IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received 5 February 2024, accepted 2 March 2024, date of publication 11 March 2024, date of current version 20 March 2024.

Digital Object Identifier 10.1109/ACCESS.2024.3375392

RESEARCH ARTICLE

LiquidListener: Supporting Ubiquitous Liquid Volume Sensing via Singing Sounds

TAEWON KIM[1], SEOHYEON PARK[1], SANGEUN OH[2], AND HYOSU KIM[1], (Member, IEEE)
[1]School of Computer Science and Engineering, Chung-Ang University, Seoul 06974, South Korea
[2]Department of Software and Computer Engineering, Ajou University, Seoul 16499, South Korea

Corresponding author: Hyosu Kim (hskimhello@cau.ac.kr)

This work was supported in part by the Chung-Ang University Research Grants in 2022 and in part by the Bio and Medical Technology Development Program of the National Research Foundation (NRF) funded by the Korean Government (MSIT) under Grant RS-2023-00222910.

This work involved human subjects or animals in its research. The authors confirm that all human/animal subject research procedures and protocols are exempt from review board approval.

ABSTRACT This work proposes LiquidListener, a novel liquid volume sensing method for containers. Specifically, it enables the ubiquitous measurement of liquid volume not available in existing work due to *i*) dependencies on dedicated sensing hardware (e.g., capacity sensors) and containers (e.g., transparent containers) and *ii*) a high training intensity. A key enabler of LiquidListener is listening to *singing sounds*. When a user taps a container using solid objects, such as pens and teaspoons, the container vibrates freely and produces a singing sound. As the container is filled with more liquid, the pitch of the sound decreases. Based on this relationship, we develop acoustic-based liquid volume sensing algorithms that support the precise measurement of liquid volume while using only a smartphone and requiring minimal user effort for calibration. The extensive experiments demonstrate that LiquidListener can support high accuracy with an average error ratio of 2.3% in sensing the liquid volume in various containers. In addition, the experimental results indicate that it can still maintain a similar level of accuracy in diverse and dynamically changing environments, even without additional calibration.

INDEX TERMS Acoustic sensing, mobile healthcare services, smarthome applications, ubiquitous liquid volume sensing.

I. INTRODUCTION

Water is indispensable in our lives and is deeply related to human activities, such as drinking, cooking, and cultivation. Therefore, there has been an increasing attention to liquid volume sensing techniques [1], [2], [3], [4], [5], that can benefit our daily lives. These techniques support in situ measurements of liquid volume in containers and eventually facilitate the development of many valuable applications for smart home environments and mobile healthcare. For example, they can help track and manage the daily intake of water or medicine [6], [7] and assist in monitoring and refilling the inventory of various liquids [8]. In addition, it would be possible to cook food anywhere (e.g., even in outdoor environments) without needing to prepare measuring cups.

The associate editor coordinating the review of this manuscript and approving it for publication was Mohammad Zia Ur Rahman.

However, none of the existing techniques can fully support ubiquitous liquid volume sensing due to their limitations in deployability and usability. First, many of them [4], [9], [10], [11], [12], [13], [14], [15], [16], [17], [18], [19] rely on dedicated sensing hardware, such as capacitive sensors, pressure sensors, ultrasonic sensors, electrodes and radio-frequency identification tags. Others [20], [21] have leveraged cameras but support liquid volume measurement only for transparent containers filled with opaque liquids. Moreover, one of the most recent studies, LiquidSense [1], works with commercial Wi-Fi access points for liquid-level sensing. However, it still requires installing additional hardware (e.g., transducers) in a container and suffers from low usability due to its high training intensity.

This work introduces LiquidListener, a novel method that enables in situ liquid volume measurements with high accuracy, deployability, and usability. The key enabler of

© 2024 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

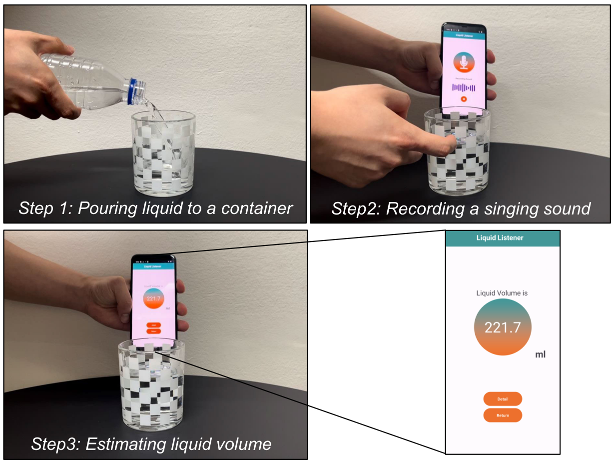

FIGURE 1. **Usage scenario of LiquidListener, measuring the amount of liquid in a container using only a smartphone.**

LiquidListener is to leverage *singing sounds*. When an external force is applied to a container (e.g., when tapping a cup), the force causes the container to deform, vibrate, and produce a singing sound. The pitch of the sound is determined by the physical properties of the container and the amount of liquid in it. For example, previous field studies [22], [23] have observed that a singing sound is emitted at a lower frequency as the liquid level in the container increases. This phenomenon occurs because the added liquid participates in the vibration of the container and eventually changes the overall mass of the oscillating body.

