

Exploring the Efficacy of Sparse, General-Purpose Sensor Constellations for Wide-Area Activity Sensing

GIERAD LAPUT, Carnegie Mellon University

CHRIS HARRISON, Carnegie Mellon University

Future smart homes, offices, stores and many other environments will increasingly be monitored by distributed sensors, supporting rich, context-sensitive applications. There are two opposing instrumentation approaches. On one end is full sensor saturation, where every object of interest is tagged with a sensor. On the other end, we can imagine a hypothetical, omniscient sensor capable of detecting events throughout an entire building from one location. Neither approach is currently practical, and thus we explore the middle ground between these two extremes: a sparse constellation of sensors working together to provide the benefits of full saturation, but without the social, aesthetic, maintenance and financial drawbacks. More specifically, we target a density of one sensor per room (and less), which means the average home could achieve full coverage with perhaps ten sensors. We quantify and characterize the performance of sparse sensor constellations through deployments across three environments and 67 unique activities. Our results illuminate accuracy implications across key spatial configurations important for enabling more practical, wide-area activity sensing.

CCS Concepts: • **Human-Centered Computing → Ubiquitous and mobile computing; ubiquitous and mobile devices;**

Additional Key Words and Phrases: Ubiquitous sensing; Internet-of-Things; IoT; Smart Environments;

ACM Reference Format:

Gierad Laput and Chris Harrison. 2019. Exploring the Efficacy of Sparse, General-Purpose Sensor Constellations for Wide-Area Activity Sensing. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 2, Article 55 (June 2019) 19 pages.

DOI: 10.1145/3328926

1 INTRODUCTION

“Smart” environments hold great promise to improve our home and work lives by inferring activity and events, supporting context-sensitive agents and responsive interfaces. Today, there are many consumer offerings (e.g., [6][27][37][40]) and research systems (e.g., [4][38][39]) that offer sensor tags for sensing human environments. Most often, these sensors are designed to attach to objects and appliances of interest. For example, Sen.se’s Mother product [37] or Notion’s “all-in-one sensors” [29] can be stuck to a coffee maker, pillbox, or doorway to monitor use. This approach is flexible and can be very accurate, as sensors are physically coupled to the item they are tasked to monitor. However, battery power and the inherent one-sensor-to-one-object relationship is an obstacle to scalability—a house or office might have hundreds of objects of interest.

Author’s addresses: Carnegie Mellon University, Pittsburgh, PA, USA, gierad.laput@cs.cmu.edu; chris.harrison@cs.cmu.edu

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

Copyright © ACM 2019 2474-9567/2019/6-55 \$15.00
<https://doi.org/10.1145/3328926>

Rather than tagging every object, it would be ideal if a single, omniscient sensor could detect everything ongoing in a building. Unfortunately, no such technology exists at the moment. The closest we have to this vision are infrastructure-mediated sensing schemes, which can monitor a wide area from a single instrumented point. For example, HydroSense [14] used a pressure sensor attached to one point on the plumbing to detect showers, toilets flushing, faucets running, and laundry cycles. Although solving the problem of wide-area sensing from a single point is highly practical, flexibility is reduced, as only a certain class of activity is detectable (e.g., water use). To capture other types of events, researchers have examined infrastructure such as HVACs [33], electrical lines [16][18], gas lines [9] and lighting [17].

More recently, researchers have put forward general-purpose sensors, with a stated goal of “one sensor per room” [24]. Using either a camera [23] or a multitude of low-level sensors on a single board [24], these aim to digitize a wide range of environmental facets in a room with no direct instrumentation of objects. Today, devices such as Google Home, Amazon Alexa or Apple HomePod are commercial examples that operate at single room scale, and have the requisite sound sensing capability for activity recognition. These approaches can offer all of the benefits of full sensor saturation, but without the aesthetic, maintenance, and cost drawbacks. However, this prior research has only investigated single room sensing, and did not consider sensing events in neighboring spaces, nor leverage potentially complementary signals if more than one sensor was deployed in e.g., a home (Figure 1). Moreover, prior research has generally been confined to a lab setting and not authentic use environments, and the range of sensed classes has been small. In this work, we directly address these open questions, move beyond lab studies, and seek to quantify four key questions:

Q1: What is the accuracy difference when sensing activities from *inside* a room vs. from *outside* of it?

Q2: Is there a benefit in using data from *another sensor*, for example, a sensor in a *nearby* room?

Q3: What is the effect of using *all available sensor data* at a location (e.g., *every sensor* in a house)?

Q4: What events, if any, are detectable when *only sensor data from outside* of a room is available.

To answer these questions, we crafted a study procedure that allowed us to evaluate different spatial configurations of sensors: (1) in room sensing, (2) in room + nearest room sensing, (3) all room sensing, (4) only-nearest-room sensing, and (5) all-but-in-room sensing. These conditions, in turn, serve as experimental proxies for the questions above. To quantify performance, we ran a deployment study capturing real-world data using thirty general-purpose sensors distributed at three distinct use locations (small business, residential and institutional). At these locations, we studied 100 sensed events (67 unique classes). We also investigate how and what types of signals propagate in typical environments, as well as sensing robustness across different rooms and locations. To facilitate future research, we make our study dataset and evaluation scripts available to download: <https://github.com/FIGLAB/constellations>.

At a high level, we found that in-room sensing (which is the most pervasive contemporary practice) can benefit from sensor data from adjacent rooms. While not a surprising result, we are the first to quantify it in a comprehensive study across 30 rooms and 100 sensed events. We show that in-room accuracy (94.6%) can be improved by 1.8% when leveraging sensor data from a neighboring room (96.4%). Although this accuracy gain appears modest, it is a meaningful (and statistically significant) reduction in classification error. Often one of the most significant challenges in moving technology from the lab and into the real world is closing the gap between 95 and 99% accuracy. Other results are less intuitive, for example, utilizing all sensors in a house can actually harm accuracy due to feature explosion and added noise. We also show that sensing without any sensor in a room (but perhaps surrounded by sensors) remains a stubborn research challenge. As before, we quantify many of these widely held assumptions in order to underscore where research challenges remain.

2 RELATED WORK

We focus on three areas most immediately related to our work: distributed sensing systems, infrastructure-mediated sensing, and general-purpose ubiquitous sensing.

2.1 Distributed Environmental Sensing

Starting from early work on “smart dust” [1] in the late 1990s, researchers have been working on deployments of many networked sensors, forming distributed sensing systems. This approach aims to improve sensing coverage and fidelity across a wide variety of application scenarios [35][36][41]. One of the most popular platforms used by the sensor networking community was the Berkeley Mote and its variants [10][12]. Other approaches rely on heterogeneous distributed sensors, such as those by Logan et al. [26], Cook et al. [10], and Suryadevara et al. [38]. These pioneering research platforms most often rely on batteries and low power RF, which places a cap on computation and network bandwidth. Also related are distributed sensing systems that use mobile phones as sensing platforms (sometimes in concert with WLAN information) to track movement and activities of users [3][22][42].

More recent sensor boards (e.g., [4][6][27][37][40]) have become smaller and more computationally capable, but in general, the overall utility is highly dependent on the quality of the sensor distribution. Achieving sensing saturation often implies a sizable, dense deployment, typically dozens of sensors for a room like a kitchen. This approach can be expensive, with sensors often costing tens of dollars. Furthermore, as sensor density increases, there is a danger of becoming socially and aesthetically obtrusive, especially in contexts such as homes or public spaces (see e.g., [7][20][39] for discussion of these issue).

2.2 Infrastructure-Mediated Sensing

To minimize deployment cost and intrusiveness, researchers have studied the installation of sensors at strategic probe points with rich signals. A canonical example is work by Patel et. al [32], where appliance usage and electrical activity in a home can be detected from a single plug-in sensor. Similar techniques have also been investigated by Abott [1], Hart [18][19], Gupta [16] and others. In these instances, a single sensor can observe activities across an entire house by taking advantage of shared electrical infrastructure.

As mentioned in the introduction, this infrastructure-centric approach has also been applied to other facets, including plumbing [13][14], HVAC [32], gas lines [9] and lighting [17]. By attaching geophones to a building’s superstructure, it is also possible to detect human activity and a range of other environmental events [30]. In these instantiations of infrastructure-mediated sensing, one or a few sensors were installed at strategic probe points, circumventing the need for dense instrumentation. However, one limitation of this approach is that sensors are constrained by the class of infrastructure they are coupled to. For example, a plumbing-attached sensor can detect sink and shower use, but not microwave or TV use. Thus, to achieve more general-purpose sensing, multiple systems would need to be combined – an area of open research.

