

Indoor Smartphone SLAM with Learned Echoic Location Features

Wenjie Luo*
Singtel Cognitive and AI Lab for
Enterprises
Nanyang Technological University
Singapore

Qun Song*
ERI@N, Interdisciplinary Graduate
School
Nanyang Technological University
Singapore

Zhenyu Yan
Department of Information
Engineering
The Chinese University of Hong Kong
Hong Kong SAR, China

Rui Tan*
Singtel Cognitive and AI Lab for
Enterprises
Nanyang Technological University
Singapore

Guosheng Lin
School of Computer Science and
Engineering
Nanyang Technological University
Singapore

Abstract

Indoor self-localization is a highly demanded system function for smartphones. The current solutions based on inertial, radio frequency, and geomagnetic sensing may have degraded performance when their limiting factors take effect. In this paper, we present a new indoor simultaneous localization and mapping (SLAM) system that utilizes the smartphone's built-in audio hardware and inertial measurement unit (IMU). Our system uses a smartphone's loudspeaker to emit near-inaudible chirps and then the microphone to record the acoustic echoes from the indoor environment. Our profiling measurements show that the echoes carry location information with sub-meter granularity. To enable SLAM, we apply contrastive learning to construct an echoic location feature (ELF) extractor, such that the loop closures on the smartphone's trajectory can be accurately detected from the associated ELF trace. The detection results effectively regulate the IMU-based trajectory reconstruction. Extensive experiments show that our ELF-based SLAM achieves median localization errors of 0.1 m, 0.53 m, and 0.4 m on the reconstructed trajectories in a living room, an office, and a shopping mall, and outperforms the Wi-Fi and geomagnetic SLAM systems.

CCS Concepts

• **Human-centered computing** → Ubiquitous and mobile computing systems and tools; Smartphones; • **Computing methodologies** → Neural networks.

Keywords

Simultaneous localization and mapping, acoustic sensing

*Also with School of Computer Science and Engineering, Nanyang Technological University, Singapore.

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.

SenSys '22, November 6–9, 2022, Boston, MA, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9886-2/22/11...\$15.00

<https://doi.org/10.1145/3560905.3568510>

ACM Reference Format:

Wenjie Luo, Qun Song, Zhenyu Yan, Rui Tan, and Guosheng Lin. 2022. Indoor Smartphone SLAM with Learned Echoic Location Features. In *ACM Conference on Embedded Networked Sensor Systems (SenSys '22)*, November 6–9, 2022, Boston, MA, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3560905.3568510>

1 Introduction

Location sensing is a fundamental service of mobile operating systems. As of May 2022, 73% of the top 100 free Android apps on Google Play [2] require location information. In the last two decades, research has exploited various smartphone's built-in sensing modalities for indoor location sensing, which include Wi-Fi [10, 57], GSM [29], FM radio [16], visible light [59], imaging [23], acoustic background [51], and geomagnetism [27]. However, these sensing modalities have their own limiting factors. For instance, radio frequency (RF) signals are susceptible to electromagnetic noises and transceivers' automatic gain controls. Visible light sensing suffers blockage. Visual imaging may generate privacy concerns in certain spaces and times. The acoustic background only gives room-level granularity. Hence, identifying new modalities from smartphones' common built-in hardware to enrich the location sensing toolbox has been an interest of research.

Exploiting smartphone's built-in audio hardware for active location sensing receives increasing research [33, 43, 45, 48, 52, 62]. The active sensing uses smartphone's loudspeaker to emit excitation sounds in the target indoor space and microphone to capture the echoes from the surroundings. The echoes carry information that depicts how a sound emitted from a specific location reverberates in the indoor space. The existing approaches can be classified into two categories. The first category, namely, *analytic approach* [33, 43, 62], estimates the smartphone's location by analyzing the processes of the sound reflections by nearby surfaces (e.g., walls). The image source modeling (ISM) is a prevailing analysis technique. However, when the surroundings are complex (e.g., irregular surfaces, many nearby objects with complex 3D structures, etc), accurate ISM becomes intractable. Thus, the existing studies following the analytic approach often make simplifying assumptions that the major reflectors are at most two nearby walls [33, 43, 62]. The second category, namely, *fingerprint approach* [45, 48, 52], uses the echoes captured by the smartphone at a certain location as the fingerprint of the location and then applies supervised learning to build a location

recognition model. However, the blanket process of collecting labeled fingerprints at spatially fine-grained locations to form the training dataset imposes a high overhead. Thus, the existing studies only target room-level localization [45, 48] or recognize a limited number of locations (11 closed locations in [52]).

Nevertheless, the fingerprint approach has the potential in offering good generalizability because it does not make specific assumptions about the surroundings. Moreover, the high spatial resolution achieved in [52] encourages further investigation of whether satisfactory resolutions can be maintained when the number of fingerprinted locations increases. To develop a preliminary understanding, in this paper, we conduct a blanket process of fingerprinting 128 locations using excitation chirps with near-inaudible frequencies in a $16 \times 28 \text{ m}^2$ space hosting tens of cubicles and several meeting areas. Our analysis by applying supervised learning to the collected data shows that the fingerprint approach can maintain sub-meter localization accuracy. The results suggest that the acoustic echo is a promising modality for designing useful indoor localization services for off-the-shelf smartphones.

To unleash the fingerprint approach from the blanket labeled training data collection process, in this paper, we aim to design a simultaneous localization and mapping (SLAM) system based on the smartphone’s inertial measurement unit (IMU) data and the acoustic echoes. Specifically, when the user carrying the smartphone moves in the space, both IMU data and acoustic echo data are collected. When the user returns to a location visited earlier, the movement trajectory produces a *loop closure*. If the loop closures can be accurately detected using the echo data, the IMU-based dead reckoning result, which is deviation-prone, can be rectified to yield a more accurate trajectory reconstruction. Accordingly, the reconstructed trajectory and the associated echo fingerprints form a *trajectory map*. The trajectory map only contains the echoes from the space covered by the trajectory. The extension to the whole indoor space relies on the *map superimposition* of many users’ trajectory maps.

As a key to SLAM, loop closure detection requires an effective embedding such that any two acoustic echo samples captured at the same location are close in the embedded space, while those collected at different locations are apart. However, finding an effective embedding by which a certain similarity metric can effectively signal loop closures is challenging. Our tests show that the generic acoustic features, e.g., power spectral density, spectrogram, and principal component analysis, cannot effectively discriminate locations. Thus, we resort to learning an effective embedding. A straw-man proposal is to train a deep neural network (DNN) using labeled echo data collected at sparse locations and use the output of a hidden layer as the embedding. However, such embedding for recognizing the pre-defined locations in general becomes ineffective for the locations out of the training dataset such as the arbitrary locations on the user’s trajectory in SLAM. Contrastive learning (CL), on the other hand, exploits the self-supervised learning technique to learn effective representations from unlabeled data. Applied to our SLAM problem, it can learn from the unlabeled echo data collected on the user’s trajectory and only requires the information of whether two echo samples are collected at close locations. Such coarse proximity information can be easily derived using heuristics. Thus, we apply CL with a cosine similarity-based contrastive loss function to train an embedding DNN that outputs a representation called

echoic location feature (ELF). As such, the cosine similarity between any two echo samples’ ELFs effectively signals whether the two samples are collected at the same/close locations (i.e., loop closure).

To realize the ELF-based SLAM, we make the following three designs. First, we design a trajectory-level CL procedure to learn the trajectory-specific ELFs for loop closure detection. It consists of the *pre-training* step, which uses extensive synthetic echoes to build a basic ELF extractor, and the *fine-tuning* step, which adapts the basic extractor to a target indoor space using unlabeled echoes collected on the user trajectory. Second, we design a clustering-based approach to curate the loop closures. This is because, although the similarity values among the ELFs can signal most of the true positive loop closures, they can also cause false positives. The curation exploits prior knowledge about the user’s movement to effectively remove the false positives. Third, we design a floor-level CL procedure to superimpose the trajectory maps from many users to form a single *floor map*. The procedure reconciles the differences among the ELFs from different trajectory maps at the same location due to the non-negligible impact of smartphone orientation on echo data.

The contributions of this paper are summarized as follows:

- We conduct extensive profiling measurements to investigate the acoustic echo’s spatial distinctness under a supervised learning setting. Our results suggest a sub-meter spatial resolution limit and show that acoustic echo is promising for indoor location sensing with sub-meter accuracy.
- We design ELF with CL to detect loop closures for SLAM. We further design ELF-SLAM, a smartphone SLAM system using IMU and ELF data. We also apply CL to reconcile and superimpose the maps from many users.
- We conduct extensive experiments in a living room, an office, and a shopping mall. ELF-SLAM achieves sub-meter mapping and localization accuracy and outperforms the Wi-Fi and geomagnetic SLAMs. We also study the allowable intensities and/or needed mitigation for various practical affecting factors, including nearby people, audible noises, and space layout changes.

The rest of this paper is organized as follows. §2 reviews related work. §3 presents preliminaries. §4 presents the measurement study. §5 presents the design of ELF-SLAM. §6 presents the evaluation results. §7 discusses several issues. §8 concludes this paper.

2 Related Work

■ **Acoustics-based indoor localization and mapping:** The ubiquity of audio speakers and microphones on consumer electronics has attracted research interest in acoustics-based indoor localization. The infrastructure-based approaches deploy dedicated sound beacons or receivers to localize a mobile receiver or beacon. However, the overhead of deploying the infrastructure is undesirable. Our review focuses on infrastructure-free approaches, which are summarized in Table 1. The *analytic approach* analyzes the acoustic signal propagation processes. It passively senses the sounds from the source or actively emits acoustic signals and analyzes the echoes. VoLoc [47] uses a smart speaker to localize a user based on the angle-of-arrival (AoA) of the user’s voice and the wall reflection. In EchoSpot [33], a device emits near-inaudible acoustic signals and analyzes the times of flight of the signals reflected off the human

Table 1: Infrastructure-free acoustic indoor localization and mapping.

Approach	Study	Objective	Presumption on reflectors	Labeled training data	Resolution	Sensing technique		
						Mode	Sensing signal	Duration
Analytic	VoLoc [47]	Localization	Yes	No	Decimeters	Passive	Human voice	15 cmds
	EchoSpot [33]						18-23kHz FMCW ¹	0.2 s
	BatMapper [62]	Mapping				Active	8-16kHz chirps	0.04 s
	SAMS [43]						11-21kHz FMCW	0.03 s
Fingerprint	SurroundSense [9]	Localization	No	Yes	Semantic locations	Passive	Acoustic background	60 s
	Batphone [51]				(Sub-)room-level		0-24kHz MLS ²	0.68 s
	RoomSense [45]					Centimeter	Active	20kHz tone
	DeepRoom [48]	Location tagging		11-22kHz chirp	0.42 s			
	EchoTag [52]			Localization & mapping	No			Sub-meter
	ELF-SLAM							

¹ Frequency-modulated continuous-wave ² Maximum Length Sequence

body and a nearby wall for user spotting. These studies [33, 47] require presumption on the sound reflectors for triangulation.