Fig. 1 illustrates how LiquidListener measures the amount of liquid in containers using singing sounds. Once a user fills a cup with liquid, LiquidListener first asks the user to tap the cup with any solid object (e.g., teaspoons, pens, and fingernails). Then, it collects singing sounds using the microphone in the user's smartphone and finally estimates the liquid volume by analyzing the pitch of the sound. LiquidListener has a considerable advantage over existing work in terms of deployability because it can turn any container into a measuring tool using only a commercial smartphone (equipped with a microphone) and a tapping tool.

However, several challenges exist to realize the precise and usable liquid volume measurement with singing sounds. First, the pitch of the singing sounds is also affected by the container's physical properties, such as size, thickness, and material. Therefore, LiquidListener should perform a precalibration to calculate the container-dependent parameters to avoid losing usability. Further, different containers have different physical characteristics, requiring individual calibration for each container. Second, users might use LiquidListener in diverse environments, especially with dynamically changing noise. This noise is often mixed with singing sounds in audio recordings, causing errors in liquid volume sensing.

We address these challenges by exploring the relationship of singing sounds with liquid volumes and using these observations in designing LiquidListener as follows:

- We design a simple two-point calibration method that precisely estimates container-dependent parameters with little user effort. Specifically, based on the phenomenon that the pitch of singing sounds is inversely proportional to the liquid volume in a container, we compute the parameters using data collected only when a container is empty or full. Moreover, once a container is calibrated, there is no need to perform further calibration for the container because the parameters are virtually constant regardless of other environmental factors, such as liquid types and tapping tools.
- We precisely estimate the pitch of a singing sound even from noisy audio recordings. To this end, taking a cue from the observation of the time and frequency characteristics of singing sounds, we design two noise cancellation schemes: irrelevant frequency filtering and early part removal. In addition, we further improve accuracy in pitch estimation through a high-resolution frequency analysis.
- We derive a new formula to calculate the amount of liquid in a container using the estimated pitch of the singing sound and the calibrated container-dependent parameters.

We evaluate the performance of LiquidListener via extensive experiments with the prototype implementation running on commercial Android smartphones. The evaluation results demonstrate that LiquidListener can accurately measure the volume of liquids in containers without compromising deployability and usability. For example, an average error ratio of 2.3% is achieved on various containers with a simple calibration, which takes less than 1 min, using only commercial smartphones for recording and everyday objects for tapping. In addition, LiquidListener maintains high accuracy in diverse situations and dynamically changing environments without conducting additional calibrations.

In conclusion, the contributions of this work are as follows:

- To the best of our knowledge, this is the first attempt to leverage singing sounds to enable ubiquitous liquid volume sensing for containers.
- We design LiquidListener, a novel acoustic-based liquid volume sensing method that achieves a high degree of accuracy with the use of a smartphone and tapping tools, requiring minimal effort for calibration.
- We implement a prototype of LiquidListener[1] and verify its effectiveness through extensive experiments in diverse real-world scenarios.

II. RELATED WORK

A. LIQUID VOLUME SENSING FOR THE INDUSTRY

In the field of industrial applications, several studies have been proposed for sensing the volume of liquids contained in large tanks. One representative method is employing capacitive sensors attached to the inside of a container [9], [10], [11]. The liquid volume can be estimated by measuring

[1] See https://youtu.be/vwanwsEre-I for our demo video.

the capacitance of the entire container via capacitive sensors because the capacitance varies depending on the amount of liquid in the container. However, this method is unsuitable for everyday containers because it requires the capacitive sensors to be immersed directly in liquid, resulting in the pollution of edible liquid. Another method is to estimate the liquid volume by measuring the pressure exerted by the liquid at the bottom of the container [12], [13]. However, this method can only be used when the container has a flat and wide enough bottom to attach pressure sensors; thus, it is not possible to apply this method to containers with various bottom shapes. Furthermore, capacitive-based and pressure-based methods have low deployability because they require sensors dedicated to the industrial field.

B. LIQUID VOLUME SENSING FOR DAILY CONTAINERS

Many studies have been conducted to estimate liquid volume in containers. This studies can be divided into the following four categories:

