

Poster: Maestro – An Ambient Sensing Platform With Active Learning To Enable Smart Applications

Tahiya Chowdhury^{1*}, Murtadha Aldeer^{2*}, Shantanu Laghate¹, Justin Yu², Qizhen Ding¹, Joseph Florentine², Jorge Ortiz²

¹ECE Department, Rutgers University, USA ²WINLAB, Rutgers University, USA

tahiya.chowdhury,shantanu.laghate, justin.yu1,qizhen.ding,joseph.florentine,jorge.ortiz@rutgers.edu, maldeer@winlab.rutgers.edu

Abstract

Smart ambient sensing applications are built on classifiers that are trained to detect different events in a physical space. Training these requires human labeling in controlled experiments. However, controlled experiments only capture specific experimental settings and application-specific labels. We aim to build a framework for data collection and active labeling that 1) reduces the number of labels necessary to maximize event coverage and 2) continuously learns the underlying distribution of different events. System *Maestro* is a data collection and labeling framework that senses the environment across 5 different ambient sensors producing 18 channel measurements. Maestro includes a web interface for continuous labeling and applies active learning with label propagation to minimize the number of necessary labels. We present the results of an initial deployment in a student apartment, where Maestro continuously learns to count occupants and progressively learns to identify activities of daily living. Our preliminary results show that we can achieve accuracy >95% for these applications, with <10% labeled examples.

1 Introduction

The emergence of the Internet-of-Things (IoT) has made the vision of smart buildings a reality. A smart building environment is equipped with a range of sensors and captures the dynamics of different activities in the environment. These data can be used to infer physical context by extracting characteristic properties for a variety of machine learning (ML) applications, such as occupancy detection [5], activity monitoring [1], and environment control [2].

However, current ML techniques do not generalize well

* Both are primary co-authors.

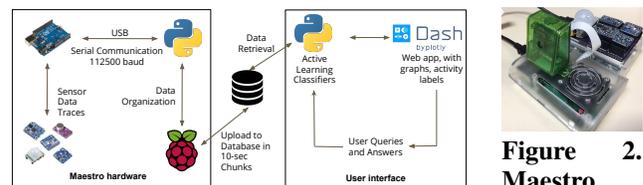


Figure 1. Overview of Maestro

Figure 2. Maestro prototype

across different spaces as generalization typically requires a large amount of labeled data across different spaces [6], which poses several challenges for ubiquitous-computing based sensing platforms. First, the ML component of such platforms provides sub-optimal performance without adequate training data and the data is proportional to number of users in the system. Poor performing systems fail to attract more users, creating a dilemma for system performance. Second, current dominant ML methods learn in isolation; given a labeled dataset, it learns to solve a well-defined, narrow task. However, effective learning takes place over time and can enable learning with fewer examples by re-using accumulated knowledge to label novel, unseen examples. Third, most ML objectives revolve around detecting and classifying raw sensor streams to physical events between humans and the environment. Hence, the low-level sensor data (e.g. vibration) can detect low level events (such as vibration of paper-towel dispenser) but cannot leverage this detection to a high-level human-understandable concept (e.g. 400 dispensing events meaning the dispenser would require a refill). For IoT systems to find success in ubiquitous computing, the “Semantic Gap” between raw data streams and semantic labels must be bridged [7].

Prior work [3], has employed static and mobile sensors to monitor occupants’ activities. Laput et al. [4] creates a custom sensor and uses manual training and clustering methods to obtain synthetic labels for appliance usage. In this work, we present Maestro, an IoT-system designed for rapid deployment and data collection for ML applications to address these challenges. Maestro allows data collection from multiple sensor-boxes simultaneously to a centralized server. It offers an user-centric, online labeling capability to the user via a web application, provides visualization and context of data streams, to help users assign semantic labels. Thus, Maestro can be a key solution for crowd-sourced label acquisition

Table 1. Results of occupancy counting

Method	Acc.	Prec.	Rec.	F1 score
SV Machine	0.76	0.78	0.76	0.76
Random Forest	0.89	0.90	0.89	0.89
Ensemble	0.75	0.86	0.75	0.77

Table 2. Results of activity recognition

Method	Acc.	Prec.	Rec.	F1 score
SV Machine	0.96	0.96	0.96	0.96
Random Forest	0.98	0.98	0.99	0.98
Ensemble	0.97	0.97	0.97	0.97

from many users to build custom ML applications. Further, it can use active learning, which allows to specifically query labels for most informative samples, and can thus reduce the labeling effort.