2.3 General-Purpose Sensing

A single sensor that can detect a wide range of different signals and events might be described as “general-purpose.” Computer vision has perhaps come closest to achieving general-purpose sensing, as cameras offer rich, non-contact signals, which can be processed through e.g., machine learning to yield sensor-like feeds. Achieving human-level abstractions and accuracy is a persistent challenge, leading to the creation of mixed CV- and crowd-powered systems (e.g., [5][15][25]). For example, Zensors [23] aimed to turn smartphone cameras into “universal” sensors, using crowd answers to bootstrap machine learning classifiers. While computer vision

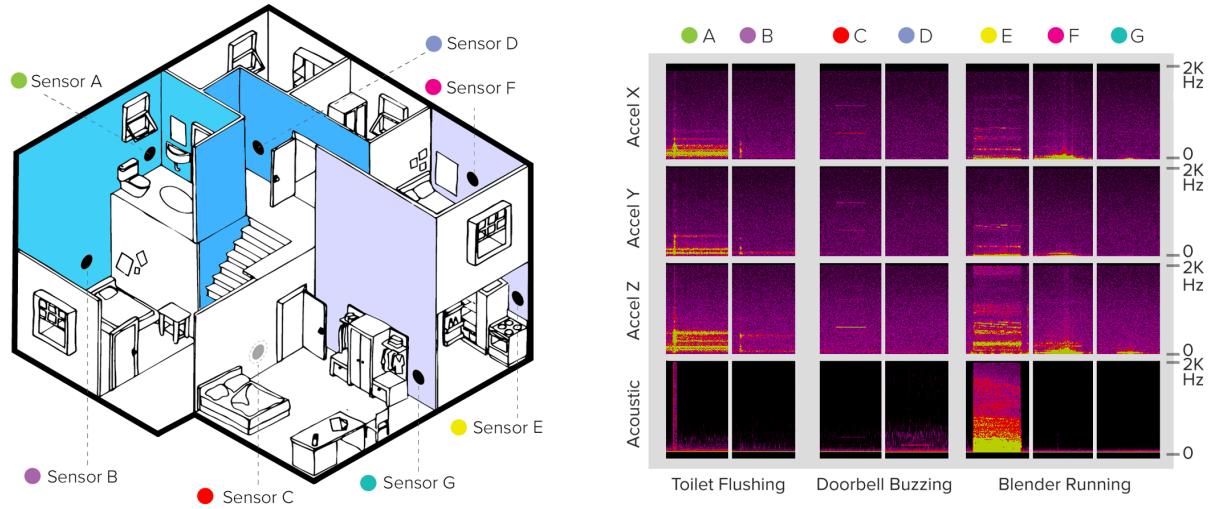


Fig. 1. An example of a sparse constellation setup, wherein a small group of sensors, one per room or less, can detect a wide range of events happening across e.g., an entire home. Sensors that share e.g., a common wall (Sensor A & B, Sensor E & F) or a contiguous space (Sensor C & D) can be used to correlate and confirm events, using signals such as changes in vibration, sound, motion or lighting.

approaches are powerful, cameras have been widely studied for their high level of social intrusiveness [7][20]. For this reason, Synthetic Sensors [24] explored general purpose sensing without cameras, using a custom board capturing 18 distinct sensor channels. We utilize the latter hardware as a vehicle for our later deployment study. Perhaps most critically, previous work on general-purpose sensing has focused on room-scale questions, assessing capabilities of a single sensing device in isolation, but not how two or more sensors could work together, extending coverage and boosting accuracy.

In this work, we build on prior work (specifically [24]), and pursue an investigation (and ultimately quantify) on how *sparse constellations* of general-purpose sensors can work cooperatively to sense events across large environments. In addition to studying densities of “one sensor per room” (speculated, but not tested in [24]), we also explore cases where *no* sensor is present, and classification must instead rely on other sensors in the building. If possible, this would allow for constellation densities of “one sensor per room or less”, making deployments even more practical and expanding sensed-area coverage without additional cost. We also move significantly beyond prior work in terms of scale (e.g., number of contexts, sensed facets, and sensor placement locations) and ecological validity (out of the lab and into real-world environments with e.g., sensor placement wherever wall outlets exist, ambient noise, foot traffic, and a much greater variety of real-world appliances).

3 PROPAGATION OF EVENTS

Physical events manifest as sound, vibration, illumination, temperature and other forms of energy, which are rarely confined to one location. Instead, they propagate outwards, into interior spaces through openings and doors, across walls, and along structural members [30]. Therefore, it is possible for a general-purpose sensor placed in one location to indirectly detect events in another. In theory, this should enable sensor constellations with a sparsity of less than one per room. Furthermore, other sensors can offer confirmatory signals, improving accuracy and robustness.

Figure 1 offers an example of this effect with a sparse sensor constellation deployed in a home. Here, Sensor A (bathroom) and B (lower bedroom) share a common wall, Sensor C (first floor hallway) and D (second floor hallway) share an open space, while Sensor E (kitchen), F (upper bedroom) and G (master bedroom) are simply proximate to one another. When an event occurs, such as a toilet being flushed, it produces characteristic sound and vibrations (Figure 1, Right; Sensors A & B). These physical manifestations are readily captured by the accelerometer and microphone of Sensor A (Figure 1, Sensor A), located in the room, and also by the accelerometer of Sensor B (Figure 1, Sensor B) in the room below and adjacent to the plumbing stack. A similar effect can be seen when the doorbell buzzes (Figure 1, Sensors C & D), or when a blender is running in the kitchen (Figure 1, Sensors E, F & G). Depending on the interior layout, changes in lighting, temperature, motion and other non-vibroacoustic channels can be useful in supporting robust recognition of event.

4 DEPLOYMENT HARDWARE AND SOFTWARE

For deployment, we take advantage of the general-purpose sensing hardware developed in [24], shown in Figure 2. This WiFi-connected board offers an array of low-level sensors (18 channels in total: 3-axis accelerometer, microphone, electromagnetic interference, temperature, humidity, barometer, illumination, RGB light color, infrared motion, 3-axis magnetometer, non-contact thermal, and WiFi RSSI), sampled at high rates and buffered locally (see [24] for full details).

Ten times a second, the board computes and transmits a feature vector derived from its buffered on-board sensors. Specifically, it computes seven statistical features (min, max, sum, mean, standard deviation, range, and centroid) for all 18 channels. Additionally, for its high-sample-rate channels (microphone, electromagnetic interference, and three accelerometer axes), the board also sends 128-length, real-valued FFTs. In total, the board computes 1207 features [24]. This featurized data not only reduces bandwidth, but also denatures data to help preserve privacy. Per sensor data throughput is ~20KBps. One improvement to the sensor firmware was the addition of Network Time Protocol (NTP), which the boards use to accurately timestamp all data transmissions, facilitating multi-sensor data fusion.

To receive data, we built a backend to handle incoming sensor streams (which connect to the Internet over WiFi). Transmitted data is decoded on the server and can be serialized to disk for offline processing or forwarded to a delegate application for real-time visualization or analysis (see Video Figure). The backend takes care of all data synchronization using sensor packet timestamps. For each incoming sensor stream, the backend performs a mean windowing ($w=5$, analogous to average pooling [2]) to reduce noise. Finally, the backend assembles all window-averaged features from each sensor in a constellation into a superset of synchronized features (*i.e.*, 10 sensors \times 1207 features), which serves as input to one or more endpoints for machine learning.

5 CLASSIFIERS

We decided to build separate models per-room (instead a single model per building), as this offers a few advantages. First, per-room machine learning models are easier to train, minimizing combinatorial feature explosion. Second, a per-room model better supports simultaneous event detection across rooms (*i.e.*, models for room A and room B can both trigger independent events). Third, a single building-wide model is harder to scale; a system designer who wants to add two more sensors (*i.e.*, two new rooms) to a building-wide model that was originally trained with data from ten sensors will need to retrain the entire model. Finally, for our implementation, we used a random forest ensemble classifier [8] for its efficient training given the large number of event classifiers we wished to include in our evaluation. Our random forest model was implemented using sci-kit learn [34], keeping most default parameters except for: number of estimators = 20 and max features = 500. As we will

describe in our Results section, we collected data for one week in each location (ten sensor boards), and use day-fold, cross-validation to train and test these models post hoc. Separating the data in this fashion prevents temporal overfitting (e.g., signals collected adjacent in time tend to look similar).



Fig. 2. The general-purpose sensing board [20] used for our sparse constellation deployments.

6 DEPLOYMENT STUDY

Our backend and fleet of synchronized sensors provided a unique technical vehicle to investigate the efficacy of sparse, general-purpose sensor constellations in different spatial configurations. We now describe where we deployed our sensors, what they were tasked with sensing, and our data collection procedure.