- *Vision-based*: Some studies have analyzed the RGB images of containers taken with off-the-shelf cameras to estimate their liquid volume. Playful Bottle [20] offers a simple method of capturing a transparent bottle with scale marks through a smartphone camera and calculating the liquid volume from the captured image. Another study [21] presented a deep-learning technique that can estimate the liquid level of a container using convolutional neural networks (CNNs) trained using the images of containers with liquid. However, these methods have low deployability because they can only be used for transparent containers with opaque liquids, and their performance is greatly affected by ambient lighting environments. In addition, the CNN approach also has low usability because it requires numerous of training data to construct the CNN model.
- *Vibration-based*: Another approach is to employ vibrations generated in a container. The VibeBin [24] technique senses the content level of a waste bin by measuring the resonant vibration caused due to the impact of a vibrating motor on the container surface. Similarly, Ryu et . [25] proposed a smaller actuator and sensor made of a special material called polyvinylidene fluoride for measuring the resonant frequency of vibration. However, these approaches have low deployability because they require dedicated hardware to generate and detect vibrations.
- *Wireless signal-based*: Several methods aim to sense the liquid level by propagating wireless signals to a container. LiveTag [17] estimates the liquid level by attaching a thin metal tag to the surface of a container and measuring the attenuation of the Wi-Fi signals reflected by the tag. However, this method has the constraint that it cannot be used in metal containers that completely block the transmission of Wi-Fi signals. Nakagawa et al. [4] proposed a contactless method to measure the liquid level based on the absorption of the propagated millimeter waves in the liquid. However, this method is unsuitable for household settings because it requires special hardware (i.e., a millimeter Doppler sensor and piezoelectric vibrator). LiquidSense [1] measures the resonance frequency of the container using Wi-Fi signals and a transducer attached to the container. Then, it uses the resonance frequency as a feature to construct a support vector machine model to predict the liquid level. However, this system has limited usability because it requires considerable training data and cannot be used outdoors without a Wi-Fi access point.
- *Acoustic-based*: Similarly to LiquidListener, many studies have been proposed to measure liquid volumes using acoustic sensing. However, they use different types of acoustic sounds differently, rather than the singing sound of the container. Some studies [14], [15] have measured the liquid volume using ultrasonic waves for higher accuracy. However, these methods require dedicated hardware, such as Time-of-Flight sensors, a multiple-input, multiple-output transducer array, and ultrasonic sensors. In addition, PursingNet [16] has a sensing system that can predict the liquid volume via a deep neural network trained with sounds collected when liquid is poured into the container. Nevertheless, this system has low usability because it can only predict the volume when additional liquid is poured into the container, and it requires a high training intensity.

C. LIQUID IDENTIFICATION TECHNIQUES

Many studies have recently been proposed to perform liquid identification for containers. Fundamentally, they aim to identify what type of liquid is contained in a container or whether it contains other impurities through various sensors. Some studies [26], [27], [28], [29], [30] have presented vision-based approaches. These approaches extract the unique features of each liquid by monitoring the physical phenomena, such as surface tension, sloshing motion, air bubbles of liquids, and Brownian motion. Another class of work [31], [32] has supported precise contactless liquid identification by leveraging radio-frequency signals. Other [33], [34], [35] have suggested methods to identify the liquid type by measuring the acoustic impedance from acoustic signals passing through the liquid or the vibration resistance force induced by the viscosity of the liquid. However, whether these physical features are effective and practical indicators for measuring liquid volumes has not been thoroughly investigated.

For example, Vi-Liquid [35] infers the liquid volume by measuring its resonance frequency induced by a smartphone vibration during the identification process, but it has several limitations in practice. First, this method requires a smartphone to be attached to a specific position on the container for high accuracy; thus, the liquid volume can only be measured in containers specially designed to hold the smartphone (low

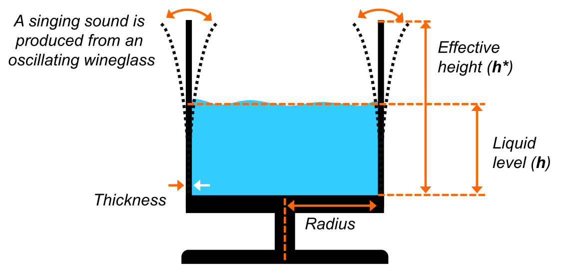

FIGURE 2. Generation of a singing sound from a wineglass. The pitch of the singing sound is determined by the physical properties of the liquid (e.g., its level and density) and the container (e.g., its density, thickness, radius, and height).

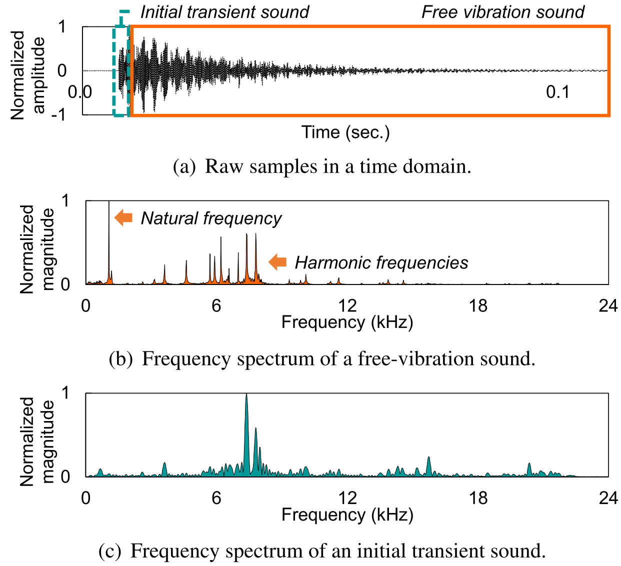

FIGURE 3. Typical structure of a singing sound measured with an empty glass.

deployability). Second, this method must construct a database by collecting a significant amount of frequency information for various liquid volumes for measurement (low usability).