2 System Overview

As shown in Figure 1, Maestro consists of two main modules. A sensing module, which is responsible for acquiring data from the sensors. While the database and web application module supports data storage and real-time data visualization via a web-based user interface. A prototype is shown in Figure 2. Maestro is powered by a Raspberry Pi3[®], and the data collection is enabled through an Arduino Uno[®] connected to various off-the-shelf sensors (IMU, PIR, color and illumination sensor, audio sensor, pressure, humidity, and temperature sensor). This combination of sensors results in a total of 18 data points per unit time. The device is configured to upload chunks of data to a centralized database running on our data processing server. A camera is attached to the Raspberry Pi to help the user while labeling, so that a video feed can be shown during the active learning phase. For data storage, we used TimescaleDB to implement our storage database. We also develop a user interface for data visualization and label querying for active learning. The visualizer allows the user to specify a particular box, a start and end time, and what sensor channels of the box to visualize. The active learning presents unlabeled examples to the users and queries the label associated with it. The web interface then displays the graph of all channels specific to the queried examples alongside a GIF of the corresponding camera feed for user reference. The user observes the feed and graph and submits their label for that example, which feeds back into the learning algorithm for further training.

3 Experiments and Preliminary Results

• **Experiments.** We conducted experiments in a typical household of a family of 5 members. One of the family members acted as the experimenter and annotated the data during the experiment period which we use as ground-truth. We focus on 2 problems in our experiments: occupancy counting and activity recognition. Our analysis performed on the raw sensor data which were transformed to values between (0, 1) as a pre-processing step. We use 75/25 train/test split. For occupancy counting, the experimenter annotated the data using corresponding occupancy labels during the experiment period. For activity recognition, we select data for 3 activities that are typical in a household environment: vacuum cleaning, exercising, and typing on a keyboard.

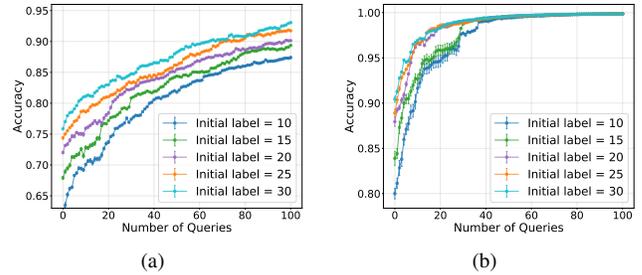


Figure 3. Effect of active learning strategy in (a) occupancy counting and (b) activity recognition

• **Preliminary Results.** We use different ML algorithms appropriate for multi-class classification (Tables 1 and 2). We report accuracy, precision, recall, and F1 score to outline the performance of each classifier. Our results in Table 1 show that Random Forest classifier performs the best in occupancy counting. However, in activity recognition applications, both Random Forest and Ensemble classifier achieve best performance (Table 2). More close observation on the results reveals that only 2-2.5% were misclassified.

We further report the classification performance on the same applications when employing active learning. We explore whether our method can still provide good performance if we lower the required number of labeled data for the classifier. We train initial classifier using a small set of initial labeled training data (ranges between 10-30 samples). The later training examples are specifically queried by sampling method from active learning and can thus decrease the number of labels required and user labeling effort. Based on the selected query strategy, active learning framework iteratively queries for more examples to be labeled and retrains the classifier. We experimented with various query strategies used in the literature and margin sampling performs the best. We allow the framework to query for labeled examples and retrain until it reaches near-convergence. Figure 3 shows the results using active learning. The Y-axis shows the accuracy improvement (with error bar) of the classifier equipped with active learning, against the number of queries it demanded (X axis). For both applications, the performance improved with new labels available from queries.

4 References

- [1] M. Aldeer, et al. Patientsense: Patient discrimination from in-bottle sensors data. In *Proceedings of the 16th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*, MobiQuitous '19, page 143–152, 2019.
- [2] B. Dong, et al. A review of smart building sensing system for better indoor environment control. *Energy and Buildings*, 199:29–46, 2019.
- [3] J. Jiang, et al. Makesense: An iot testbed for social research of indoor activities. *ACM Trans. Internet Things*, 1(3), June 2020.
- [4] G. Laput, et al. Synthetic sensors: Towards general-purpose sensing. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 3986–3999, 2017.
- [5] L. Rueda, et al. A comprehensive review of approaches to building occupancy detection. *Building and Environment*, page 106966, 2020.
- [6] M. Ruta, et al. Machine learning in the internet of things: a semantic-enhanced approach. *Semantic Web*, 10(1):183–204, 2019.
- [7] R. Yus, et al. Abstracting interactions with iot devices towards a semantic vision of smart spaces. In *Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, BuildSys '19, page 91–100, 2019.