The analytic approach can be also applied for indoor mapping by estimating the distances of the smartphone to the nearby surfaces (e.g., sidewalls, ceilings, and floors) [43, 62]. To build the wall contour map, the studies [43, 62] require the user to walk straight along the walls. They detect up to two walls simultaneously. In this work, we aim to build the maps on arbitrary trajectories with loop closures and the complete floor maps in indoor spaces. In addition, we aim to be able to address complex spaces such as a lab with dense cubicles and a shopping mall with dense stores. The work [30] considers collated omni-directional speaker and microphone mounted on a robot and analyzes echoes' arrival times for indoor mapping. The evaluation in [30] is only based on simulations. In real world, complex surroundings (e.g., many nearby objects) may bring challenges to the echo arrival identification.

Vis-à-vis analytic approach, the *fingerprint approach* collects training samples from different indoor locations and trains a machine learning model for online location inference [9, 45, 48, 51, 52]. Early studies [9, 45, 51] apply shallow learning and require either long recording times that may cause privacy concerns or full-spectrum recording that is susceptible to interference like human voice [48]. The work [48] applies deep learning to reduce the requirements on recording time and spectrum usage. These studies [9, 45, 48, 51] address semantic or room-level localization. The work [52] uses echoic fingerprints to tag up to 11 locations with centimeters resolution. When the fingerprint approach is applied for the whole indoor space, the blanket process of collecting labeled fingerprints at many locations causes high overhead.

Acoustic sensing has also been used for other applications like inter-smartphone ranging [42, 61], finger [40, 58] and body [41] tracking, gesture recognition [26], speaker authentication [13], breathing rate estimation [55], and lung function monitoring [49].

■ **SLAM:** SLAM constructs the map of an environment in terms of a certain signal and localizes the user device simultaneously. Any SLAM solution consists of two components: *front-end signal processing* and *back-end pose-graph optimization*. Here, we discuss existing SLAM solutions according to their used sensing modalities and the used loop closure detection methods. Radar SLAM [36] and Lidar SLAM [18] are based on point clouds generated by

radar and high-profile lidar, which are unavailable on smartphones. Visual SLAM [50] uses the camera to capture images for landmark detection and map construction. The imaging may introduce privacy concerns. The imaging may introduce privacy concerns. Wi-Fi SLAM [8, 21] employs the received signal strength indicators (RSSIs) from nearby Wi-Fi access points. However, Wi-Fi RSSI can be time-varying. Geomagnetic SLAM [54] exploits the spatially varying magnetic field. Electromagnetic radiation (EMR) SLAM [34] can use the smartphone's earphone as a side-channel sensor to sense the EMR from the alternating current power network. However, the side-channel sensing may experience weak signal strength when the earphone is away from the powerlines. This paper's evaluation will employ geomagnetic, EMR, and Wi-Fi SLAMs as the main baselines. Acoustic SLAM [20] constructs a map based on the AoAs from multiple infrastructural sound sources. Our ELF-SLAM is a new SLAM solution based on infrastructure-free acoustic echoes.

Different sensing modalities need specific loop closure detection methods. The DNN-based loop closure detection on Lidar and Visual SLAM [12, 60] often outperforms the hand-crafted solutions. Wi-Fi SLAM [21] applies the Gaussian process latent variable model to determine the locations of the Wi-Fi RSSI signal. Both the geomagnetic and the EMR SLAM systems employ dynamic time warping (DTW) for loop closure detection. Different from these sensing modalities, generic acoustic features are ineffective to determine locations. Thus, we apply CL to learn a new acoustic feature embedding called ELF to effectively signal loop closures.

3 Preliminaries

■ **Excitation signal:** The commonly used excitation signals include single tone, frequency hopping spectrum spread (FHSS), and chirp [11]. Single tone facilitates detecting Doppler shift and is usually adopted in Doppler ranging and tracking [26]. However, in this work, we prefer a multi-frequency excitation signal such that the echoes carry richer information about the surroundings for location fingerprinting. FHSS signal has rapidly hopping frequency in a wide band. As it better copes with self-interference, it is suitable for short-range active sensing [58]. When exciting sizable indoor spaces, the self-interference can be easily avoided by limiting the excitation signal duration. Therefore, in this work, we adopt chirp

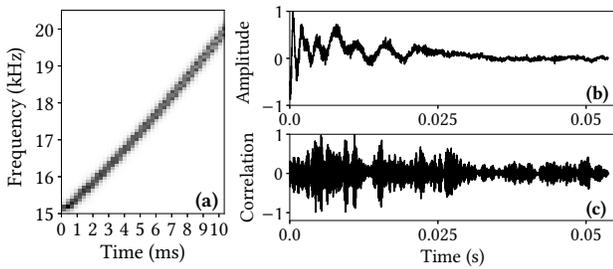


Figure 1: (a) Excitation chirp's spectrogram; (b) Received signal in time domain; (c) Pearson correlation between received signal and excitation signal template.

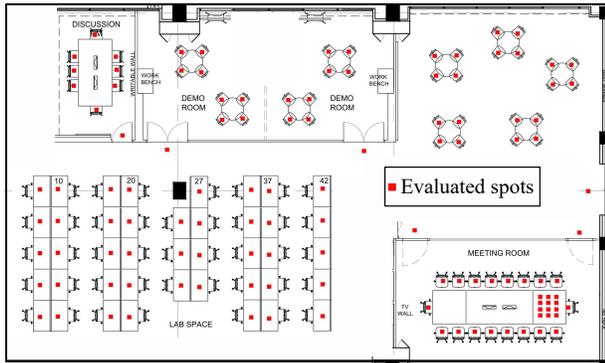


Figure 2: Floor plan of the lab space.

with frequency varying with time continuously that can induce information-richer echoes, compared with FHSS.

As commodity smartphones support audio sampling at 44.1 or 48 ksp/s, they can capture acoustic frequencies up to 22.05 or 24 kHz. To reduce annoyance to human user, the excitation signal should be in the inaudible or near-inaudible frequency range. Although nominal audible frequency is up to 20 kHz, the audible limit of average adults is usually 15 to 17 kHz [1]. Signals with frequencies higher than 20 kHz often suffer from drastic distortions due to the smartphone audio hardware's nonlinearity in the inaudible band [48]. In this paper, we adopt a near-inaudible logarithmic chirp sweeping the 15–20 kHz band within 10 ms as the excitation signal. Fig. 1a shows the spectrogram of the excitation chirp. The chirp uses a relatively wide band (i.e., 5 kHz) for the benefit of pulse compression [32], which helps capture fine-grained spatial features.

■ **Echo extraction:** We develop a program to use the smartphone's loudspeaker to emit the excitation signal and microphone to record the echo at 44.1 ksp/s for 100 ms. We assume the smartphone is held around 30 to 40 cm in front of the chest. Note that the data collection cannot be unobtrusive, e.g., put the smartphone in a pocket. In the received data, we discard the first 10 ms due to the propagation via the direct path from the loudspeaker to the microphone. We also discard the subsequent 1 ms data, which usually contains the first reflection from the human body that is around 30 – 40 cm away from the microphone. The subsequent 50 ms data, which are collectively referred to as *echo trace* and illustrated in Fig. 1b, are used for SLAM. Fig. 1c shows the Pearson correlation

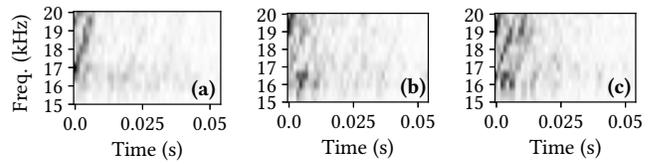


Figure 3: Echo spectrograms obtained at three locations.

between the received signal shown in Fig. 1b and the chirp template. The peaks in Fig. 1c indicate echoes.

4 Measurement Study

To gain insights for the system design, we conduct a set of measurement experiments in a $16 \times 28 \text{ m}^2$ lab space as shown in Fig. 2. We use a Google Pixel 4 smartphone to excite the space at 128 spots indicated by red squares in Fig. 2 and collect 1,700 echo traces at each spot. This section presents the analysis results for these traces.

4.1 Spatial Distinctness of Echoes

We apply short-time Fourier transform (STFT) on the echo data to extract the spectrogram. Specifically, we slide a 96-point Hann window with 48 points of overlap on echo data, resulting in a 49×48 spectrogram. We further discard the frequency bins below 15 kHz, yielding a 12×48 image as the final result. Fig. 3 shows the echo spectrograms collected at three different spots. The frequency dimension of a spectrogram is from 15 to 20 kHz. The differences among the spectrograms suggest that echoes vary across locations. We also conduct experiments to show that chirp excitation brings more spatial distinctness compared to single-tone excitation. The details are omitted due to space limits and can be found in [35].

To understand the distinctness limit, we use the task of supervised learning-based localization to investigate the achievable spatial resolution and scalability with respect to the number of spots.

4.1.1 Spatial resolution. For each location, the spectrograms of the 1,700 echo traces are divided into 1,500 training samples and 200 test samples. The ground truth labels correspond to the spot's location. To understand how the inter-spot distance affects the localization accuracy, we divide the 128 spots into multiple groups with different densities. As a result, the average inter-spot distances of the groups range from 0.25 m to 3 m. For each group, we use echo spectrograms to train a ResNet-18 DNN to classify the spots. We measure both the spot recognition accuracy and the mean localization error. Fig. 4a shows the measured results versus the average inter-spot distance. For each group, we repeat the evaluation process 20 times and plot the error bars. The recognition accuracy remains at around 90%. The mean localization error increases with the inter-spot distance and remains at the sub-meter level. The results suggest that the acoustic echoes can achieve sub-meter spatial resolution.

4.1.2 Scalability. We increase the number of spots handled by a single ResNet-18 model (denoted by k) to understand the scalability. For each setting of k , we randomly draw k spots from the 128 spots, train and test a ResNet-18 model. Note that the selected spots become denser for a larger k . We repeat the process 20 times for each k setting. Fig. 4b shows the results. The recognition accuracy

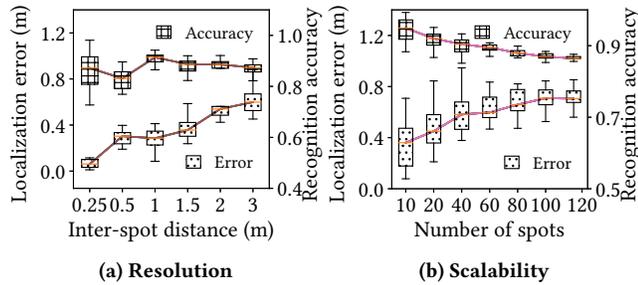


Figure 4: Distinctness limit of acoustic echoes.