D. SUMMARY

The existing studies in the above categories have limitations in either deployability, usability or both. In contrast, LiquidListener can support in situ liquid volume measurements with higher deployability and usability by *i*) capturing the singing sound of a container using a commercial smartphone microphone, *ii*) not contacting the liquid directly, and *iii*) requiring a low calibration effort.

III. PRELIMINARIES

This section describes the preliminary knowledge on which LiquidListener is based. Then, it explores the feasibility of leveraging this knowledge for liquid volume sensing.

A. WINEGLASS ACOUSTICS

We consider a glass harp, a musical instrument comprising wine glasses. Each glass in the harp is filled with a different amount of liquid and produces a singing sound with a different pitch when a wet finger runs around the rim of the glass.

Such a relationship between the pitch of a singing sound and the liquid volume in a wineglass has been intensively explored in previous studies [22], [23]. A straightforward analysis is to model a singing glass as a spring-mass system. When a wineglass is struck or rubbed by a moistened finger, it vibrates freely, emitting a singing sound. The frequency of the free vibration can be calculated as follows:

$$f = \frac{1}{2\pi}\sqrt{\frac{k}{m}} \tag{1}$$

where m is the overall mass of the system and k is the spring constant, which depends on the physical properties of an oscillating body (i.e., the glass itself). Therefore, adding liquid to the glass does not affect k but does affect m. The overall mass increases with the liquid, causing a decrease in the frequency f and the pitch of the emitted singing sound.

French [23] conducted a deeper analysis of the singing glass effect with consideration of damping. Fig. 2 presents a cylindrical wineglass with a height of h^*. French derived a general formula for how the pitch of a singing sound could vary with the amount of liquid in the cylindrical wineglass as follows:

$$\left(\frac{f_0}{f}\right)^2 \approx 1 + \rho_l a \left(\frac{h}{h^*}\right)^4 \tag{2}$$

where h denotes the current liquid level in the glass, f_0 represents the pitch measured when the glass is empty, ρ_l is the density of the liquid, and a indicates a container-dependent parameter determined by the density, thickness, and radius of the glass. Chen et al. [22] further generalized (2) with an experimental analysis as follows:

$$\left(\frac{f_0}{f}\right)^2 \approx 1 + \rho_l a \left(\frac{h}{h^*}\right)^{n+1} \tag{3}$$

where n is approximately 2. This formula indicates that $(f_0/f)^2$ is linearly proportional to $(h/h^*)^{n+1}$. As the liquid level (h) in the glass increases, the frequency f decreases.

B. FEASIBILITY OBSERVATIONS

Based on the preliminary knowledge, we investigated the characteristics of singing sounds in diverse real-world environments. To this end, we conducted the following experiments. A single user was asked to fill a certain amount of water in containers, such as glasses, ceramic mugs, and stainless cups. Then, the user tapped them with a stainless steel teaspoon in a quiet classroom (30 – 40 dBA). During that time, we collected audio data using the built-in microphone of a Google Pixel 4 smartphone with a sampling rate of 48 kHz. Note that commercial smartphones are equipped with microphones that can capture audio data in an audible frequency range of 0 to approximately 20 kHz [36]. Therefore,

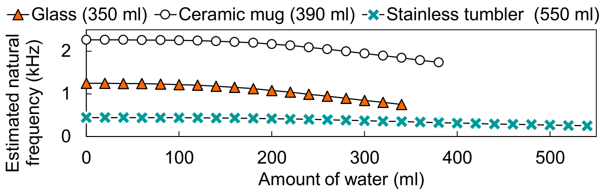

FIGURE 4. Variations in the pitch of a singing sound (i.e., the natural frequency of a container) over the amount of water in the container.

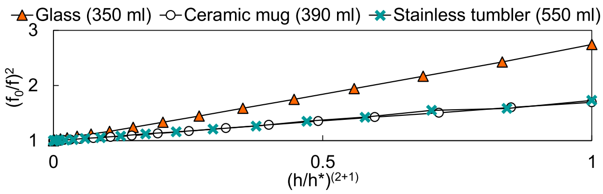

FIGURE 5. Linear relationship between $(f_0/f)^2$ and $(h/h^*)^{n+1}$, where n is set to 2, in diverse containers.

according to the Nyquist theorem [37], the 48 kHz sampling rate selection allows us to observe the characteristics of singing sounds in all the possible measurable frequencies.

Fig. 3(a) illustrates how a singing sound changes over time. When a container is impacted, it deforms due to the applied force. This sudden deformation causes the generation of an initial transient sound of a few milliseconds from the contact area. The force propagates to the entire body, and the container freely vibrates at its natural frequency for several tens of milliseconds, generating a free-vibration sound. Thus, as depicted in Fig. 3(b), the free-vibration sound has distinctive peak values at the natural and harmonic frequencies of the container. The initial transient sound has different frequency characteristics than the free-vibration sound (Fig 3(c)) because the initial transient sound's characteristics primarily depend on the applied external force signal.