6.1 Study Locations

As a test bed, we selected three distinct locations: a *home* (Figure 3 top left), a non-profit *institution* (Figure 3 top right) and a small *business* (Figure 3 bottom). These locations captured diversity in size, layout, construction, function, and appliances available to sense. According to municipal records, the home, institution and business were last renovated, inspected and brought up to code in 2015, 2014 and 2010 respectively. The home and institution both followed the 2012 International Building Code (IBC) [21], while the business followed the 2009 IBC (both of which provide, e.g., standards for electrical outlet spacing). All three structures are typical northeastern US construction, with loadbearing masonry exteriors. The home and business used wood floor joists, while the institutional building used steel beams. For walls, the home used wood studs, while the business and institution locations used metal studs. All three used drywall for wall surfaces. Occupants in each location were briefed about the deployment and signed participation waivers, but otherwise went about their usual activities. Deployed sensors were connected to the local WiFi, which punched through the NAT to our backend.

6.2 Sensed Events

In order for our results to be representative across a wide range of uses, we sought to develop a large and diverse set of sensed events. We started by asking building occupants about environmental facets a “smart building” might wish to know about. From this list, our research team curated a final set, selecting those plausible to sense

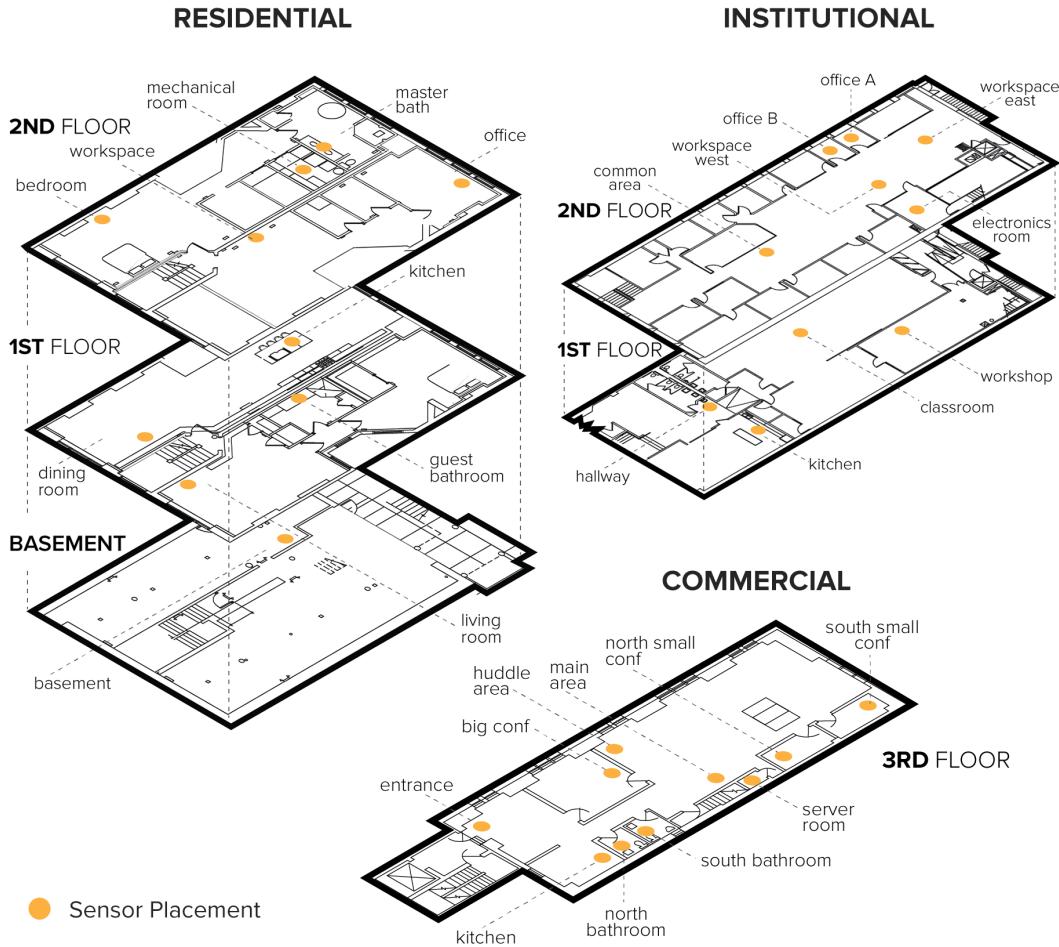


Fig. 3. Floor plan and sensor placements at our three test contexts: Home, Institution, and Business.

with our hardware. In total, we included 100 sensed events, encompassing 67 unique classes, across our three test locations, summarized in Table 1.

6.3 Value of Binary Sensing

We focus on binary event detection (*i.e.*, active or not) for several reasons. Foremost, this paper is first and foremost a study contribution – many research teams have deployed homogenous arrays of sensors for binary activity detection (*e.g.*, [39] – see also Related Work), and we wished our results to offer compatible insights. Second, binary events are logical building blocks for enabling higher level recognition, such as human activity sensing. For example, [24] fed “1st order” binary event classifications into “2nd order” classifiers, which could track values such as count, consumption and state. These outputs, in turn, can then be passed to even higher order sensors, building in

Table 1. Our deployment results broken out by evaluation condition (colorized columns) and sensed event (rows). L column denotes acoustically loud events; V for vibrational events, and S for events requiring line-of-sight.