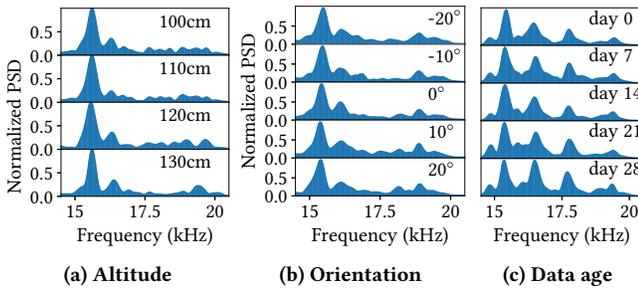


Figure 5: Power spectral densities (PSDs) of the acoustic echoes when several factors vary.

gradually decreases with k and becomes flat when k exceeds 100. This result is consistent with the intuition that the complexity of learning using a DNN increases with the number of classes. The mean localization error increases with k but remains under 1 m and thus at the level of the spatial resolution limit shown in §4.1.1. The above results suggest that the ResNet-18 model does not present a bottleneck when the number of spots is up to 128.

4.2 Robustness of Acoustic Echoes

This section investigates the impacts of several potential affecting factors on acoustic echoes.

4.2.1 Altitude. A researcher holds the phone at different altitudes to simulate the cases where users of different heights hold the phone with natural arm gestures for indoor navigation. As typical adult heights are within 150–194 cm [3], we test at the altitudes from 100 cm to 130 cm (i.e., two thirds of user height). Fig. 5a shows the power spectral densities (PSDs) of the acoustic echoes at different altitudes of the same spot. We can see that the altitude variations of less than 30 cm introduce little impact. Hence, the acoustic echoes are insensitive to variations of user height and hand altitude.

4.2.2 Phone orientation. As a smartphone’s loudspeaker and microphone are not perfectly omnidirectional, phone orientation may affect the received signal at a spot. We evaluate the stability of the received echoes when phone orientation varies. Fig. 5b shows the echo PSDs when the phone has an orientation deviation of -20° to 20° from a certain direction. The results show that the echoes change slightly with less than 40° orientation deviation. When the orientation deviation is larger, the acoustic echoes exhibit larger differences. This suggests that the impact of phone orientation

needs to be properly dealt with when there are multiple echo traces collected at the same spot but in different phone orientations. We will present the solution to this issue in §5.5.

4.2.3 Temporal stability. We collect acoustic echoes at a certain location over a one-month period. In this period, the indoor layout has no significant changes. From Fig. 5c, the echo PSDs are consistent over the time period. In practice, the constructed floor map can be updated whenever a user contributes a trajectory map. The continuous update helps address the potential aging issue over longer periods of time. In §4, we will evaluate the impact of significant changes of the indoor space layout (e.g., furniture re-arrangement) on our system and a mitigation approach beyond map update.

5 Design of ELF-SLAM

From the measurement study, the acoustic echoes exhibit sub-meter spatial distinctness, which is the basis of the fingerprint approach. To unleash the fingerprint approach from laborious labeled training data collection, this section presents the design of ELF-SLAM based on acoustic echoes and IMU data captured by a smartphone during movements. In this section, §5.1 overviews the design of ELF-SLAM. §5.2 presents the graph-based SLAM formulation. §5.3 introduces ELF for loop closure detection. §5.4 presents a clustering-based approach for loop closure curation. §5.5 presents the trajectory map superimposition. §5.6 presents ELF-based localization.

5.1 Approach Overview

As illustrated in Fig. 6, the mapping phase of ELF-SLAM consists of two major components: *trajectory map construction* and *trajectory map superimposition*, where the former focuses on a single trajectory and the latter combines all available trajectories. The trajectories only cover a portion of the indoor space. However, when sufficient trajectories are collected, the combined map can cover the popular locations in the indoor space.

■ **Trajectory map construction:** A user holds a smartphone and moves within the target indoor space to collect the acoustic echoes and IMU data simultaneously on the movement trajectory. The IMU data is used to reconstruct the user’s trajectory via dead reckoning, while the acoustic echo data is used to detect loop closures on the user’s trajectory. The detected loop closures provide critical regulation for the dead reckoning to deal with its long-run drifting problem. Since the generic acoustic features, such as PSD, spectrogram, etc. are ineffective for loop closure detection, ELF-SLAM applies CL to learn a custom and trajectory-specific ELF for loop closure detection. This trajectory-level CL consists of model *pre-training* using synthetic data from a room acoustics simulator and *fine-tuning* using real data collected by the user from the target indoor space. ELF-SLAM detects loop closures based on a custom similarity metric called echo sequence similarity (ESS) between two sequences of ELF traces. Then, a clustering-based approach removes the false positive detection results to curate the loop closures. Lastly, a graph-based SLAM algorithm constructs an accurate trajectory map of ELFs for the user.

■ **Trajectory map superimposition:** When multiple trajectory maps are available (e.g., through crowdsensing), they are superimposed to generate a floor map. The superimposition reconciles different trajectory maps’ ELFs that are collected at the same spot

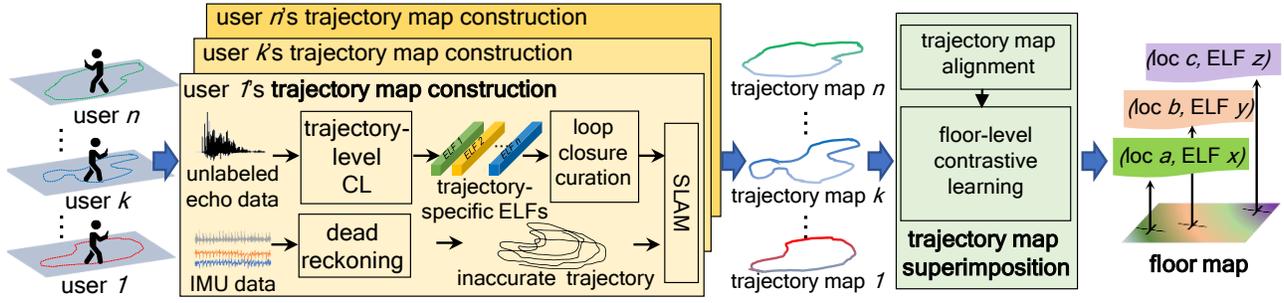


Figure 6: Overview of the mapping phase of ELF-SLAM.

but in different phone orientations. To achieve this, different users' trajectory maps are first aligned into a common coordinate system. The alignment can be achieved based on known initial positions of the users' trajectories (e.g., the entrance of the space) and/or prior knowledge about the accessible passages of the target indoor space [39]. Then, we apply the floor-level CL to train a floor-wide ELF extractor using the acoustic data from all the trajectory maps. The floor map consists of the floor-level ELFs at all spots covered by all the trajectory maps, where each spot is associated with a single floor-level ELF.

Once a map is available, a smartphone can be localized based on its captured echoes in response to a few chirps emitted by the phone. Specifically, the smartphone extracts the ELFs of the echoes and compares them against the map to estimate its location.

5.2 Graph-based SLAM Formulation

Graph-based SLAM [25] constructs a graph whose nodes represent the mobile's poses and edges represent the kinetic constraints relating two poses. When using IMU data to establish the kinetic constraints, loop closure is a vital regulation in combating the long-run drifting problem of IMU-based odometry [44]. In this paper, by letting \mathbf{x}_k denote the node (i.e., location) corresponding to the k^{th} detected footstep, the acoustic echo trace captured between the k^{th} and $(k+1)^{\text{th}}$ footsteps is the measurement associated with the node \mathbf{x}_k and used to detect whether \mathbf{x}_k is at the same location as any previous node (i.e., loop closure detection). The edge connecting two nodes is associated with the IMU-based odometry. The user trajectory is estimated via the graph-based optimization after the loop closures are identified. The estimation method is as follows. For a total of N detected footsteps, let $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ denote the sequence of nodes that describes the user trajectory and $\mathbf{u}_{i,j}$ denote the edge constraint between nodes \mathbf{x}_i and \mathbf{x}_j based on the IMU data. Let C denote the set of footstep index pairs of the detected loop closures. The essence of the trajectory reconstruction can be described by the following optimization problem:

$$\mathbf{X}^* = \operatorname{argmin}_{\mathbf{X}} \sum_{\forall i \in [1, \dots, N-1]} \|f(\mathbf{x}_i, \mathbf{u}_{i,i+1}) - \mathbf{x}_{i+1}\|^2 + \sum_{\forall (i,j) \in C} \|f(\mathbf{x}_i, \mathbf{u}_{i,j}) - \mathbf{x}_j\|^2,$$

where $\|\mathbf{x} - \mathbf{y}\|$ denotes the Euclidean distance between two locations \mathbf{x} and \mathbf{y} , $f(\mathbf{x}_i, \mathbf{u}_{i,j})$ represents the prediction of \mathbf{x}_j based on \mathbf{x}_i and $\mathbf{u}_{i,j}$. In this paper, we implement the SLAM algorithm based on the general graph optimization framework in [31], which also addresses the uncertainty of the prediction $f(\mathbf{x}_i, \mathbf{u}_{i,j})$.

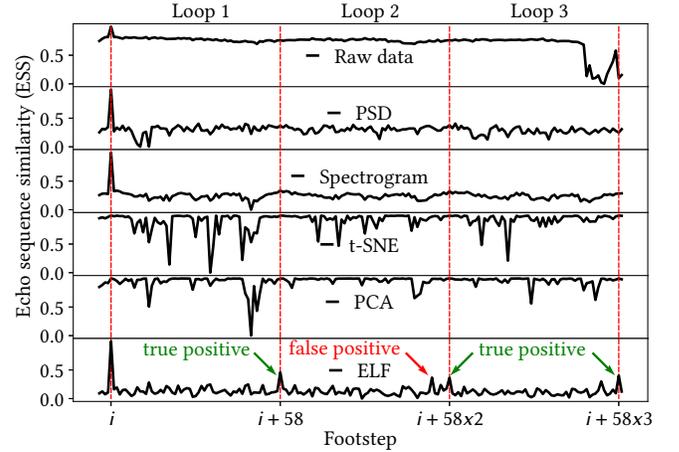


Figure 7: ESS traces with respect to footstep i . Peaks indicate loop closures at footstep $i + 58$, $i + 58 \times 2$, and $i + 58 \times 3$.

5.3 ELF For Loop Closure Detection

Identifying an effective feature for loop closure detection is critical to SLAM. In this section, we first demonstrate the ineffectiveness of the generic features. Then, we propose using CL to construct a learning-based feature.