As explained in Section III-A, singing sounds produced from the same container can have different characteristics depending on the amount of liquid in the container. Specifically, the pitch decreases as the amount of liquid in the container increases (Fig. 4). This trend can also be observed with nonglass and noncylindrical containers (i.e., any container). In addition, Fig. 5 demonstrates that $(f_0/f)^2$ has a linear relationship with $(h/h^*)^{n+1}$ as predicted in (3). The slope between the two values differs depending on the container because each container has nonidentical physical properties (e.g., size, shape, and material). Such a linear relationship can be an excellent hint for measuring liquid volumes in containers.

IV. OVERVIEW OF LIQUIDLISTENER

LiquidListener aims to support an accurate, deployable, and usable liquid volume measurement for containers. This goal can be achieved by leveraging the preliminary observations on singing sounds in designing LiquidListener, as illustrated in Fig. 6.

- *Noise-robust pitch detection.* LiquidListener collects audio data using a microphone, detects the singing sound from the collected recordings, and estimates its pitch through a frequency analysis. To improve robustness against noise, LiquidListener uses two methods: *i*) irrelevant frequency filtering and *ii*) early part removal. The estimated pitch is used for calibration (in the calibration phase) or liquid volume measurement (in the measurement phase).
- *Simple two-point calibration.* In the calibration phase, LiquidListener computes container-dependent parameters for liquid volume estimation. It attempts to minimize the time and effort for the calibration (i.e., to maintain a high degree of usability) by *i*) asking users to collect singing sounds only when the container is empty and full and *ii*) requiring no additional calibration for previously calibrated containers.
- *Liquid volume calculation.* LiquidListener calculates the liquid volume in the container using the estimated pitch and calibrated parameters. Specifically, this calculation is based on the formula for the liquid volume derived from (3).

Regarding deployment conditions, LiquidListener assumes the following two deployment conditions: *i*) a tapping tool and a microphone must be available to produce a singing sound from a container and record the sound, respectively, and *ii*) some properties of both the container and liquid (e.g., the container capacity and liquid density) are given in advance.

The first assumption can be very easily satisfied. Users can make noticeable singing sounds by leveraging everyday solid objects (e.g., pens, chopsticks, teaspoons, and even fingernails) as a tapping tool. In addition, most commercial smartphones are already equipped with microphones that can be used for recording. For the second condition, users can use the information about the containers and liquids provided by vendors or previous field studies. For example, most manufacturers offer detailed information about their containers, including their capacity. In addition, the density of common liquids, such as water, milk, and coke, is already well studied. To further improve deployability, we can construct a database of such already-known information and provide it to the users. Thus, LiquidListener can be deployed in everyday environments. Such information might be erroneous or unavailable in the worst case. We discuss its effect on the performance of LiquidListener in Sections VII-D and VIII.

V. LIQUIDLISTENER DESIGN

This section explains how LiquidListener enables ubiquitous liquid volume sensing with high accuracy and usability in diverse real-world environments.

A. NOISE-ROBUST PITCH DETECTION

Given audio recordings, LiquidListener attempts to detect a singing sound and estimate its pitch. However, precise

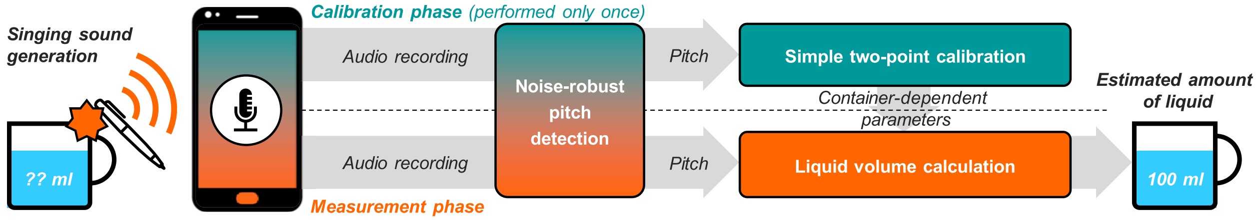

FIGURE 6. Overall architecture of LiquidListener.

TABLE I. Specification of diverse daily containers.

	Type	Cylindrical	Handle	Vol. (mL)	f^L (Hz)	f^H (Hz)
G1	Glass	O	X	320	1407	2011
G2	Glass	X	O	360	722	1249
G3	Glass	O	O	450	684	1079
G4	Glass	O	X	1150	433	796
C1	Ceramic mug	O	O	390	1716	2265
C2	Ceramic teacup	X	O	230	2038	2535
C3	Ceramic bowl	X	X	320	1140	1539
C4	Ceramic bowl	X	X	1000	970	1406
S1	Stainless tumbler	O	X	410	475	1020
S2	Stainless tumbler	X	X	550	248	442
S3	Stainless cup	O	O	440	294	517
S4	Stainless cup	O	O	340	325	543
S5	Stainless cup	X	X	170	412	645
ST	Stoneware cup	X	O	360	1694	2181
P1	Plastic cup	X	X	540	184	405
P2	Plastic cup	X	O	480	305	542

Note: f^L and f^H for each container were measured when the container is full and empty, respectively.

pitch estimation is challenging due to noise, which can be classified into three categories. The first two types are ambient noise from the surrounding environment and the sound produced by the handles and support for containers. When a container is struck, these parts vibrate at a different frequency than the container's main body due to differences in the physical shape, size, and thickness. The last type is an initial transient sound generated from a container's contact area. As explained in Section III-B, the initial sound has frequency characteristics that are less relevant to the natural frequency of a container, disturbing the accurate measurement of the pitch of a singing sound.