Room	Sensed Event	L	V	S	In-Room Only	In Room+ Nearest	All Rooms	Nearest Only	All But In-Room	Room	Sensed Event	L	V	S	In-Room Only	In Room+ Nearest	All Rooms	Nearest Only	All But In-Room
Office	Door Closed	✓	✓		99%	100%	96%	82%	33%	Office B	Door Closed	✓	✓		100%	100%	98%	100%	93%
	Printer On	✓			99%	100%	100%	32%	26%		Door Knock	✓	✓		91%	92%	96%	91%	58%
	Space Heater Running			✓	100%	100%	100%	7%	5%		Fridge Closed			✓	84%	87%	92%	84%	56%
	Basement		Back Door Closed	✓	✓	100%	100%	100%	24%	46%	Coffee Grinding	✓	✓		100%	100%	100%	92%	77%
	Bedroom	Door Closed	✓	✓	86%	98%	96%	96%	89%	Drawer Actuated	✓			99%	97%	99%	22%	7%	
	Workspace	Chair Rolling	✓	✓	98%	98%	95%	55%	38%	Faucet Running	✓	✓		100%	100%	100%	35%	15%	
		Computer On	✓		94%	97%	86%	0%	0%	Microwave Running	✓			99%	100%	100%	84%	89%	
		Music Playing	✓		99%	99%	97%	53%	34%	Mouse Moving On Table	✓			84%	95%	92%	30%	9%	
	Dining Room	Door Closed	✓	✓	100%	100%	100%	81%	79%	Objects Placed On Table	✓			90%	92%	88%	27%	7%	
	Guest	Door Closed	✓	✓	96%	95%	96%	2%	18%	Typing On Table	✓			94%	97%	96%	25%	25%	
Bathroom	Left Faucet Running	✓	✓		90%	92%	84%	15%	3%	Lab East	Writing On Table	✓			98%	98%	99%	18%	39%
	Right Faucet Running	✓	✓		79%	88%	80%	9%	1%		Card Swiped	✓			100%	99%	99%	65%	53%
	Shower On	✓	✓		98%	99%	97%	14%	42%		Closet Door Closed	✓	✓		85%	95%	88%	91%	73%
	Toilet Flushing	✓	✓		98%	99%	99%	28%	12%		Door Closed	✓	✓		91%	91%	89%	76%	76%
	Kitchen	Blender On	✓	✓	100%	100%	99%	100%	73%		Air Compressor Running	✓	✓		100%	100%	100%	96%	74%
		Coffee Grinder On	✓		100%	100%	100%	100%	66%		Exhaust Fan On	✓	✓		100%	100%	100%	99%	93%
		Dishwasher On	✓	✓	93%	95%	98%	9%	3%		Laser Exhaust Running	✓	✓		100%	100%	100%	99%	77%
		Faucet On	✓	✓	92%	91%	91%	55%	15%		Mitre Saw Running	✓			100%	100%	100%	86%	57%
		Garbage Disposal On	✓	✓	100%	100%	100%	46%	58%		North Door Closed	✓			98%	97%	86%	94%	85%
		Left Burner On	✓		98%	98%	98%	19%	22%		South Door Closed	✓	✓		99%	99%	99%	100%	83%
Living Room	Right Burner On	✓			98%	98%	98%	4%	4%	Lab West	Front Door Closed	✓	✓		92%	94%	95%	81%	81%
	Toaster Finished	✓			98%	98%	99%	87%	11%		Front Door Buzzer	✓	✓		71%	75%	75%	86%	86%
	Fireplace On	✓			97%	97%	97%	56%	38%		Whiteboard In Use	✓			100%	100%	100%	73%	100%
	TV Playing	✓			97%	97%	98%	49%	44%		Dishwasher Running	✓	✓		98%	100%	100%	100%	100%
	Master Bath	Electric Toothbrush On	✓		98%	98%	98%	55%	23%		Espresso Running	✓	✓		70%	95%	86%	96%	75%
	Hair Dryer On	✓			100%	100%	100%	73%	49%		Faucet Running	✓	✓		65%	95%	94%	86%	85%
	Left Faucet Running	✓	✓		93%	100%	98%	100%	100%		Garbage Disposal On	✓	✓		94%	100%	100%	100%	100%
	Right Faucet Running	✓	✓		88%	98%	94%	86%	7%		Kettle Running	✓	✓		84%	98%	91%	62%	49%
	Shower Running	✓	✓		90%	95%	96%	27%	25%		Keurig Running	✓	✓		95%	93%	95%	68%	46%
	Toilet Flushing	✓	✓		83%	74%	69%	6%	0%		Microwave Running	✓	✓		100%	100%	100%	86%	95%
Mechanical Room	Furnace Running	✓	✓		100%	100%	100%	100%	96%	Business location	Toaster Finished	✓			85%	90%	91%	93%	81%
	Water Heater Running	✓	✓		100%	100%	100%	97%	65%		Paper Shredder On	✓	✓		100%	100%	100%	99%	97%
	Laundry Dryer Running	✓	✓		100%	100%	100%	100%	88%		Printer Drawer Closed	✓			99%	97%	97%	86%	91%
	Utility Sink Running	✓	✓		96%	97%	98%	85%	46%		Printer Printing	✓	✓		100%	100%	100%	88%	63%
	Office A	Door Closed	✓	✓		96%	96%	96%	95%	90%	Scale Key Pressed	✓			91%	95%	94%	54%	9%
Classroom	Door Knock	✓	✓		99%	99%	99%	40%	25%	Institution Location	Door Closed	✓	✓		98%	97%	98%	98%	97%
	Phone Ringing	✓			100%	100%	99%	48%	48%		Faucet Running	✓	✓		96%	100%	100%	55%	46%
	Lecture Talking	✓			99%	99%	100%	11%	32%		Toilet Flushing	✓	✓		93%	100%	99%	89%	75%
	Partition Drag	✓	✓		100%	100%	100%	100%	99%		Server Room	Door Closed	✓		100%	98%	100%	97%	97%
	Projector Running	✓			100%	100%	100%	54%	41%		North Conf	Door Closed	✓		89%	93%	93%	88%	28%
Common Area	Coffee Brewing	✓	✓		99%	100%	100%	31%	35%	All Locations	Whiteboard In Use	✓			87%	92%	93%	45%	53%
	Microwave Running	✓	✓		100%	100%	100%	71%	26%		Door Closed	✓	✓		83%	83%	84%	61%	57%
	Toaster Finished	✓			100%	99%	100%	4%	2%		Whiteboard In Use	✓			99%	99%	99%	55%	41%
	Water Fountain Running	✓			97%	99%	98%	6%	18%		Door Closed	✓	✓		84%	80%	81%	75%	70%
	West Door Closed	✓	✓		97%	100%	94%	98%	79%		Faucet Running	✓	✓		88%	92%	84%	80%	78%
Electronics Room	Drill Press On	✓	✓		86%	86%	86%	44%	33%	Bathroom A	Paper Towel Dispensed	✓			77%	87%	73%	56%	35%
	Extractor Fan On	✓	✓		83%	86%	85%	10%	14%		Toilet Flushing	✓	✓		82%	98%	92%	96%	91%
	Staple Gun Fired	✓			89%	95%	97%	15%	57%		Urinal Flushing	✓	✓		90%	100%	99%	100%	100%
	Water Fountain Running	✓	✓		100%	100%	100%	35%	19%		Background Class				96%	97%	97%	67%	60%
Hallway	Chiller Running	✓	✓		100%	100%	100%	79%	56%	Bathroom B	Average Accuracy				95%	96%	96%	63%	54%
	Door Closed	✓	✓		100%	100%	100%	17%	27%		Standard Deviation				7.1%	5.3%	6.2%	31.4%	30.1%
	Water Fountain Running	✓	✓		100%	100%	100%	35%	19%										

Legend: L = Loud, V = Vibrational, S = Line-of-Sight

abstraction and sophistication. For example, knowing the refrigerator was just opened, that the stove was turned on, and cabinets are being opened, etc. could be passed to a “making dinner” classifier. Critically, small errors accumulate up the sensor “food chain” – a sensing system that is 90% accurate at each abstraction level is less than 75% accurate by the third order. Thus, the quality of such high-level sensing is wholly dependent on underlying (often binary) detections, and thus this is where we focus in this paper.

6.4 Sensor Placement

We deployed ten sensor tags in each study location, at a density of one per room. We define a room as a contiguous area separated physically or functionally by walls, flooring, elevation change, walkways, furniture or other demarcations. Sensor placement followed a strict procedure, to avoid experimenter bias:

- 1) Using a pre-determined list of desired sensed events, find the spatial centroid of the events in the room.
- 2) Plug the sensor into the closest (*i.e.*, Euclidian distance) wall electrical outlet to the centroid.

We believe this placement procedure is both logical and realistic, and we further confirmed with occupants that this is where they would have placed the sensor as well. In some cases, this led to suboptimal placement (e.g., sensor placed farther away from more subtle signals). However, we do not view this as an experimental confound, but rather as ecological validity. We did not know *a priori* which event were going to be harder to sense than others, and more importantly, it is unrealistic to expect end users to move sensors around to improve accuracy. As a result, our study results offer a realistic assessment of accuracy should such sensors be deployed in a logical, but perhaps not optimal way. Moreover, we purposely included a large number of sensed events (100) as a way to assess population level accuracy, which is more robust than drawing conclusions on individual sensor placements or specific events.

6.5 Procedure

We initially considered collecting data serendipitously, as occupants triggered events in the course of their days, which we would label post hoc by e.g., listening to audio. However, there were three significant drawbacks to this procedure. Foremost was privacy – capturing data that facilitated accurate labeling required a high degree of privacy invasion (either a camera or audio feed – neither of which our pilot participants would accept for more than a brief deployment period). Second, even with video and/or audio, we often found it challenging to know exactly what event was happening (faucet running vs. toilet bowl refilling; coffee maker vs. microwave running) due to e.g., visual occlusion and ambiguous sounds. This made building a reliable ground truth for evaluation difficult and laborious. Finally, we found that many interesting events happened infrequently (e.g., paper shredder, bench grinder, toaster), which precluded building a sufficient data set for either training or analysis.

As a necessary experimental compromise, we decided to visit each of the 30 rooms (three locations, ten rooms each) once per day for a week and manually trigger events (e.g., ran the microwave, flushed the toilet). For transient events, such as a door closing, we repeated the action several times over a ten second period. The order we visited locations on any given day was randomized, and then within location, room order was randomized, and then finally within rooms, event performance order was randomized.

All data was captured during regular hours, with occupants going about their normal routines. Throughout data collection, we did not control events happening in other rooms (e.g., we collected microwave events in the kitchen, while someone could be washing a cup, flushing a toilet, or printing a document). It was common for captured data to have background noise, including e.g., human chatter, road noise from outside, HVAC and occupants walking around. This noise is purposely part of our data, and one of the reasons why we ran a deployment as opposed to an entirely controlled lab study (like that in [24], which offers a different and complementary experimental result). We also collected a “background” class (i.e., ambient noise, no event happening) for each room. To facilitate the tedious task of labeling events, we built an annotation tool that supported playback and the ability to overlay synchronized sensor data. Data labeling was completed immediately after performing (known-order) events in a room. In total, this procedure yielded ~640K labeled instances, representing ~17.8 hours of data.

6 OPEN DATASET AND CODE

We make our full dataset and study scripts freely available at github.com/figlab/constellations. The dataset contains roughly 640K labeled data instances, segmented by event, day, room and location. Each data instance contains 1207 denatured features, reported by all ten sensor boards deployed at a location. These are synchronized to within 50ms. Our hope is for researchers and practitioners to study, replicate, and train more



Fig. 4. Far left, hypothetical layout with eight rooms, with signal of interest in Room 4. To the right, pictorial representations of our five spatial conditions (purple rooms = sensor contributing to recognition).

complex machine learning models leveraging sparse sensor data, and to foster research in general-purpose sensor deployments.

7 RESULTS

To investigate our four research questions, we developed five spatial configuration conditions that serve as probes to quantify system performance: (1) in room sensing, (2) in room + nearest room sensing, (3) all room sensing, (4) only-nearest-room sensing, and (5) all-but-in-room sensing (illustrated in Figure 4). Using our collected data, we built machine learning models for each room (30×30) × each spatial configuration condition (5).