5.3.1 Ineffectiveness of generic features. We conduct a controlled experiment to evaluate several generic acoustic features. A researcher walks four rounds by following 58 markers pasted on the floor of the lab shown in Fig. 2 and uses a phone to collect acoustic data. As such, each round consists of exact same 58 footsteps. We compute the following features of the echo data: PSD, spectrogram, t-SNE, and principal component analysis (PCA). Then, we compute the similarity between the features collected at footstep i in the first round with those at all footsteps in the four rounds. The similarity is measured by the echo sequence similarity (ESS), which is defined as follows. For two footsteps i and j at which K_i and K_j features are collected, the ESS between them is obtained by averaging the $K_i \times K_j$ pair-wise cosine similarity among the two sets of features. Fig. 7 shows the resulting ESS traces with footstep i (where $i = 4$) in the first round and j being all footsteps of the four rounds sequentially. In this controlled experiment, with respect to the footstep i , loop closures are formed at the footsteps $i + 58$,

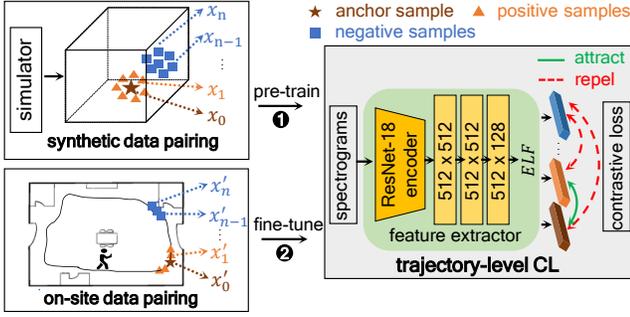
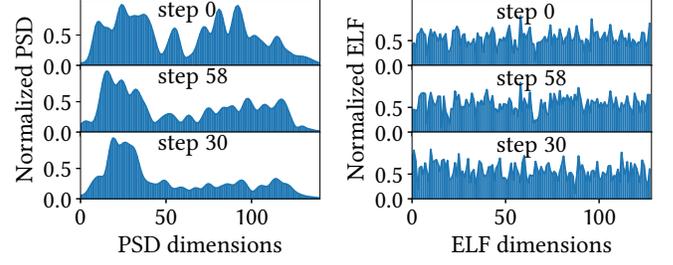


Figure 8: Trajectory-level CL to learn trajectory-specific ELF.

$i + 58 \times 2$, and $i + 58 \times 3$. If the used feature is effective, ESS peaks should be observed at these footsteps. However, from the plots in the first five rows of Fig. 7, no salient peaks are observed at these footsteps. This suggests that the raw data and the used generic feature extraction techniques are ineffective for our loop closure detection problem. Note that although t-SNE is effective for finding feature embeddings of clustered data [7], it is ineffective on the echo samples collected in the spatial continuum that do not exhibit clustered patterns. Note that the ESS peaks are also not observed under other similarity metrics, e.g., those related to Euclidean and Manhattan distances, etc.

5.3.2 Learning-based ELF. The ineffectiveness of generic features motivates us to apply CL to construct a custom feature, i.e., ELF. In what follows, we present our CL design to construct the ELF for loop closure detection. Fig. 8 depicts the workflow. It consists of three steps: *data pairing*, *model pre-training*, and *model fine-tuning*.

Data pairing constructs positive/negative data pairs needed by CL. In image recognition tasks, the positive samples are constructed by introducing spatial perturbations such as resizing, cropping, and blurring, which do not erase the information needed for image recognition. However, such spatial perturbations are not applicable to the echo data, because they destruct the subtle structural information embedded in the echo signal that is related to the smartphone’s location. The design of our data pairing approach is based on the observation that the echoes are similar if collected at close locations and distinct if collected at locations apart. Thus, we construct positive pairs of echoes collected at close locations and negative pairs of echoes at locations apart. Specifically, we treat two consecutive echoes as a positive pair. For each training step, we randomly sample a training batch of 256 such positive pairs from the entire sequence of echoes collected by a certain user. According to our design in §4, the time gap between two consecutive echoes is 0.1 s. As the average human walking speed is 5 km/h, the locations for collecting two consecutive echoes are separated by 0.14 m on average. This average separation is slightly lower than the achievable spatial resolution of the echo modality as evaluated in §4. Thus, viewing two consecutive echoes during the user’s movement as a positive pair is a good heuristic. Then, the cross pairing of the members of the 256 positive pairs gives the negative pairs. Note that the chance of wrongly forming a negative pair with two echo samples collected at a loop closure is low. Such limited wrong forming of negative pairs will not devastate the CL.



(a) $\text{sim}(0,58)=0.40$, $\text{sim}(0,30)=0.42$ (b) $\text{sim}(0,58)=0.61$, $\text{sim}(0,30)=0.22$

Figure 9: (a) PSDs (b) ELFs visualization at steps 0, 30, and 58.

Model pre-training exploits self-supervised learning to build a basic ELF extractor, which will be specialized by the model fine-tuning step presented later. Self-supervised learning often requires abundant unlabeled training data for feature representation learning. To reduce the overhead of data collection in real environments, we use `pyroomacoustic` [46], which is a room acoustics simulator, to generate abundant synthetic training data. Specifically, we simulate a smartphone with a speaker and a microphone separated by 15 cm and generate massive echo data at fine-grained grid points in various simulated rooms with different shapes and sizes. The synthetic echo samples are used to train a feature extractor DNN by CL. We adopt a feature extractor architecture from [15], which consists of a ResNet-18 encoder and a 3-layer projection head. The extractor maps the input spectrogram of an echo trace to a 128-dimensional ELF. As the locations of the synthetic echoes are controlled and thus known, we treat the synthetic echoes collected at locations separated by less than 20 cm as positive pairs, and the rest as negative pairs. The following contrastive loss adopted from [14] is minimized during the model pre-training:

$$\ell_{i,j} = -\log \frac{\exp(\text{sim}(z_i, z_j) / \tau)}{\sum_{k=1}^{2M} \mathbf{1}_{[k \neq i]} \exp(\text{sim}(z_i, z_k) / \tau)},$$

where $\mathbf{1}_{[k \neq i]} \in \{0, 1\}$ is evaluated to 1 if and only if $k \neq i$, $\text{sim}(\cdot, \cdot)$ denotes cosine similarity, z is the extracted feature vector, i and j indicate a positive pair, M is batch size, and τ is the temperature parameter. With the above contrastive loss, the pre-training increases the feature similarity for echoes at close locations and decreases the feature similarity for those at locations apart. Thus, it builds a feature extractor that discriminates the echoes’ locations under the cosine similarity metric. As a result, the loop closure detection can be implemented by comparing the ELFs produced by the feature extractor in terms of the cosine similarity. The necessity of the model pre-training will be evaluated in §6.

Model fine-tuning uses a small amount of unlabeled data collected by users in a specific target space to adapt the pre-trained model to capture the environment-specific characteristics. We follow the same self-supervised CL procedure as described above and construct the data pairs to fine-tune the feature extractor. The resulted feature extractor can generate trajectory-specific ELFs for loop closure detection in the target space.

5.3.3 Loop closure detection using ELF. The last row of Fig. 7 shows the ESS trace computed using ELF. It shows peaks at footstep $i + 58$,

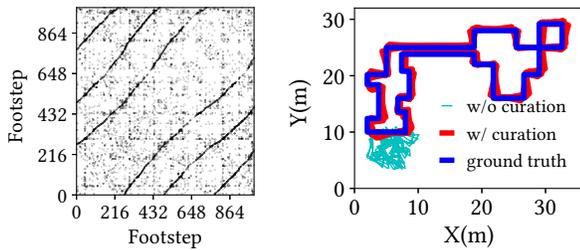


Figure 10: ESS matrix with true loop closures forming trend curves.

Figure 11: Reconstructed trajectories with/without loop closure curation.

$i + 58 \times 2$, and $i + 58 \times 3$ as marked by the green arrows. Fig. 9 shows the cosine similarities using echo features from steps 0, 30 and 58. For PSD, the cosine similarity between loop closure steps 0 and 58 is 0.4, and between the non-loop closure steps 0 and 30 is 0.42. PSD cannot differentiate the loop closure and non-loop closure data. Thus, PSD is ineffective for loop closure detection. For ELF, the cosine similarity between loop closure steps 0 and 58 is 0.61, and between the non-loop closure steps 0 and 30 is 0.22. The ELFs can see significant differences on the loop closure and non-loop closure data. Thus, ELF is effective for loop closure detection.

However, we also observe an unexpected peak close to the footstep $i + 58 \times 2$ as marked by a red arrow, which may lead to a false positive loop closure detection in Fig. 7. Unfortunately, the SLAM is often sensitive to false positive loop closures – some false positives can degrade the SLAM performance [34]. Thus, a loop closure curation algorithm is needed to remove the false positives.

5.4 Loop Closure Curation

5.4.1 Approach design. We propose a clustering-based loop closure curation approach that is based on an ESS matrix defined as follows.

ESS matrix: Consider a user’s trajectory consisting of N footsteps. The pair-wise ESSs between any two footsteps form a $(N - 1) \times (N - 1)$ ESS matrix, where the $(i, j)^{\text{th}}$ element is the ELF-based ESS between the footsteps i and j . Thus, the ESS matrix is symmetric. A large ESS suggests a high similarity between the two involved footsteps’ ELFs, signaling a potential loop closure. Our extensive experiments show that a threshold value of 0.4 for ESS can identify most true positive loop closures while capturing an acceptably low number of false positives to be removed shortly. The ESS matrix is binarized using the threshold, where the positive elements represent candidate loop closures. Fig. 10 shows an example of the binarized ESS matrix constructed using the ELFs collected in a shopping mall. The x- and y-axis represent the footstep index. The black dots in the ESS matrix represent the positive elements.

Clustering-based approach for loop closure curation: The goal of loop closure curation is to remove the false positives from the binarized ESS matrix. Due to the user movement, the true positives in the binarized ESS matrix form trend curves. For instance, consider an ideal case in which the user walks at a constant speed, the true loop closures of footsteps $0, 1, \dots$, and 10 are footsteps $0 + L, 1 + L, \dots$, and $10 + L$, where L is the loop length. As a result, the $(0, 0 + L)^{\text{th}}$, $(1, 1 + L)^{\text{th}}$, \dots , and $(10, 10 + L)^{\text{th}}$ elements of the binarized ESS matrix should be positives and form a trend line. In

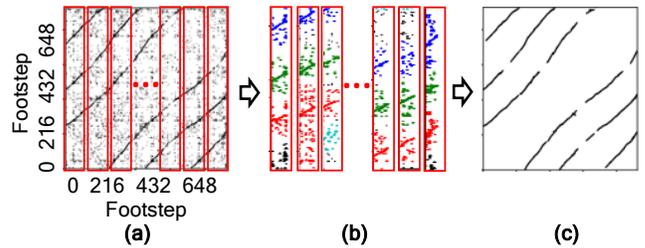


Figure 12: Clustering-based approach for loop closure curation: (a) ESS matrix slicing, (b) clustering in each slice, and (c) concatenated line regression results.

contrast, the false positives tend to appear at random positions in the ESS matrix, as shown in Fig. 10. Based on this observation, we propose a clustering-based approach to isolate the trend curves formed by the true positives from the scattered false positives.