LiquidListener improves the measurement's robustness against these noise types using the following methods: *i) irrelevant frequency filtering* and *ii) early part removal*.

1) IRRELEVANT FREQUENCY FILTERING

LiquidListener first cancels out noisy frequencies using a bandpass filter between f^L and f^H Hz. The required cutoff frequencies might vary depending on the container (Table I). This variation is primarily because of differences in the mass, density, and stiffness of the containers. For example, stainless containers (with high density) and plastic containers (with low stiffness) have a lower natural frequency (e.g., $<$ 1 kHz) than others. Based on these empirical observations, we initially set f^L and f^H to 0.1 and 1 kHz

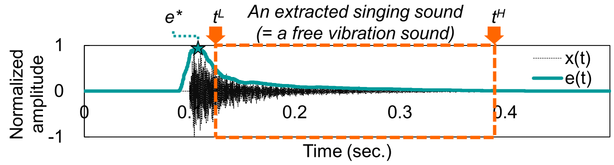

FIGURE 7. Singing sound segmentation, discarding the early part of the singing sound.

(for stainless and plastic containers) or 0.4 and 2.5 kHz (for others), respectively. Next, LiquidListener conducts calibration using these initial values and determines the cutoff frequencies optimized for each container (for more details, see Section V-C). Once f^L and f^H are determined, it constructs a Butterworth band pass filter with an order of 4 and denoises audio recordings by applying the filter.

2) SINGING SOUND SEGMENTATION WITH EARLY PART REMOVAL

A singing sound is extracted from the filtered recordings. LiquidListener computes $e(t)$, an accumulated sound energy level at time t, as follows:

$$e(t) = \sum_{i=t}^{t+l} x(i)^2 \tag{4}$$

where $x(t)$ is the filtered audio signal at time t and l, a window size, is empirically set to 1 ms. LiquidListener then takes audio samples between $[t^L, t^H]$. We determine the time instants t^L and t^H as the last time instant at which $e(t)$ is larger than $e^* \times 0.5$ and $e^* \times 0.01$, respectively, where e^* is the maximum level among $e(t)$. In other words, the early part of the singing sound (i.e., the initial transient sound) is excluded, and only the latter part (the free-vibration sound) is extracted for the pitch estimation (Fig. 7). This helps increase accuracy in pitch estimation.

3) FINE-GRAINED PITCH ESTIMATION

LiquidListener estimates the pitch of the extracted singing sound via a frequency analysis. In particular, to improve accuracy in pitch estimation, we analyze the frequency characteristics of the sound in a fine-grained manner, especially using a high-resolution fast Fourier transform (FFT). LiquidListener first applies a Hamming window to the extracted signal, extends the signal's length to 1 s by padding

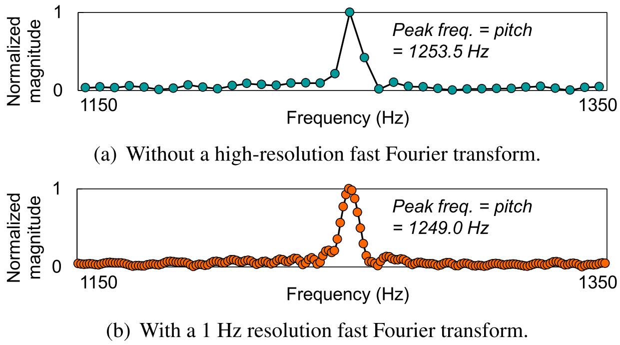

FIGURE 8. Accuracy improvement in pitch estimation using high-resolution frequency analysis.

zeros, and takes the FFT outputs from the extended signal. This simple zero padding method enables LiquidListener to detect the pitch in a frequency resolution of 1 Hz. Finally, the pitch of the singing sound, denoted as $\hat{f}$, is estimated as the frequency at which the sound has the maximum magnitude value (Fig. 8).

The pitch estimation accuracy can be further improved by asking a user to tap a container several times and using the set of singing sounds for the estimation. For example, suppose that a container is struck n^T times. LiquidListener first detects n^T singing sounds and extracts an FFT spectrum for each sound. Then, it combines them by computing their average FFT spectrum and measures $\hat{f}$. This average spectrum reduces the effect of outliers, allowing more precise pitch estimation. Increasing n^T might cause a decrease in usability. However, the experiments reveal that LiquidListener can provide high accuracy even when n^T is 2 or 3, that is, not compromising usability much (Section VII-C).