7.1 Analysis Methodology

To evaluate classification performance, we ran day-fold, cross-validation analyses for each model (see Section 4 for details on machine learning). In this day-fold validation scheme, we procedurally choose one day out of our week-long deployment as a test dataset (e.g., Day 3), and use the remaining days (e.g., Days 1-2 and 4-7) for training data. Separating the data in this fashion prevents temporal overfitting. We repeat this process seven times, for all days in the week, and then average the results. This method enables us to quantify the effect of several constellation configurations across variations in location and time. Finally, we note that although we do not specifically discuss “false positives”, our per-room models incorporate a “background” event that serves as a “negative” class for mitigating false event triggers.

7.1 In-Room Sensing Accuracy

In this condition, we evaluated classification accuracy using only data collected from the sensor operating *in the room* where an activity occurred (Figure 4). Due to physical proximity, in-room sensors are the most likely to provide robust classification, and thus serve as a baseline (Q1). Across all 100 sensed events (in 30 rooms in three locations), our event classifiers achieved a mean accuracy of 94.6% (SD 7.1%). In-room accuracy was 95.9% (SD 5.4%) at the Home location, 95.8% (SD 5.7%) at the Institution location, and 91.9% (SD 9.2%) at the Business location. 80 out of the 100 sensed events achieved accuracies greater than 90%, while 33 were 100% accurate. Table 1, “In-Room Only” column, lists the individual accuracies for all 100 sensed events. These baseline accuracies are consistent with results depicted in prior work [24], indicating that indeed, physical events in a room can be reasonably sensed from the vantage point of a single general-purpose sensor.

7.2 In-Room + Nearest Room Sensing

Next, we wished to investigate whether other sensors can offer complementary and confirmatory data to improve classification accuracies beyond that of what in-room sensors can provide alone (Q2). To do this, we trained, and then tested, our models using data from sensors *in the room* plus the *nearest room* (Figure 4). For example, to detect events in a bedroom, we utilize the bedroom sensor as well as the sensor in an adjoining

bathroom. Across all events and rooms, we found an average classification accuracy of 96.4% (SD 5.3%). This is a significant gain of 1.8% over using only in-room sensor data. This yields a statistically significant result ($p<0.001$) using a paired t-test ($n=100$) against the baseline condition (*i.e.*, the in-room configuration). Broken down by location, the Home achieved 96.8% accuracy (a gain of 0.9%), Institution had 96.9% accuracy (+1.1%), and Business had 95.5% (+3.6%). Table 1, “In-Room + Nearest” column, provides full results.

Although these improvements in accuracy may seem small – single percent increases – it is important to note they constitute a significant decrease in error (roughly a third). Closing the gap between 95% and 99% accuracy is often a substantial challenge. However, it also can make the difference between a technology that sits in the literature vs. meaningfully deployed in the real world.

7.3 All Room Sensing

Next, we took our second research question (Q2) to its full conclusion to observe the effect on sensing accuracy if *all* sensors in a building are utilized to detect events (Q3; Figure 4). For example, to detect events in a kitchen, we use sensor data from every room, including the kitchen. In this configuration, the average accuracy across all 100 events was 95.7% (SD 6.2%), a 1.1% increase from in-room sensing.

This too is a significant improvement over in room sensing (paired t-test; $p<.01$). The improvement was consistent across our test locations: Home was 95.8% (+0.1%), Institution was 96.9% (+1.1%), and Business was 94.1% (+2.2%).

Though better than in-room accuracy, it was significantly worse than in room + nearest room sensing (paired t-test; $p<.05$). We suspect this dampening of performance is chiefly due to the substantial increase in machine learning features (from ten sensors, vs. just one or two), most of which likely contain no information power, introduce unwanted noise, and may lead to overfitting of some classifiers. Please also see Table 1, “All Rooms” column, for individual event results.

7.4 Nearest Room Only Sensing

The results discussed thus far suggest improved recognition accuracy when in-room sensors are complimented with other sensors, either nearby or throughout a building. This a positive result for the notion of deploying a constellation of sensors at a sparsity of roughly one per room. However, we also wished to investigate if it was possible to support recognition of events at densities of *less than one per room* (Q4), potentially making deployments lower cost and even more practical. To start, we examined a *nearest room only* spatial configuration (Figure 4), where we include sensor data from the nearest room, but *not* the sensor in the room where the event actually occurred (Q1).

Across all rooms and events, *nearest-room-only* sensing accuracy is 63.4% (STD 31.4%), a 31.2% decrease from our baseline, *in-room* sensing condition, which is significantly worse (paired t-test; $p<.001$). Home location drops to 55.2% accuracy (-40.7%), Institution drops to 57.6% (-38.2%), and Business drops to 79.1% (-12.8%). For the latter location, we suspect a more compact physical footprint helped buffer losses, as the nearest sensor was often still fairly close, offering a greater chance to capture signals of interest. More interestingly, the drop in accuracy is not uniform across all sensed events. Indeed, 21 of our 100 events had unchanged accuracies or even improved in accuracy. If we consider an accuracy drop of less than 10% acceptable, 36 of our sensed events qualify, despite not having a sensor in the room where the event occurred, which is a promising result. Specifically, it suggests that for most events, in room sensing is required to achieve high accuracy, however there are some classes of events that produce physical manifestations that are detectable beyond the immediate room, allowing for remote detection and thus sparser sensor distributions, a topic we revisit in later discussion.

7.5 All But-In-Room

Building on the previous section, we also investigated the feasibility of supporting events when no sensor is present in a room (Q1), but leveraging *all other* sensors available at a location (Q3/Q4). In this *all but in-room* condition (Figure 4), average sensing accuracy degrades to 53.8% (STD 30.1%), a 40.8% drop-off from in room sensing. For the Home location, overall accuracy is 42.5% (-53.4%), 47.8% at the Institution location (-48.0%), and 73.3% for the Business location (-18.6%). This result is worse than just relying on the single nearest sensor (marginal significance, $p=0.11$). Similar to the *all rooms* condition, we suspect this accuracy degradation versus our *nearest-room-only* condition is from overfitting of the models due to excessive features, most of which are from sensors unable to capture any useful signal.

7.6 Learning Curves

We conducted a supplemental analysis to understand how different sensor configurations might impact the training of classifiers. For this, we computed per-day cumulative learning curves, which illustrate accuracy gains as more training data becomes available. More specifically, we use data collected on day 1, and test it against data from days 2-7, simulating classification accuracy as though the model only had one day to learn before being deployed. For day 2, we train on data from days 1 and 2, and test on the remaining data (days 3-7). We repeat this process for all days in our training set, up to day 6, which is tested on day 7. The results (Figure 5) mirror our previous results, with *in room*, *in room + nearest room* and *all room* spatial configurations performing the best. For these three conditions, accuracy appears to have a positive trajectory on day 7, suggesting deployments longer than a week might yield even stronger accuracies. Next best is training on *nearest-room-only* data, followed by *all-but-in-room* data, with both accuracies appearing to plateau by around day 3.

8 DISCUSSION AND IMPLICATIONS

Synthetic Sensors [24] – the project most related to this research – demonstrated 96.0% accuracy across 38 sensed events at a single controlled location. Our larger and considerably more ecologically valid study yielded a similar 94.6% accuracy over 100 sensed events. Moreover, 54 sensed events achieved accuracies $\geq 99\%$,

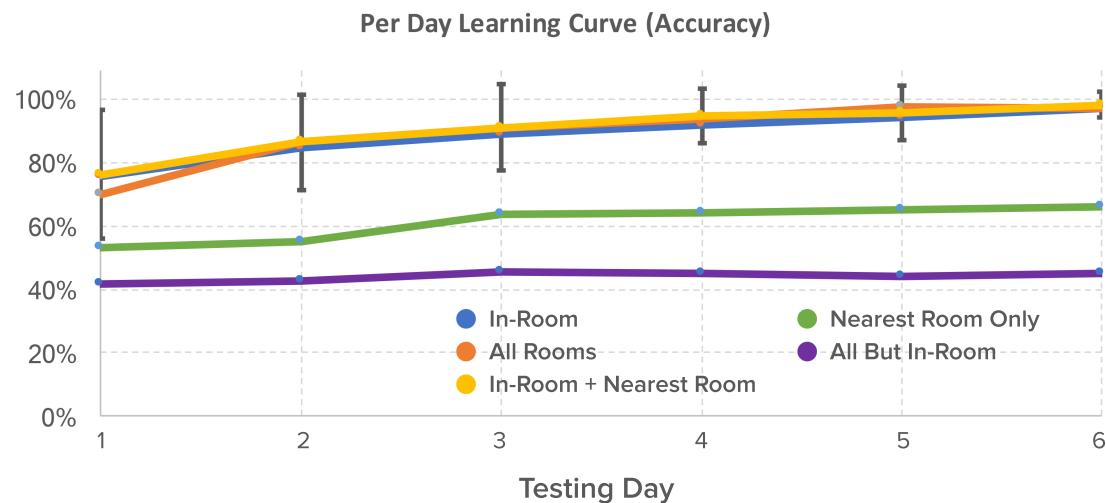


Fig. 5. Accuracy with increasing days of training data for different sensor spatial configurations. Error bars plotted for “in-room + nearest room” (representative), illustrating how accuracy variance decreases over time.

reinforcing the central premise of prior research that room-scale, general-purpose sensing can achieve useful accuracies *without* direct instrumentation of objects.