First, we divide the binarized ESS matrix into multiple slices, as illustrated in Fig. 12a. With the slicing, it is easier to identify the true positive clusters in each slice. In our implementation, we set the slice width to be 16 footsteps. Then, for the positives in each slice, we apply the DBSCAN clustering algorithm [19] to identify the number of loops and divide the positives into multiple clusters. This is illustrated by Fig. 12b, where the clusters are differentiated by colors. Although some false positives are classified by DBSCAN as outliers, the remaining false positives close to the trend curves are still in the clusters. To remove these remaining false positives, for the points in each cluster, we apply the random sample consensus (RANSAC) [22] linear regression algorithm to detect a line approximating the trend curve segment. RANSAC is a preferred regression algorithm when many outliers are present. Concatenation of the regressed line segments across all slices gives the clean trend curves. Fig. 12c shows the concatenated results, in which the trend curves formed by the positives are effectively isolated from the scattered noises as shown in Fig. 12a. Lastly, we further curate the loop closure candidates based on the ESS matrix’s symmetric property. Since the negative impact of a false positive on SLAM outweighs that of a false negative, we apply a strategy of only retaining the positives that conform to the symmetric property. Specifically, if the positive at the $(i, j)^{\text{th}}$ position of the ESS matrix has no counterpart positive at the $(j, i)^{\text{th}}$ position, the positive is viewed false and excluded.

5.4.2 Effectiveness of loop closure curation. To demonstrate the impact of the false positives on SLAM, we use all positives in Fig. 10 as the loop closure information to construct the trajectory map. The plot labeled “w/o curation” in Fig. 11 shows the trajectory of the constructed map. We can see that the false positives devastate the trajectory reconstruction. The plot labeled “w/ curation” in Fig. 11 shows the reconstructed trajectory using the curated loop closures. The new trajectory highly resembles the ground truth. The result demonstrates the effectiveness of the proposed clustering-based loop closure curation.

5.5 Trajectory Map Superimposition

The trajectory map constructed from a single user’s data only contains the echo data on a specific trajectory. For real applications, it is desirable to combine many trajectory maps to form a floor map

that covers most/all accessible locations. We assume that each trajectory map’s initial position relative to the indoor space is known. For instance, the SLAM mobile app may prompt the user to start the process from the entrance of the indoor space. When there are multiple entrances, location tagging [52] can be used to recognize the actual entrance. With the known initial position, the trajectory maps can be collated onto a common coordinate system. However, the trajectory maps crossing a certain spot from different entering directions may have different echo data for the spot, due to the dependency of echo data on phone orientation as shown in §4. Thus, we need to reconcile such differences.

We propose a floor-level CL approach to train a unified feature extractor for map superimposition. It shares the same model pre-training workflow as the trajectory-level CL except the data pairing approach for model fine-tuning. Specifically, the echo data collected at the same location regardless of the phone orientation are treated as positive pairs, whereas those collected from different locations are treated as negative pairs. The feature extractor trained via the floor-level CL generates the floor-level ELF’s covering all trajectory maps. As the quality of the floor map is related to its spatial coverage, this floor-level CL approach needs to scale well with the number of locations. In §6, we evaluate this approach in handling 4,000 different locations with four phone orientations at each location.

When falling back to the scheme of learning a location recognition model in the supervised manner, a possible approach to mitigate the echo data’s sensitivity on phone orientation is to form a training dataset with echo data and location labels (regardless of orientation) from all the trajectory maps. In §6, we will compare the localization performance of this supervised learning approach with our approach based on the floor map.

5.6 Localization

Once a map (either a trajectory map or floor map) is constructed, a smartphone’s location can be determined after capturing the echoes in response to the chirps. We consider two localization approaches, i.e., *one-shot localization* and *trajectory localization*, which are suitable for the scenarios where the user stands still and moves, respectively. In the former, an ELF sequence containing multiple consecutive echoes collected at a spot is matched against the map in terms of the ESS to determine the location. In the latter, both the ELF sequence and the IMU data during the user’s movement over a short time period are used for localization. Specifically, we apply dead reckoning to the IMU data to estimate the user’s trajectory, and then apply a curve matching algorithm [17] to find the candidate segments in the map that resemble the user’s trajectory. The candidate segment that has the largest average ESS from the captured ELF sequence is the output of the trajectory localization.

6 Performance Evaluation

6.1 Experiment Setup

Evaluation environments: We evaluate ELF-SLAM in three indoor environments: a living room (60 m²), an office (360 m²), and a shopping mall (2,000 m²). The floor plans are shown in Fig. 13a. To conduct comparative evaluation side by side, we employ the SLAM systems using two smartphone’s built-in sensing modalities, i.e., Wi-Fi RSSI and geomagnetism, as the baselines. This is the same as

Table 2: Mapping error statistics.

Modality	Living room			Office			Mall		
	\bar{x} ¹	\bar{x} ²	Q3 ³	\bar{x}	\bar{x}	Q3	\bar{x}	\bar{x}	Q3
ELF	0.10	0.10	0.14	0.63	0.63	0.80	0.45	0.53	0.69
ELF w/o pre-train	0.73	0.82	1.25	1.69	1.68	1.94	1.16	1.14	1.42
Wi-Fi	0.44	0.45	0.55	1.52	1.54	2.06	1.24	1.26	1.54
Geomag	0.56	0.55	0.64	1.14	1.24	1.82	0.79	0.81	1.05

¹ Median error ² Mean error ³ Third quartile of the error

the evaluation methodology adopted in [34] that studies powerline EMR SLAM. Note that we also compare the results of ELF-SLAM and EMR SLAM. To implement Wi-Fi SLAM, we deploy five Wi-Fi access points (APs) in the living room and office, as illustrated by the stars in Fig. 13a, that provide sufficient Wi-Fi coverage. The large shopping mall has dense APs deployed by the tenants. The number of Wi-Fi APs observable is around 5 to 10 when conducting experiments in the mall. Random people are walking in the shopping mall when we collect the data.

Data collection: We develop an Android app and run it on a Google Pixel 4 smartphone to collect acoustic echoes, Wi-Fi RSSI, geomagnetic field signals, and IMU data. The app uses the `WifiManager` Android API to scan the surrounding Wi-Fi APs and collect RSSI data at a sampling rate of 0.8 sps. We do not sample Wi-Fi channel state information (CSI), because CSI sampling requires rooting a phone [28]. The app also samples the phone’s built-in magnetometer at 50 sps. During data collection, the user holds the phone with one hand around 30 to 40 cm in front of the chest and walks on a marked trajectory for multiple rounds in each of the evaluation environments. Note that the purpose of trajectory marking is to obtain the location ground truth.

Loop closure detection for baseline modalities: For Wi-Fi SLAM, we use the Euclidean distance between two Wi-Fi RSSI vectors for loop closure detection [56]. For geomagnetic SLAM, we first normalize the triaxial magnetic data and then use dynamic time warping distance as the metric for loop closure detection [54]. We also apply the loop closure curation and graph-based optimization algorithms described in §5 to the baseline SLAM systems.

6.2 Trajectory Map Construction Performance

Fig. 13a shows the floor plans and ground-truth trajectories. Fig. 13b shows the trajectories reconstructed via IMU-based dead reckoning, which deviate heavily from the ground truth due to the long-run drift problem. Fig. 13c, 13d, and 13e show the trajectories reconstructed by the SLAM systems using ELF, Wi-Fi RSSI, and geomagnetism. The trajectories reconstructed by ELF-SLAM are the closest to the ground truth. Table 2 lists the detailed mapping error statistics of the three modalities. ELF-SLAM achieves sub-meter mapping accuracy in all three environments, whereas Wi-Fi SLAM and geomagnetic SLAM’s mapping errors increase in the large indoor space, i.e., office and mall. In [34], EMR SLAM using the smartphone earphone as the side-channel sensor yields about 1 m to 2 m median mapping errors in the evaluated office and lab spaces. Thus, ELF-SLAM outperforms Wi-Fi SLAM, geomagnetic SLAM, and EMR SLAM in map construction.

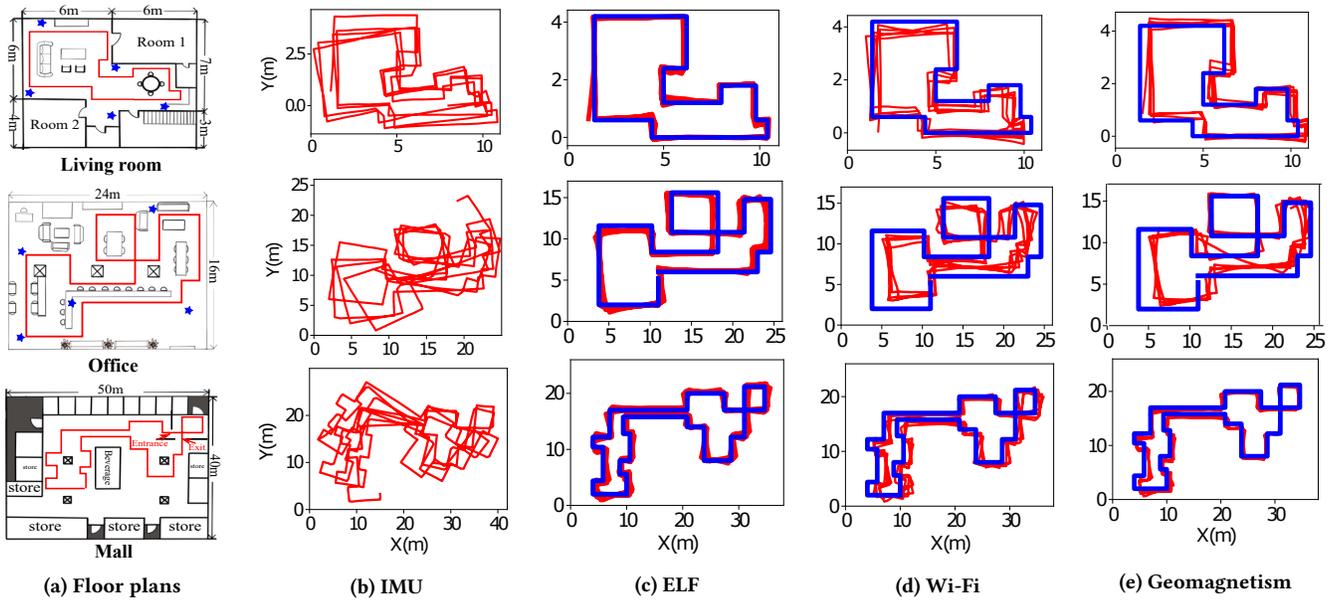


Figure 13: Floor plans and trajectory reconstruction results in three evaluation environments.

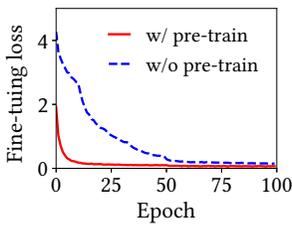


Figure 14: Trend of model finetuning loss.

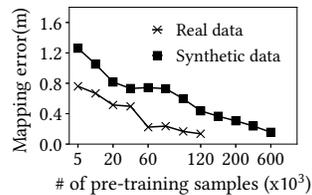


Figure 15: Impact of type and amount of pre-training data.