B. LIQUID VOLUME CALCULATION

Once the pitch is estimated, LiquidListener calculates the amount of liquid in a container based on the relationship observed in (3). Specifically, we derived the approximate formula for a liquid volume v as follows:

$$v \approx v^* \frac{h}{h^*} \tag{5}$$

where v^* is the container's maximum capacity. This approximation might be reasonable only for cylindrical containers. However, through the experimental analysis, we observed that even for non-cylindrical containers, LiquidListener can still support an accurate estimate of liquid volume under this approximation (for more details, See Section VII-B). LiquidListener obtains the estimated liquid volume, denoted as $\hat{v}$, with the following equation derived from (3) and (5):

$$\hat{v} = v^* \left(\frac{\left(f_0/\hat{f}\right)^2 - 1}{\rho_l^M a} \right)^{-1/(n^C+1)} \tag{6}$$

where container-dependent parameters, such as f_0 and a, are determined in the calibration phase, and n^C is set to 2 based on empirical observations in the previous research [22]. As discussed in Section IV, we assumed that v^* and ρ_l^M (the density of the liquid used for measurement) are given in advance.

C. SIMPLE TWO-POINT CALIBRATION

The frequency characteristics of singing sounds are affected by physical properties and the liquid amount, implying that container-dependent parameters (e.g., a and f_0 in (6)) must be estimated in advance by performing user-involved calibration. To minimize a user's calibration effort, we leverage the relationship observed in Section III: As a container is filled with more liquid, the singing sound is produced at a lower frequency. This relationship means a singing sound has the highest pitch f^H (or lowest pitch f^L) when a container is empty (or full). Based on this, LiquidListener conducts a simple two-point calibration for a given container. First, it collects singing sounds when the container is empty and full and then estimates f^H and f^L from the sounds using the proposed pitch estimation method. Last, the container-dependent constant a in (6) is computed as follows:

$$a = \frac{(f_0/f^L)^2 - 1}{\rho_l^B} \tag{7}$$

where f_0 equals f^H and ρ_l^B, denotes the density of the liquid used for calibration and is given in advance. These parameters are kept for future use (i.e., the measurement phase). Human error might occur during calibration. For example, users might pour more or less liquid than a container's known capacity when measuring f^L. However, we empirically observed that these errors are less than 10 mL in usual cases and have a negligible effect on the accuracy of liquid volume measurement. For example, in the presence of these human errors, LiquidListener achieves high accuracy, with an average error ratio of 2.3% (Section VII-B).

As mentioned in Section V-B, n^C in (6) is empirically set to 2 by default. However, it can be further optimized during the calibration phase if singing sounds are collected with a partially filled container. For example, if we collect n^S singing sounds for a specific container, LiquidListener finds a and n^C that minimize the following mean squared error (ϵ):

$$\epsilon = \frac{1}{n^S} \sum_{i=1}^{n^S} (v_i - \hat{v}_i)^2 \tag{8}$$

where $\hat{v}_i$ denotes the liquid volume estimated with the ith singing sound using (6) and v_i represents the ground truth.

VI. IMPLEMENTATION

We implemented a prototype of LiquidListener as a mobile application running on commercial Android smartphones. It collects audio data with a sampling rate of 48 kHz and the `UNPROCESSED` audio source to avoid using vendor-specific preprocessing techniques. The application runs the proposed liquid volume sensing algorithms as follows:

FIGURE 9. Containers used in the experiments. Table I provides their detailed specifications.

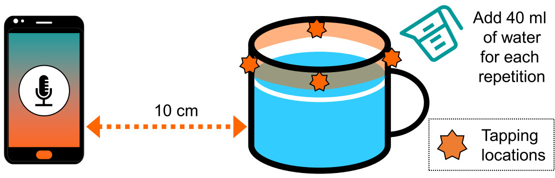

FIGURE 10. Default evaluation setup.

- In the calibration phase, the application initially asks the user to enter the liquid type (to obtain its density from a database) and the container's name and capacity. The application computes container-dependent parameters as proposed in Section V-C and stores the user-type and calibrated parameters for future use. This calibration process is requested only once for a specific container. In the measurement phase, the user selects the liquid and the container. In particular, the container is selected among those whose parameters have been precalibrated.
- The application loads its calibrated parameters and estimates the liquid volume in the container by analyzing singing sounds.

VII. EVALUATION

We evaluated the performance of LiquidListener in terms of accuracy, deployability, usability, and robustness by answering the following questions:

- How accurately does LiquidListener estimate the amount of liquid in a container regardless of the *i*) container shape and material, *ii*) liquid type, *iii*) tapping tool, and *iv*) smartphone model?
- How much effort should a user expend to measure liquid volume precisely?
- How robustly does LiquidListener work in diverse environments and situations (e.g., in the presence of noise)?

