. Chereshevnev and A. Kertész-Farkas, "HuGaDB: Human Gait Database for Activity Recognition from Wearable Inertial Sensor Networks," in *International Conference on Analysis of Images, Social Networks and Texts*. Springer, 2017, pp. 131–141.
- [36] L. M. Candanedo and V. Feldheim, "Accurate occupancy detection of an office room from light, temperature, humidity and co2 measurements using statistical learning models," *Energy and Buildings*, vol. 112, pp. 28–39, 2016.
- [37] S. Madden, "Intel Berkeley Research Lab Dataset," <http://db.csail.mit.edu/labdata/labdata.html>, 2004, accessed: 2019-01-30.