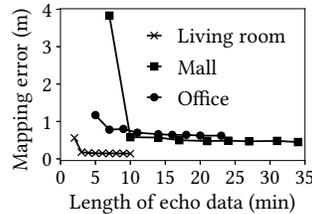


Figure 16: Impact of finetuning data volume.

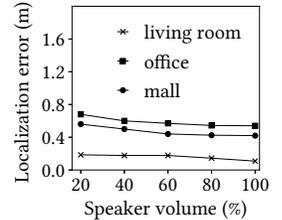


Figure 17: Impact of speaker volume.

6.3 Effectiveness of CL Pre-training

We investigate the necessity of the CL pre-training by comparing the SLAM performance using the ELF extractors learned with and without the model pre-training. The row “ELF w/o pre-train” in Table 2 shows the case without model pre-training. Compared with the result with model pre-training (row “ELF”), the median errors increase to 0.73 m, 1.16 m, and 1.69 m, respectively. Fig. 14 shows the fine-tuning loss. The fine-tuning loss on the pre-trained model converges more quickly and is lower than that on the model without pre-training. The above results show that the model pre-training improves the efficiency of CL.

We also investigate the model pre-training’s requirement on data. We use two types of data for the pre-training, i.e., the synthetic data generated by the room acoustics simulator and real data collected from the different spaces. For the latter, we use real echoes collected in the lab space, office, and shopping mall for model pre-training. Then, we perform fine-tuning and evaluation in the living room. Fig. 15 shows the median mapping error versus the amount of pre-training data. The mapping errors decrease with pre-training data volume for both types of data. If real data is used for pre-training, the median mapping error converges to about 0.13 m when data volume is more than 120 K. If synthetic data is used, the error

converges to about 0.15 m when data volume is more than 600 K. This shows that both the real data and synthetic data are useful for model pre-training. Although real data shows better efficiency in supporting the pre-training, synthetic data can be generated at massive scales and thus suffice for the pre-training.

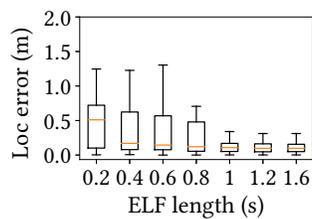
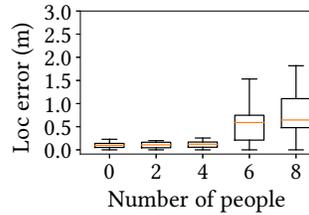
6.4 Impact of Finetuning Data Volume

We investigate the impact of the finetuning data volume on the performance of trajectory map construction. We gradually decrease the used data volume for model finetuning and show the trajectory map construction results in Fig. 16. In the living room, the median mapping error remains around 0.15 m when the used data volume decreases from 10 min to 3 min. The error increases to 0.56 m when the used data volume is 2 min. The increased error is due to no loop closure formed and the estimated trajectory cannot be corrected. Similarly, the mapping errors remain low when the used data volume is larger than 7 min and 10 min where loop closures can be detected via the ELF’s in the office and the mall, respectively. The results show that the ELF-SLAM only requires a few minutes of data for model finetuning as long as there is loop closure formed during the walking.

Table 3: Localization error statistics.

Modality	Living room			Office			Mall		
	\tilde{x}^1	\tilde{x}^2	Q3 ³	\tilde{x}	\tilde{x}	Q3	\tilde{x}	\tilde{x}	Q3
One-shot localization									
ELF	0.10	0.29	0.14	0.54	0.60	0.80	0.42	0.79	0.67
Wi-Fi	1.67	2.17	3.30	3.44	4.27	6.16	3.04	3.86	5.22
Geomag	1.06	2.31	3.93	2.19	3.95	5.38	12.48	13.36	19.28
Trajectory localization									
ELF	0.10	0.22	0.64	0.41	0.54	0.97	0.47	0.53	0.86
Wi-Fi	0.54	0.730	1.13	1.74	1.86	3.09	1.46	1.90	3.73
Mag	0.56	0.56	0.78	1.75	1.81	2.39	8.70	8.29	14.59

¹ Median error ² Mean error ³ Third quartile of the error

**Figure 18: ELF sequence length.****Figure 19: Moving people.**

6.5 Localization Performance

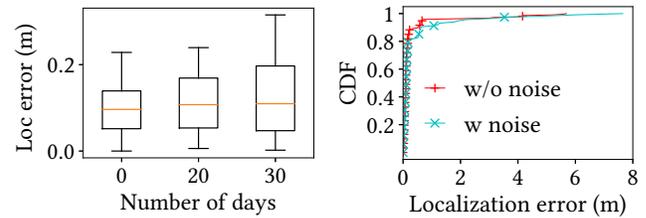
We evaluate both the one-shot localization and trajectory localization of the three sensing modalities. Table 3 lists the localization error statistics. For one-shot localization, ELF-SLAM achieves sub-meter median error in the three environments and outperforms both Wi-Fi SLAM and geomagnetic SLAM. For trajectory localization, each short trajectory consists of 8 consecutive footsteps. For Wi-Fi and geomagnetic SLAMs, the trajectory localization errors are less than the one-shot localization errors. For ELF-SLAM, trajectory localization does not bring much accuracy improvement over the one-shot localization, because the latter has already achieved a high localization accuracy close to the modality’s spatial resolution. However, we will shortly show in §6.6 that the trajectory localization brings performance improvement when ELF-SLAM is affected by various affecting factors.

6.6 Sensitivity Analysis for Localization

We conduct experiments mainly in the living room to evaluate the sensitivity of ELF-SLAM to various factors. By default, we consider one-shot localization.

6.6.1 Speaker volume. We evaluate the localization with various settings for the smartphone’s loudspeaker volume in emitting the chirps. Fig. 17 shows the results. When the volume decreases from 100% (i.e., the highest volume) to 20%, the median localization errors in the living room, office, and shopping mall increase from 0.1 m, 0.54 m, and 0.42 m to 0.18 m, 0.68 m, and 0.56 m, respectively. Note that with 20% loudspeaker volume, on the audible frequency band, the smartphone’s sound is soft and becomes nearly imperceptible in environments with normal noise levels. Thus, ELF-SLAM maintains sub-meter accuracy when the chirp emission is soft.

6.6.2 ELF sequence length. Fig. 18 shows the localization errors when we vary the length of the ELF sequence used for computing

**Figure 20: Aging.****Figure 21: Audible noises.**

ESS from 0.2 s to 1.6 s. A boxplot shows the localization error distribution. The horizontal line in each boxplot shows the median. We can see that the localization error decreases with the ELF sequence length and becomes flat when the sequence length is more than 1 s. Note that at human’s average walking speed, the duration between two consecutive footsteps is about 0.6 s, which results in an ELF sequence length of 0.6 s as well. From Fig. 18, at this length setting, the one-shot localization median error is around 0.1 m. Thus, ELF-SLAM performs well when the user walks at a normal speed.

6.6.3 Nearby moving people. Human bodies can reflect the excitation chirp and generate echoes irrelevant to ELF. Thus, we evaluate the impact of the nearby moving people on one-shot localization. We ask multiple volunteers to walk freely in the living room and talk to each other during the localization phase. Fig. 19 shows the localization error versus the number of nearby moving people. The localization error remains low when the number of people is up to 4. Note that the tested area is only about 60 m². When there are 6 and 8 moving people, whose crowd density is similar to that in the shopping mall during peak hours, the median localization errors increase to 0.57 m and 0.6 m. Nevertheless, the errors remain at the sub-meter level. Thus, ELF-based localization can tolerate nearby moving people to a certain extent.

6.6.4 Map aging. We investigate whether the map constructed by ELF-SLAM ages. Specifically, at day 0, we use ELF-SLAM to construct a trajectory map. Then, we evaluate the ELF-based localization performance multiple times during a one-month period. Fig. 20 shows the results. The median localization errors are 0.10 m, 0.11 m, and 0.12 m at day 0, 20, and 30, respectively. This suggests that the constructed map does not have salient aging issue. In practice, a map can be continuously updated using the latest data contributed by users, to mitigate any potential aging issue.

6.6.5 Audible noises. We evaluate the robustness of ELF-based localization against audible noises. We use a laptop computer to play video clips of different contents (music, speech, etc) from Youtube to generate the noises. From Fig. 21, the noises have little impact on the localization performance. This is because our system operates within the near-inaudible frequency band. Thus, audible noises have negligible impact on ELF-based localization.

6.6.6 Space layout changes. The layout changes of the target space may have impact on the chirp reverberation processes. Thus, we deliberately change the furniture locations in the living room to evaluate such impact. Fig. 22a illustrates how the furniture objects are moved. Specifically, we move five objects including a dining table, a tea table, a TV cabinet, and two sofas. We move one object at a time. Fig. 22b shows the localization error versus the number of

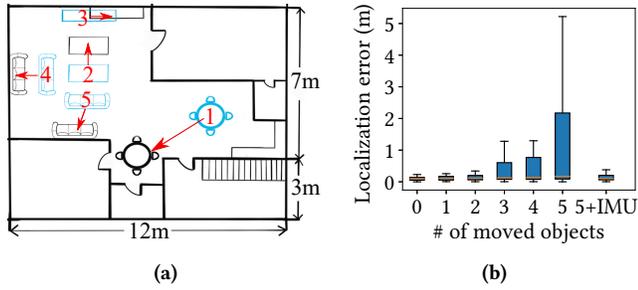


Figure 22: (a) Movements of furniture objects in a living room; (b) the corresponding localization performance.

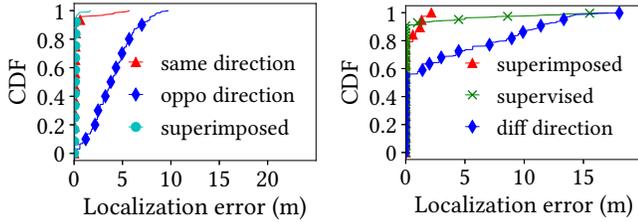


Figure 23: Trajectory map superimposition on small-scale real echoes.

Figure 24: Trajectory map superimposition on large-scale synthetic echoes.

moved objects. The error remains low when the number of moved objects is less than 5. When all the 5 objects are moved, the mean localization error increases to 1.3 m. If we apply the trajectory localization, the mean localization error decreases to 0.3 m as labeled by “5+IMU” in Fig. 22b. Therefore, the trajectory localization improves the robustness of ELF-based localization against the layout changes. In practice, a map can be continuously updated using the latest data from the users to mitigate the impact of layout changes.

6.7 Trajectory Map Superimposition

We evaluate the performance of the one-shot localization on the map superimposed from trajectory maps as described in §5.5.