8.1 Benefit of Constellations & Reinforcement

Moving beyond prior work, we quantified for the first time the benefit of leveraging other general-purpose sensors deployed at a location. In particular, we tested two ways to supplement data from in-room sensors: leveraging data from *one nearby sensor*, and also leveraging data from *all sensors* at a location. In general, we found this generally improves event recognition, by around 1.5% in our study. In some cases, the gain is significant, and perhaps key to unlocking usable accuracies. For example, the *faucet running* event in the Business Bathroom sees *in-room* accuracy of 65.2% jump to 94.6% when adding *nearest-room* data (in this case, due to plumbing inducing vibrations on an adjoining wall). Another example in the Business location is the *espresso machine running* event, which jumps from 69.8% to 95.0% (also due to vibrations induced in a neighboring room). These two examples, and 14 other events, improve by at least 5% in absolute accuracy when allowing classifiers to utilize nearest-room sensor data, demonstrating the value of distributed sensor fusion. Moreover, this benefit is mutual: deploying in adjacent rooms *reinforces* recognition in both rooms.

8.2 More Sensor Data is Not Always Better

One of our initial hypotheses was that more sensor data would be better, and even though some sensors may be far away, they would still offer weak, but useful signals to support recognition. This did not appear to be the case in our deployment, as our results show that *in room + nearest room* sensing was significantly better than leveraging all available data (*i.e.*, *all rooms* condition). As noted previously, we believe this is due to the substantial growth in machine learning features, which add little or no information power, impeding model training. We speculate this could be overcome with more training data, where perhaps weak and noisy signals could ultimately improve classification. However, end users will demand robust classification “out of the box”, and so it may be that deployments have to start with in-room sensing, slowly add nearest-room data, and finally transition to whole house data.

8.3 Benefit of Proximity

Our results also serve to confirm that proximity to physical events is vital to enable robust classification when using a general-purpose sensor (our three best spatial conditions all took advantage of data from an in-room sensor). Importantly, this detriment cannot be overcome simply by having more sensors elsewhere, as illustrated by our *nearest-room-only* and *all-but-in-room* conditions. For example, despite offering classifiers nine times the sensor data, including from the nearest room, *all-but-in-room* sensing is 40.8% worse on average than *in room* sensing (53.8% vs. 94.6%) – accuracies untenable for supporting end-user applications.

8.4 Overcoming Lack of Proximity

Interestingly, the accuracy drop when having to rely on out-of-room sensors is bi-modal, with some event classes essentially operating at chance (*i.e.*, no useful signal to operate on), while roughly a third of sensed events saw little or no degradation in accuracy. We suspected this behavior was correlated with how events manifested and propagated physically. To investigate this, we categorized our event classes into three properties of interest:

- 1) Is the event acoustically *loud*? Loud sounds will travel to adjacent spaces (e.g., a door closing, a large appliance running), where they can be detected by other sensors.
- 2) Did the event cause *vibration*? Though often subtle, vibrations propagate through building structures (e.g., a toilet flushing, laundry machine spinning) with much less degradation than sound through air. Additionally, this can travel between floors, and along walls/floor/ceilings to rooms not immediately connected. With sensors plugged into wall power sockets (which are typically mounted to a wall stud), they are in an ideal position to capture structural vibrations originating elsewhere.
- 3) Does the event produce signals that require *line of sight* (i.e., a direct, unimpeded path between event and sensor). This is a property of some of our sensor's channels, such as non-contact temperature (e.g., stove burner on) and infrared motion sensing (e.g., occupancy).

We applied these property labels to our sensed events in Table 1; Figure 6 offers an accuracy comparison across spatial configurations broken out by these properties. Of the events that experience less than 10% drops in accuracy when no sensor is present in the room they occur, 90% are loud, vibrational or both. This suggests that these classes of physical event may be the use cases where in-room sensor placement can be skipped, and the sensor saved for deployed in areas with more subtle signals. We caution this effect, however, only works well for noisy or strong vibratory signals (e.g., doors closing, HVAC, motor-power appliances), and is not universal. For example, the toaster in the Home Kitchen dropped in accuracy from 97.6% *in room* accuracy to 86.6% *nearest-room-only* accuracy, despite emitting a completion chime. Although loud, and certainly audible in the adjacent room, there is nonetheless a substantial drop in accuracy, especially when there is any noise present in the adjacent room. Likewise, small vibrations, especially those local to a surface, such as typing, do not appear to have enough energy or coupling to structural members to be detected from outside of their room (or indeed even in the room, as seen in the case of *mouse moving on table*, with in-room accuracies of 84.0%).

8.5 Line of Sight

We identified five sensed events that *required* line of sight (through a feature selection analysis that showed a majority of information power was derived from sensor channels requiring line of sight). In all five cases, this was due to a reliance on non-contact thermal sensing to detect heat-related events (e.g., fireplace on). When moving from in-room sensing *with line of sight*, to the nearest-room sensor *with no line of sight*, the average accuracy went from 97.6% to 17.3% (Figure 6), a much sharper drop than the 95 events not requiring line of sight (average *nearest-room-only* accuracy of 65.2%). This reinforces the notion that some events require not just in-room sensing, but line of sight sensing, and when this is impossible (e.g., no available outlet, occlusion from furniture or occupants, or aesthetic reasons), it can render that event impossible to sense, even if leveraging all sensors in a constellation.

8.6 Redundancy

A final benefit of sparse sensor constellations is *redundancy*. In real world environments, sensors will lose connectivity, be occluded by people and furniture, get damaged, and outright fail. Ideally, we want ubiquitous sensing systems to gracefully continue when such failures occur, even if it means operating at reduced accuracies. Our results suggest this is indeed possible, by transferring the responsibility to the next nearest sensor. We can see in our *nearest-room-only* spatial condition (which offers a proxy for this failure event) that 45 of our 100

events operate at $\geq 80\%$ accuracy (22 events at $\geq 95\%$), which may be acceptable in such failure modes, as opposed to a loss of sensing entirely.

8.7 Placement Guide

Although we believe an ideal sensing system should require no special placement consideration (*i.e.*, essentially “plug and play”), there are nonetheless lightweight optimizations that can be considered to increase the chances of a successful deployment. The most immediate rule of thumb is for the sensor to be as close as possible to events of interest, as borne out by our main study results. Next, users should consider the relative signal strength of events, *e.g.*, subtle noises and vibrations are less likely to travel and must contend with background noise. If a user can hear or feel an event, that is a strong indicator of actionable SNR. Likewise, for events requiring line of sight, it may be that sensors have to move to a less desirable location in order to have visibility. This tradeoff is impossible to quantify, as spaces are so varied. With regards to signals that are able to be sensed through walls, such as plumbing and highly vibratory appliances, our results still suggest in-room placement is the strongest avenue. However, if out-of-room deployment is desired, adjacent walls are the next most reliable instrumentation point. Where relevant, we have added placement notes and recommendations to our open source corpus. We acknowledge this guide offers few concrete recommendations, largely because spaces (their layout and construction) are so varied, which precludes all but the most generic recommendations. This complexity might suggest that companion apps are needed for rapid, end user evaluation of sensor placements.

8.8 Performance across Rooms

In our study design, we included a large number of rooms – of varying size and function – so as to provide a more generalizable result that was not tightly coupled to a specific location or event type. Furthermore, as noted earlier, this provided an uncontrolled, but real-world distribution of electrical outlets for sensor deployment. Without doubt, some rooms were more favorable, while others were more challenging. With our study data and accuracy results, we can explore this distribution to better understand underlying effects. For simplicity, we use results from our *in room only* sensing condition.