6.7.1 Evaluation on a small-scale dataset. We conduct experiments in the living room. We follow the route in Fig. 13a and walk in two opposite directions to generate two different trajectory maps. Then, we apply the proposed CL approach for map superimposition. Fig. 23 shows the localization results. The plot labeled “same direction” is obtained when the smartphone during the localization phase is in the same orientation as the used map. The median localization error is 0.1 m. The curve labeled “oppo direction” is for the case when the smartphone during localization is in the opposite orientation as the trajectory map. The median localization error increases up to 3.8 m. The increased error is due to the sensitivity of ELF to large phone orientation deviations. The curve labeled “superimposed” shows the localization results using the map superimposed via the CL approach. The median localization error is 0.1 m, which is the same as the “same direction” result. This small-scale experiment shows that the map superimposed via CL can improve the ELF-based localization performance when trajectory maps with opposite directions are available.

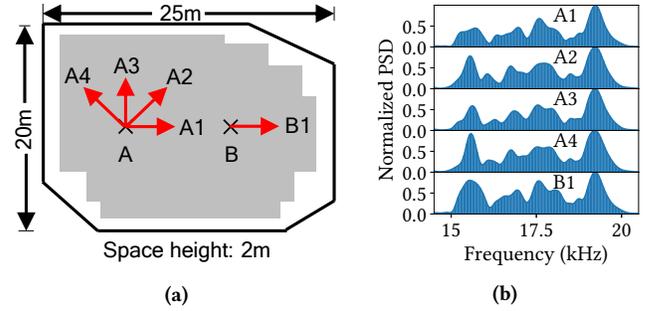


Figure 25: (a) Simulated space for extended evaluation of floor-level CL. (b) Spot A’s echo PSDs on directions 1 to 4 and Spot B’s echo PSD at direction 1.

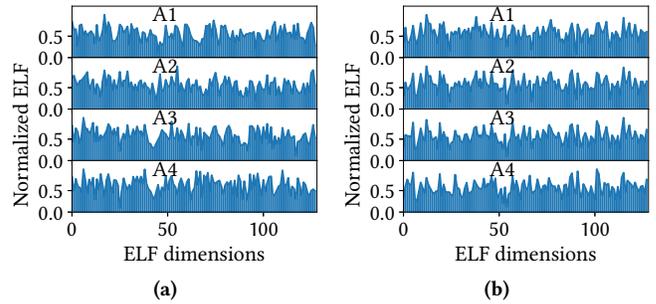


Figure 26: ELF’s average similarity is (a) 0.34 before map superimposition and (b) 0.76 after map superimposition.

6.7.2 Evaluation on a large-scale synthetic dataset. We also evaluate whether the floor-level CL can handle massive echo data from many trajectory maps. This experiment omits trajectory map construction and only focuses on evaluating the superimposition performance, in terms of the one-shot localization error using the floor map. To allow the evaluation to easily scale up, we use the `pyroomacoustic` simulator to generate a large-scale synthetic dataset for an indoor space as shown in Fig. 25a. We collect data at 4,000 spots in the grey area. At each spot, we collect echo data when the simulated smartphone points to the directions as marked by the red arrows at spot A in Fig. 25a. For each orientation, we collect 100 echo traces with random perturbations to the orientation such that the traces are slightly different. As a result, a total of 16 million echo traces are collected. The first four rows of Fig. 25b show the synthetic echoes’ PSDs for four directions at spot A. They are slightly different from each other. The last row of Fig. 25b shows the echo’s PSD at spot B. It is different from all PSDs synthesized at spot A. This shows that the used simulator can synthesize both the orientation- and location-dependent echoes. Note that, from our estimation, collecting the same amount of data in the real world requires about 400 hours of manual labor. This experiment using synthetic data focuses on evaluating the scalability of map superimposition algorithm. We leave the pilot deployment of our system that crowdsources many users’ data for map superimposition to our future work.

Fig. 26 shows the spot A’s ELF’s from directions 1 to 4 before/after map superimposition. The ELF’s average cosine similarity is 0.34

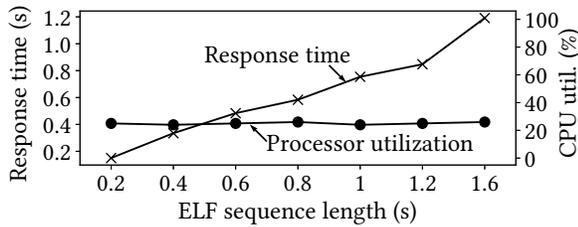


Figure 27: Model execution overhead.

before the map superimposition. This value increases to 0.76 after applying the floor-level CL for map superimposition. This result shows that map superimposition is effective in reconciling the ELF's collected in different orientations. Fig. 24 shows the localization results on the synthetic data. The plot labeled “diff direction” shows the CDF when the CL-based map superimposition is not applied and the evaluated samples are in a different phone orientation from that in the used map. The mean localization error is 3.2 m. This poor result shows the necessity of the CL-based reconciliation. The curve labeled “superimposed” shows the results obtained using the floor map constructed by the floor-level CL. The mean localization error decreases to 0.24 m. We also employ the supervised fingerprint approach as a baseline, which forms the training dataset by labeling the echoes synthesized at the same spot with the same location label and trains a DNN to classify the 4,000 spots. The CDF curve labeled “supervised” shows the results. The mean localization error is 0.56 m. The supervised fingerprint approach is inferior to the proposed solution that performs localization using the floor map.

6.8 System Overhead

We evaluate the computation overheads of the ELF extractor and real-time one-shot localization using the floor map on the Google Pixel 4 smartphone. We use Pytorch-Mobile [6] to optimize and deploy the model. The compressed model is about 96 MB.

6.8.1 App’s response time and processor utilization. When the ELF sequence length varies from 0.2 s to 1.6 s, from Fig. 27, the smartphone processor utilization remains at around 20%. The memory usage of storing 4,000 spots’ ELF’s is less than 4 MB. We also measure the app’s response time, which includes the times for extracting ELF’s and matching the ELF trace against the floor map. The response time increases from 0.18 s to 1.2 s, when the ELF sequence length varies from 0.2 s to 1.6 s. The increased response time is from the localization module, because the computation overhead of the feature matching increases with the ELF sequence length. From Fig. 18, by setting the ELF sequence length to be 0.6 s, our system achieves 0.1 m median localization error, while the corresponding measured response time is about 0.5 s. Thus, the user can get the localization result in about 1.1 s.

6.8.2 App’s network bandwidth and battery usages. To continuously transmit echo and IMU data to the cloud server for map construction, the app’s bandwidth usage is around 90 kbps. This data rate is similar to that of Advanced Audio Coding (AAC), a widely adopted standard for lossy audio compression. Note that as the localization phase of ELF-SLAM is performed locally on the phone, it requires no data transmission. We use the battery historian [5] to estimate the app’s energy usage. The app’s energy usage per hour is

around 270 mAh when the app performs localization continuously. This energy usage is similar to that of the Google Map app in continuous navigation, i.e., around 280 mAh and much lower than a visual SLAM [4], whose energy consumption is around 450 mAh. Thus, our ELF-based localization system introduces acceptable overhead.

7 Discussion

■ **Concurrent use and security.** It is common to have multiple users simultaneously use their smartphones for location sensing in same indoor space. To avoid signal collision, we will implement the carrier-sense multiple access (CSMA) protocol in the ELF-SLAM system to manage the traffic. Specifically, when a smartphone is about to transmit the chirp, it uses its microphone to detect whether there is an ongoing echoing process. If so, the smartphone will defer the chirp play. The ELF-SLAM can be vulnerable to malicious attacks. For example, attackers can deploy speakers and play the sound in the used band to mislead the system. The development of the defense mechanism will be considered in future work.

■ **Domain adaptation across different devices.** The device heterogeneity could affect the performance of ELF-SLAM. Several machine learning techniques are available to address the domain shift problem, e.g., data augmentation [38], few-shot learning [24], and adversarial learning [37], etc. We will explore an effective way to address device heterogeneity in future work.

■ **Impact of inaudible sound to human.** The experimental results in [53] show that participants experience mild side effect after 20 minutes’ exposures in frequency from 12.5 Hz to 20 kHz with sound level from 44 to 71 dB. We investigate the played chirp level on Google Pixel 4 by increasing the sound volume from 20% to 100%, and find that the corresponding sound level is between 40 - 68 dB. As shown in Fig. 16, ELF-SLAM requires less than 20 min data collection in the tested environment. Thus, ELF-SLAM will have little impact on humans within tens of minutes exposure.

8 Conclusion

This paper presents ELF-SLAM, an indoor smartphone SLAM system using learning-based ELF’s. ELF-SLAM uses a smartphone’s built-in audio hardware to emit near-inaudible chirps and record acoustic echoes in an indoor space, then uses the echoes to detect loop closures for regulating the IMU-based dead reckoning. To effectively capture loop closures, we design a trajectory-level contrastive learning procedure and apply it on the echoes to learn ELF’s. Then, we design a clustering-based approach to remove the false detection results and curate the loop closures. Lastly, we design floor-level contrastive learning to superimpose the trajectory maps such that the differences among the ELF’s collected in different phone orientations are reconciled. Our extensive experiments show that ELF-SLAM achieves sub-meter accuracy in both mapping and localization, and outperforms Wi-Fi RSSI and geomagnetic SLAMs.

Acknowledgments

This study is supported under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from Singapore Telecommunications Limited (Singtel), through Singtel Cognitive and Artificial Intelligence Lab for Enterprises (SCALE@NTU).

References

- [1] 2022. The Audible Spectrum. <https://www.ncbi.nlm.nih.gov/books/NBK10924/>.
- [2] 2022. Google Play Top Charts. <https://play.google.com/store/apps/top>.
- [3] 2022. Human Height. <https://ourworldindata.org/human-height>.
- [4] 2022. Monocular Visual Odometry. https://github.com/sunzuolei/mvo_android.
- [5] 2022. Profile battery usage with Batterystats and Batter historian. <https://developer.android.com/topic/performance/power/setup-battery-historian>.
- [6] 2022. Pytorch Mobile. <https://pytorch.org/mobile/home/>.
- [7] Sanjeev Arora, Wei Hu, and Praveesh K Kothari. 2018. An analysis of the t-sne algorithm for data visualization. In *Conference On Learning Theory*. PMLR.
- [8] Aditya Arun, Roshan Ayyalasomayajula, William Hunter, and Dinesh Bharadia. 2022. P2SLAM: Bearing based WiFi SLAM for Indoor Robots. *IEEE Robotics and Automation Letters* (2022).
- [9] Martin Azizyan, Ionut Constandache, and Romit Roy Choudhury. 2009. SurroundSense: mobile phone localization via ambience fingerprinting. In *Proceedings of the 15th annual international conference on Mobile computing and networking*. 261–272.
- [10] Paramvir Bahl and Venkata N Padmanabhan. 2000. RADAR: An in-building RF-based user location and tracking system. In *Proceedings IEEE INFOCOM 2000. Conference on computer communications. Nineteenth annual joint conference of the IEEE computer and communications societies (Cat. No. 00CH37064)*, Vol. 2. Ieee, 775–784.
- [11] Chao Cai, Rong Zheng, and Jun Luo. 2022. Ubiquitous Acoustic Sensing on Commodity IoT Devices: A Survey. *IEEE Communications Surveys & Tutorials* (2022).
- [12] Daniele Cattaneo, Matteo Vaghi, and Abhinav Valada. 2022. Lcdnet: Deep loop closure detection and point cloud registration for lidar slam. *IEEE Transactions on Robotics* (2022).
- [13] Jagmohan Chauhan, Yining Hu, Suranga Seneviratne, Archan Misra, Aruna Seneviratne, and Youngki Lee. 2017. BreathPrint: Breathing acoustics-based user authentication. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*.
- [14] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*. PMLR, 1597–1607.
- [15] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. 2020. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029* (2020).
- [16] Yin Chen, Dimitrios Lymberopoulos, Jie Liu, and Bodhi Priyantha. 2012. FM-based indoor localization. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*. 169–182.
- [17] Ming Cui, John Femiani, Jiuxiang Hu, Peter Wonka, and Anshuman Razdan. 2009. Curve matching for open 2D curves. *Pattern Recognition Letters* 30, 1 (2009), 1–10.
- [18] David Droschel and Sven Behnke. 2018. Efficient continuous-time SLAM for 3D lidar-based online mapping. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 5000–5007.
- [19] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Knowledge Discovery and Data Mining*.
- [20] Christine Evers and Patrick A Naylor. 2018. Acoustic slam. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 26, 9 (2018), 1484–1498.
- [21] Brian Ferris, Dieter Fox, and Neil D Lawrence. 2007. Wifi-slam using gaussian process latent variable models. In *IJCAI*. 2480–2485.
- [22] Martin A Fischler and Robert C Bolles. 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM* (1981).
- [23] Ruipeng Gao, Yang Tian, Fan Ye, Guojie Luo, Kaigui Bian, Yizhou Wang, Tao Wang, and Xiaoming Li. 2015. Sextant: Towards ubiquitous indoor localization service by photo-taking of the environment. *IEEE Transactions on Mobile Computing* 15, 2 (2015), 460–474.
- [24] Taesik Gong, Yeonsu Kim, Jinwoo Shin, and Sung-Ju Lee. 2019. MetaSense: few-shot adaptation to untrained conditions in deep mobile sensing. In *Proceedings of the 17th Conference on Embedded Networked Sensor Systems (SenSys)*. 110–123.
- [25] Giorgio Grisetti, Rainer Kümmerle, Cyrill Stachniss, and Wolfram Burgard. 2010. A tutorial on graph-based SLAM. *IEEE Intelligent Transportation Systems Magazine* (2010).
- [26] Sidhant Gupta, Daniel Morris, Shwetak Patel, and Desney Tan. 2012. Soundwave: using the doppler effect to sense gestures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.
- [27] Suining He and Kang G Shin. 2017. Geomagnetism for smartphone-based indoor localization: Challenges, advances, and comparisons. *ACM Computing Surveys (CSUR)* 50, 6 (2017), 1–37.
- [28] Steven M Hernandez and Eyuphan Bulut. 2020. Performing WiFi sensing with off-the-shelf smartphones. In *2020 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. IEEE, 1–3.
- [29] Jeffrey Hightower, Sunny Consolvo, Anthony LaMarca, Ian Smith, and Jeff Hughes. 2005. Learning and recognizing the places we go. In *International Conference on Ubiquitous Computing*. Springer, 159–176.
- [30] Miranda Kreković, Ivan Dokmanić, and Martin Vetterli. 2016. EchoSLAM: floor with acoustic echoes. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Ieee, 11–15.
- [31] Rainer Kümmerle, Giorgio Grisetti, Hauke Strasdat, Kurt Konolige, and Wolfram Burgard. 2011. g 2 o: A general framework for graph optimization. In *2011 IEEE International Conference on Robotics and Automation*. IEEE, 3607–3613.
- [32] Patrick Lazik and Anthony Rowe. 2012. Indoor pseudo-ranging of mobile devices using ultrasonic chirps. In *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems*. 99–112.
- [33] Jie Lian, Jiadong Lou, Li Chen, and Xu Yuan. 2021. EchoSpot: Spotting Your Locations via Acoustic Sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* (2021).
- [34] Chris Xiaoxuan Lu, Yang Li, Peijun Zhao, Changhao Chen, Linhai Xie, Hongkai Wen, Rui Tan, and Niki Trigoni. 2018. Simultaneous localization and mapping with power network electromagnetic field. In *Proceedings of the 24th annual international conference on mobile computing and networking*. 607–622.
- [35] Wenjie Luo, Qun Song, Zhenyu Yan, Rui Tan, and Guosheng Lin. 2022. Indoor Smartphone SLAM with Learned Echoic Location Features. *arXiv preprint arXiv:2210.08493* (2022).
- [36] Jan Willem Marck, Ali Mohamoud, Eric vd Houwen, and Rob van Heijster. 2013. Indoor radar SLAM A radar application for vision and GPS denied environments. In *2013 European Radar Conference*. IEEE, 471–474.
- [37] Akhil Mathur, Anton Isopoussu, Fahim Kawsar, Nadia Berthouze, and Nicholas D Lane. 2019. Mic2mic: using cycle-consistent generative adversarial networks to overcome microphone variability in speech systems. In *Proceedings of the 18th international conference on information processing in sensor networks*. 169–180.
- [38] Akhil Mathur, Tianlin Zhang, Sourav Bhattacharya, Petar Velickovic, Leonid Joffe, Nicholas D Lane, Fahim Kawsar, and Pietro Lió. 2018. Using deep data augmentation training to address software and hardware heterogeneities in wearable and smartphone sensing devices. In *2018 17th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, 200–211.
- [39] Michael Montemerlo, Sebastian Thrun, Daphne Koller, Ben Wegbreit, et al. 2002. FastSLAM: A factored solution to the floor problem. *AAAI* (2002).
- [40] Rajalakshmi Nandakumar, Vikram Iyer, Desney Tan, and Shyamnath Gollakota. 2016. Fingero: Using active sonar for fine-grained finger tracking. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*.
- [41] Rajalakshmi Nandakumar, Alex Takakuwa, Tadayoshi Kohno, and Shyamnath Gollakota. 2017. Covertband: Activity information leakage using music. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* (2017).
- [42] Chunyi Peng, Guobin Shen, Yongguang Zhang, Yanlin Li, and Kun Tan. 2007. Beepbeep: a high accuracy acoustic ranging system using cots mobile devices. In *Proceedings of the 5th international conference on Embedded networked sensor systems*.
- [43] Swadhin Pradhan, Ghufuran Baig, Wenguang Mao, Lili Qiu, Guohai Chen, and Bo Yang. 2018. Smartphone-based acoustic indoor space mapping. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* (2018).
- [44] Cliff Randell, Chris Djalllis, and Henk Muller. 2003. Personal position measurement using dead reckoning. In *Seventh IEEE International Symposium on Wearable Computers, 2003. Proceedings*. IEEE Computer Society.
- [45] Mirco Rossi, Julia Seiter, Oliver Amft, Seraina Buchmeier, and Gerhard Tröster. 2013. RoomSense: an indoor positioning system for smartphones using active sound probing. In *Proceedings of the 4th Augmented Human International Conference*. 89–95.
- [46] Robin Scheibler, Eric Bezzam, and Ivan Dokmanić. 2018. Pyroomacoustics: A python package for audio room simulation and array processing algorithms. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 351–355.
- [47] Sheng Shen, Daguang Chen, Yu-Lin Wei, Zhijian Yang, and Romit Roy Choudhury. 2020. Voice localization using nearby wall reflections. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. 1–14.
- [48] Qun Song, Chaojie Gu, and Rui Tan. 2018. Deep room recognition using inaudible echos. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* (2018).
- [49] Xingzhe Song, Boyuan Yang, Ge Yang, Ruirong Chen, Erick Forno, Wei Chen, and Wei Gao. 2020. SpiroSonic: monitoring human lung function via acoustic sensing on commodity smartphones. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*.
- [50] Takafumi Taketomi, Hideaki Uchiyama, and Sei Ikeda. 2017. Visual SLAM algorithms: A survey from 2010 to 2016. *IPSN Transactions on Computer Vision and Applications* 9, 1 (2017), 1–11.
- [51] Stephen P Tarzia, Peter A Dinda, Robert P Dick, and Gokhan Memik. 2011. Indoor localization without infrastructure using the acoustic background spectrum. In *Proceedings of the 9th international conference on Mobile systems, applications, and services*. 155–168.

- [52] Yu-Chih Tung and Kang G Shin. 2015. Echotag: Accurate infrastructure-free indoor location tagging with smartphones. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*.
- [53] A Van Wieringen. 2014. Undesirable effects as a result of short-term exposure to an ultrasonic repellent device. Part II—Exposure of volunteers. *Assignment from the Federal Public Service in Belgium for Health, Food Chain Safety and Environment* (2014).
- [54] Sen Wang, Hongkai Wen, Ronald Clark, and Niki Trigoni. 2016. Keyframe based large-scale indoor localisation using geomagnetic field and motion pattern. In *2016 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 1910–1917.
- [55] Xiangyu Xu, Jiadi Yu, Yingying Chen, Yanmin Zhu, Linghe Kong, and Minglu Li. 2019. Breathlistener: Fine-grained breathing monitoring in driving environments utilizing acoustic signals. In *Proceedings of the 17th Annual International Conference on Mobile Systems, Applications, and Services*.
- [56] Zheng Yang, Chenshu Wu, and Yunhao Liu. 2012. Locating in fingerprint space: wireless indoor localization with little human intervention. In *Proceedings of the 18th annual international conference on Mobile computing and networking*. 269–280.
- [57] Moustafa Youssef and Ashok Agrawala. 2005. The Horus WLAN location determination system. In *Proceedings of the 3rd international conference on Mobile systems, applications, and services*. 205–218.
- [58] Sangki Yun, Yi-Chao Chen, Huihuang Zheng, Lili Qiu, and Wenguang Mao. 2017. Strata: Fine-grained acoustic-based device-free tracking. In *Proceedings of the 15th annual international conference on mobile systems, applications, and services*.
- [59] Chi Zhang and Xinyu Zhang. 2016. LiTell: Robust indoor localization using unmodified light fixtures. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*. 230–242.
- [60] Xiwu Zhang, Lei Wang, and Yan Su. 2021. Visual place recognition: A survey from deep learning perspective. *Pattern Recognition* 113 (2021), 107760.
- [61] Zengbin Zhang, David Chu, Xiaomeng Chen, and Thomas Moscibroda. 2012. Swordfight: Enabling a new class of phone-to-phone action games on commodity phones. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*.
- [62] Bing Zhou, Mohammed Elbadry, Ruipeng Gao, and Fan Ye. 2017. BatMapper: Acoustic sensing based indoor floor plan construction using smartphones. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*. 42–55.