Foremost, when averaging accuracies of all classifiers operating in a room, we can see that room accuracies vary from 65.2% to 100% (mean 94.6%, SD 7.1%). Of course, rooms vary in size, function, sensor placement and

Accuracy Characteristics for Different Event Types

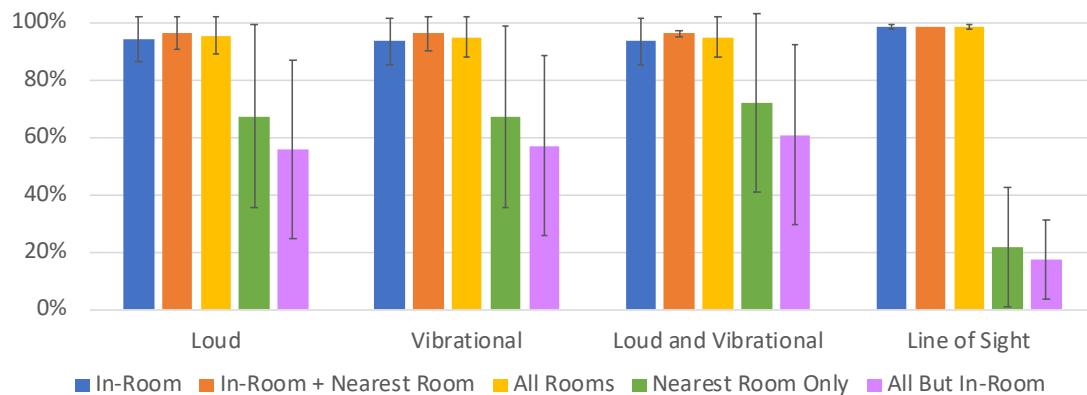


Fig. 6. Accuracy across different event properties and sensor spatial configurations.

challenge of the sensed facets, and so a more reliable test of generalizability across environments is to look at sensed events that were common across rooms. For example, we ran *door closed* classifiers in 17 rooms, which varied in accuracy from 83.2% to 100.0% (mean of 94.3%, SD 6.1%) – an extremely tight clustering which suggests this signal is robust across locations and sensor placements.

Conversely, *faucet running*, our next most common classifier (included in 8 rooms) varied in accuracy from 65.2% to 100.0% (mean of 88.8%, SD 10.2%) – a wider distribution of accuracies. This is almost certainly due to a much weaker signal (vibro-acoustically) than a door closing and latching shut. In locations where the sensor happened to be placed far from the faucet in the room, and especially when decoupled from the wall containing the plumbing stack, accuracy dropped substantially (94.6% when proximate, 80.1% when far). Thus, sensor placement does appear to be important factor in accuracy for some signals, much more so than the room they operate in, and it suggests that users should be encouraged to place sensors near areas of interest when outlets permit.

9 STUDY LIMITATIONS AND FUTURE WORK

Although we conducted a large deployment in 30 rooms across three buildings, this is still a small subset of the diverse and complex environments present in the real world. Likewise, our 67 unique sensed events do not represent all activities that users and smart buildings may wish to monitor. Although we believe our deployment study is a significant contribution to this long-standing problem domain, tackling key questions, many challenges remain before such sensing can move out into the real world.

For example, many physical events can be subtle or imperceptible to the types of sensors we deployed (e.g., is the cat sleeping, is the cooler out of water). Indeed, most of our sensed events relied on sound and vibration, but none used change in magnetic field, temperature, or humidity (despite being captured by the sensor board we used). This was partly due to the fact that the events we studied were on the milliseconds-to-seconds time-scale, whereas interesting variations in signals like temperature and humidity typically happen on the order of hours or days. Such longer-term sensing might be able to detect events such as the ingress of moisture or poor insulation, which e.g., home owners might wish to be notified about.

Additionally, events in the real world are chaotic, noisy, and can happen simultaneously, whereas the events captured in our study were mostly isolated. Although there were uncontrolled background activities happening during our study, we did not explicitly perform multiple simultaneous events, which was combinatorically explosive and logically challenging. Anecdotally, however, we can report that simultaneous recognition of events occurring in different rooms is straightforward when using any of the spatial configurations utilizing in-room sensors (where the in-room sensor dominates classification). However, recognizing multiple simultaneous events in the same room remains a significant challenge. We have examined signal processing techniques e.g., adaptive background subtraction [16] and source separation [28] as possible directions for future work.

Finally, we note that several events in our deployment were common across rooms and locations, such as detecting doors closing and faucets running. Many of these exhibited similar sound and vibrational properties regardless of location. For instance, a door closing is typified by a loud sound and vibrational impulse, whereas a faucet produces a steady white-noise-esque sound with low frequency vibrations. We believe these types of classes could be learned by a “universal” model that would require no on-site training, a challenge we hope to address in future work.

10 CONCLUSION

Many previous distributed sensing systems have envisioned densely instrumented environments. However, through clever sensing, we believe many of the same benefits can be achieved in a more practical and less obtrusive manner through sparse arrays of general-purpose sensors. In this work, we explored the efficacy of

this approach, wherein a constellation of sensors work synergistically to support sensing of physical events across multi-room environments. We found that one sensor per room can unlock high accuracies across a wide variety of sensed events, supporting and extending prior study results. More interestingly, by leveraging sensors in other rooms, accuracy can be further improved. We also found that for some classes of events, detection is possible even when no sensor is present in the room, instead relying entirely on other sensors in the building. Taken together, these results suggest that near-complete sensor saturation of everyday spaces is within reach using sparse sensor arrays.

ACKNOWLEDGMENTS

This research was generously funded by the The David and Lucile Packard Foundation and a Google Ph.D. Fellowship. We also thank Sudershan Boovaraghavan, Chen Chen, and Yuvraj Agarwal for their system improvements, as well as Qeexo for offering their office space for in-the-wild experiments. Finally, we acknowledge our reviewers and associate editors for their invaluable feedback.

REFERENCES

- [1] R.E Abbott and S.C. Hadden. 1990. Product Specification for a Nonintrusive Appliance Load Monitoring System. *EPRI Report #NI-101*, 1990.
- [2] Artem Babenko and Victor Lempitsky. 2015. Aggregating Deep Convolutional Features for Image Retrieval. In *IEEE International Conference on Computer Vision (ICCV)*, Dec. 2015, pp. 1269-1277. DOI: 10.1109/ICCV.2015.150. arXiv: 1510.07493.
- [3] Rajesh Krishna Balan, Archan Misra, and Youngki Lee. 2014. LiveLabs: building an in-situ real-time mobile experimentation testbed. In *Proceedings of the 15th Workshop on Mobile Computing Systems and Applications (HotMobile '14)*. ACM, New York, NY, USA, Article 14, 6 pages. DOI=http://dx.doi.org/10.1145/2565585.2565597
- [4] Alex Beltran, Varick L. Erickson, and Alberto E. Cerpa. 2013. ThermoSense: Occupancy Thermal Based Sensing for HVAC Control. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings (BuildSys'13)*. ACM, New York, NY, USA. Article 11, 8 pages. DOI=http://dx.doi.org/10.1145/2528282.2528301
- [5] Jeffrey P. Bigham, Chandrika Jayant, HaSmoMwenjie Ji, Greg Little, Andrew MillerCoffe Maker Micro, Robert C. Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samual White, and Tom Yeh. 2010. VizWiz: nearly real-time answers to visual questions. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology (UIST '10)*. ACM, New York, NY, USA, 333-342. DOI=http://dx.doi.org/10.1145/1866029.1866080
- [6] Bosch. Cross Domain Development Kit. Last accessed: September 16, 2018. <https://xdk.bosch-connectivity.com>
- [7] Michael Boyle, Christopher Edwards, and Saul Greenberg. 2000. The effects of filtered video on awareness and privacy. In *Proceedings of the 2000 ACM conference on Computer supported cooperative work (CSCW '00)*. ACM, New York, NY, USA, 1-10. DOI=http://dx.doi.org/10.1145/358916.358935
- [8] Leo Breiman. 2001. Random Forests. 2001. In *the Journal of Machine Learning*. Volume 45-1. 5-32. DOI: <https://doi.org/10.1023/A:1010933404324>
- [9] Gabe Cohn, Sidhant Gupta, Jon Froehlich, Eric Larson, and Shwetak N. Patel. 2010. GasSense: Appliance-Level, Single-Point Sensing of Gas Activity in the Home. In *Pervasive Computing: 8th International Conference (Pervasive 2010)*, Helsinki, Finland, May 17-20, 2010. Springer. 265-282. DOI=http://dx.doi.org/10.1007/978-3-642-12654-3_16
- [10] Diane J. Cook, Narayanan C. Krishnan, and Parisa Rashidi. 2013. Activity Discovery and Activity Recognition: A New Partnership. *IEEE Transactions on Cybernetics*, Volume: 43, Issue: 3, June 2013.
- [11] David Culler, Jason Hill, Mike Horton, Kris Pister, Robert Szewczyk, and Alec Wood. 2002. Mica: The commercialization of microsensor motes. In *Sensor Technology and Design*.

- [12] Samuel DeBruin, Bradford Campbell, and Prabal Dutta. 2013. Monjolo: an energy-harvesting energy meter architecture. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems* (SenSys '13). ACM, New York, NY, USA, Article 18, 14 pages. DOI=<http://dx.doi.org/10.1145/2517351.2517363>
- [13] James Fogarty, Carolyn Au, and Scott E. Hudson. 2006. Sensing from the basement: a feasibility study of unobtrusive and low-cost home activity recognition. In *Proceedings of the 19th annual ACM symposium on User interface software and technology* (UIST '06). ACM, New York, NY, USA, 91-100. DOI=<http://dx.doi.org/10.1145/1166253.1166269>
- [14] Jon E. Froehlich, Eric Larson, Tim Campbell, Conor Haggerty, James Fogarty, and Shwetak N. Patel. 2009. HydroSense: infrastructure-mediated single-point sensing of whole-home water activity. In *Proceedings of the 11th international conference on Ubiquitous computing* (UbiComp '09). ACM, New York, NY, USA, 235-244. DOI=<http://dx.doi.org/10.1145/1620545.1620581>
- [15] Anhong Guo, Xiang 'Anthony' Chen, Haoran Qi, Samuel White, Suman Ghosh, Chieko Asakawa, and Jeffrey P. Bigham. 2016. VizLens: A Robust and Interactive Screen Reader for Interfaces in the Real World. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (UIST '16). ACM, New York, NY, USA, 651-664. DOI: <https://doi.org/10.1145/2984511.2984518>
- [16] Sidhant Gupta, Matthew S. Reynolds, and Shwetak N. Patel. 2010. ElectriSense: single-point sensing using EMI for electrical event detection and classification in the home. In *Proceedings of the 12th ACM international conference on Ubiquitous computing* (UbiComp '10). ACM, New York, NY, USA, 139-148. DOI=<http://dx.doi.org/10.1145/1864349.1864375>
- [17] Sidhant Gupta, Ke-Yu Chen, Matthew S. Reynolds, and Shwetak N. Patel. 2011. LightWave: using compact fluorescent lights as sensors. In *Proceedings of the 13th international conference on Ubiquitous computing* (UbiComp '11). ACM, New York, NY, USA, 65-74. DOI=<http://dx.doi.org/10.1145/2030112.2030122>
- [18] George Hart. Advances in Nonintrusive Appliance Load Monitoring. In *Proc. of EPRI Inf. and Automation Conf.* '91.
- [19] George Hart. Nonintrusive Appliance Load Monitoring. 1992. In *Proceedings of the IEEE EPRI Information and Automation Technology Conference*, Washington, DC, June 26-28, 1992. 80(12), 1870-1891. DOI=<http://dx.doi.org/10.1109/MIM.2016.7477956>
- [20] Scott E. Hudson and Ian Smith. 1996. Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems. In *Proceedings of the 1996 ACM conference on Computer supported cooperative work* (CSCW '96). ACM, New York, NY, USA, 248-257. DOI=<http://dx.doi.org/10.1145/240080.240295>
- [21] International Code Council, Inc. (2011). 2012 International Building Code (IBC). Country Club Hills, IL.
- [22] Nicholas Lane, Emiliano Miluzzo, Hong Lu, Daniel Peebles, Tanzeem Choudhury, and Andrew T. Campbell. A survey of mobile phone sensing. In *IEEE Communications Magazine*, vol. 48, no. 9, pp. 140-150, Sept. 2010. DOI: <https://doi.org/10.1109/MCOM.2010.5560598>
- [23] Gierad Laput, Walter S. Lasecki, Jason Wiese, Robert Xiao, Jeffrey P. Bigham, and Chris Harrison. 2015. Zensors: Adaptive, Rapidly Deployable, Human-Intelligent Sensor Feeds. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15). ACM, New York, NY, USA, 1935-1944. DOI: <http://dx.doi.org/10.1145/2702123.2702416>
- [24] Gierad Laput, Yang Zhang, and Chris Harrison. 2017. Synthetic Sensors: Towards General-Purpose Sensing. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (CHI '17). ACM, New York, NY, USA, 3986-3999. DOI: <https://doi.org/10.1145/3025453.3025773>
- [25] Walter S. Lasecki, Young Chol Song, Henry Kautz, and Jeffrey P. Bigham. 2013. Real-time crowd labeling for deployable activity recognition. In *Proceedings of the 2013 conference on Computer supported cooperative work* (CSCW '13). ACM, New York, NY, USA, 1203-1212. DOI=<http://dx.doi.org/10.1145/2441776.2441912>
- [26] Beth Logan, Jennifer Healey, Matthai Philipose, Emmanuel Munguia Tapia, and Stephen Intille. 2007. A long-term evaluation of sensing modalities for activity recognition. In *Proceedings of the 9th international conference on Ubiquitous computing* (UbiComp '07). Springer-Verlag, Berlin, Heidelberg, 483-500.

- [27] MeSchup - Build connected products. Last accessed: September 17, 2018. <http://www.meschup.com>
- [28] Ganesh R. Naik and Wenwu Wang. 2014. Blind Source Separation: Advances in Theory, Algorithms and Applications. In *Signals and Communication Technology*, Springer 2014. ISBN 9783642550164.
- [29] Notion - Wireless Home Monitoring Made Easy. Last accessed: September 15, 2018. <https://getnotion.com>
- [30] Shijia Pan, Mostafa Mirshekari, Pei Zhang, Hae Young Noh, "Occupant traffic estimation through structural vibration sensing," In *Proc. SPIE 9803, Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems 2016*, 980306 (20 April, 2016). DOI= <http://dx.doi.org/10.1117/12.2222024>
- [31] Particle Inc, Spark Server. Last accessed: September 15, 2018. <https://github.com/spark/spark-server>
- [32] Shwetak N. Patel, Thomas Robertson, Julie A. Kientz, Matthew S. Reynolds, and Gregory D. Abowd. 2007. At the flick of a switch: detecting and classifying unique electrical events on the residential power line. In *Proceedings of the 9th international conference on Ubiquitous computing (UbiComp '07)*. Springer-Verlag, Berlin, Heidelberg, 271-288.
- [33] Shwetak N. Patel, Matthew S. Reynolds, and Gregory D. Abowd. 2009. Detecting Human Movement by Differential Air Pressure Sensing in HVAC System Ductwork: An Exploration in Infrastructure Mediated Sensing. In *Proceedings of Pervasive '08*. Springer-Verlag, Berlin, Heidelberg, 1-18. DOI=http://dx.doi.org/10.1007/978-3-540-79576-6_1
- [34] Fabian Pedregosa et. al. 2011. Scikit-learn: Machine Learning in Python. In *the Journal of Machine Learning Research*. Vol. 12. 2825-2830, 2011.
- [35] Kristofer SJ Pister, Joseph M. Kahn, and Bernhard E. Boser. 1999. Smart dust: Wireless networks of millimeter-scale sensor nodes. Highlight Article, *Electronics Research Laboratory Research Summary*, 1999.
- [36] Manuel Ruiz-Sandoval, Tomonori Nagayama, and B. F. Spencer Jr. 2006. Sensor development using Berkeley Mote platform. In *Journal of Earthquake Engineering* 10, no. 2, 2006. 289-309. DOI: <http://dx.doi.org/10.1080/13632460609350597>
- [37] Sen.se Mother. The Universal Monitoring Solution. Last accessed: September 15, 2018. <https://sen.se/mother/>
- [38] Nagender Kumar Suryadevara and Subhas Chandra Mukhopadhyay. Wireless Sensor Network Based Home Monitoring System for Wellness Determination of Elderly. In *IEEE Sensors Journal*, vol. 12, no. 6, pp. 1965-1972, June 2012. DOI: <https://doi.org/10.1109/JSEN.2011.2182341>
- [39] Emmanuel Munguia Tapia, Stephen S. Intille, and Kent Larson. 2004. Activity Recognition in the Home Using Simple and Ubiquitous Sensors. In *Pervasive Computing: Second International Conference (PERVASIVE 2004)*, Linz/Vienna, Austria, April 21-23, 2004. Springer. 158-175. DOI=http://dx.doi.org/10.1007/978-3-540-24646-6_10
- [40] Texas Instruments. SimpleLink SensorTag. Last accessed: September 14, 2018. http://www.ti.com/ww/en/wireless_connectivity/sensortag2015/gettingStarted.html
- [41] Geoffrey Werner-Allen, Konrad Lorincz, Matt Welsh, Omar Marcillo, Jeff Johnson, Mario Ruiz, and Jonathan Lees. 2006. Deploying a Wireless Sensor Network on an Active Volcano. In *IEEE Internet Computing* 10, 2 (March 2006). 18-25. DOI=<http://dx.doi.org/10.1109/MIC.2006.26>
- [42] Mengyu Zhou, Kaixin Sui, Minghua Ma, Youjian Zhao, Dan Pei, and Thomas Moscibroda. 2016. MobiCamp: a Campus-wide Testbed for Studying Mobile Physical Activities. In *Proceedings of the 3rd International on Workshop on Physical Analytics (WPA '16)*. ACM, New York, NY, USA, 1-6. DOI: <http://dx.doi.org/10.1145/2935651.2935654>

Received November 2018; revised February 2019; accepted April 2019