A. EVALUATION SETUP AND METHODOLOGY

We conducted experiments with 16 containers, each of which has various shapes, sizes, volumes, and materials (Fig. 9 and Table I for detailed specifications). During the experiments, we placed the containers on a wooden table in a quiet classroom (30 – 40 dBA), tapped their upper side using a stainless teaspoon, and collected singing sounds using the built-in microphone of a Google Pixel 4 smartphone placed 10 cm away from the containers as demonstrated in Fig. 10. We repeated this process several times while varying the amount of liquid in the containers. For each repetition, we increased the containers' liquid volume by adding 40 mL of water (with a density of 0.997 g/mL) and tapped each of the four locations 10 times. We used a pipette and kitchen scale to minimize human error in pouring a certain amount of water into the containers. Finally, for the ith container, we collected $40 \times n^R$ singing sounds, where n^R denotes the number of repetitions and equals $\lceil \frac{v_i^*}{40ml} \rceil + 1$, and v_i^* indicates the capacity of the ith container.

Using the collected data, we verified the performance of LiquidListener using the following steps. First, we assumed that a user taps a container multiple times for more precise pitch estimation as explained in Section V-A. Under this assumption, we randomly selected three singing sounds for each volume (i.e., $3 \times n^R$ sounds were selected). Among them, the data collected when the container was empty and full were used for calibration, and the rest were used for measuring the accuracy of LiquidListener. This process was repeated 1000 times for each container individually.

It should be noted that we selected our default experimental conditions (e.g., tapping location, tool, and distance) as described above. However, these conditions can play a critical role in determining the the performance of LiquidListener. So, we also conducted additional experiments under diverse environments (e.g., with various tapping tools).

We used an *error ratio* as the primary metric to evaluate the performance of LiquidListener. This metric is frequently employed in other studies related to liquid volume sensing [1], [14], [21]. The error ratio is defined as follows:

$$Error\ ratio = \frac{|Estimated\ volume - Ground\ truth|}{Total\ capacity} \quad (9)$$

where the ground truth is measured using a kitchen scale.

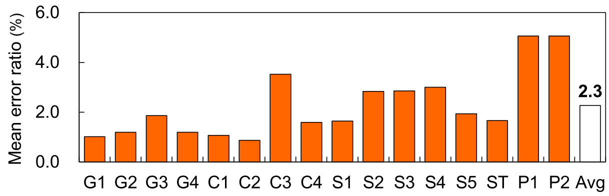

FIGURE 11. Overall accuracy of LiquidListener with diverse containers.

B. OVERALL PERFORMANCE OF LIQUIDLISTENER

Fig. 11 demonstrates that LiquidListener supports precisely measuring liquid volumes in any container (e.g., 2.3% on average). This high accuracy is primarily due to the capability of LiquidListener to precisely estimate the pitch of singing sounds using a fine-grained frequency analysis (i.e., a 1 Hz resolution FFT). As illustrated in Fig. 12, the measurement error decreases with a higher-resolution FFT. For example,

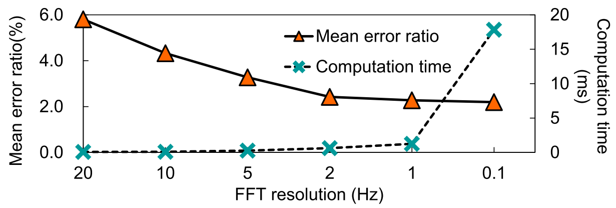

FIGURE 12. Effect of using a high-resolution fast Fourier transform(FFT) in precisely estimating a liquid volume.

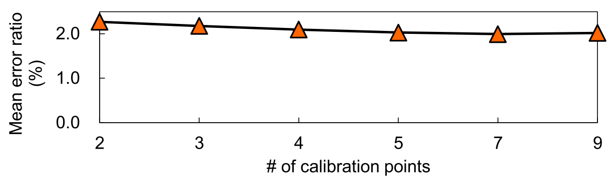

FIGURE 13. Liquid volume sensing errors with various calibration efforts.

an increase in the resolution from 20 to 1 Hz incurs a significant drop in error ratio by approximately 60%, not compromising responsiveness much (a computation time of $<$1.5 ms). The accuracy can slightly increase further using a more fine-grained analysis, requiring high computational overhead (17.8 ms with a 0.1 Hz resolution FFT).

Notably, plastic containers (P1 and P2) have higher error ratios than other containers because the plastic material usually has a higher damping ratio than others, resulting in weak vibrations induced by tapping. This phenomenon makes it challenging for LiquidListener to estimate the pitch of singing sounds accurately. Nevertheless, LiquidListener exhibits a satisfactory level of accuracy for plastic containers, with an error ratio of 4.9% or less.

C. USABILITY TEST

LiquidListener requires users to collect singing sounds by tapping a container. However, this data collection might burden users and decrease usability. Therefore, in this experiment, we verify how much user effort is required to measure liquid volumes.

1) NUMBER OF CALIBRATION POINTS

As explained in Section V-C, LiquidListener conducts calibration to compute container-dependent parameters. More specifically, it asks a user to collect a set of calibration data when a container is empty and full. With this simple two-point calibration, LiquidListener optimizes the container-dependent constant a according to (7) and achieves high measurement accuracy (e.g., an error ratio of 2.3% on average). Concerning this, some may argue that additional calibration data for a partially filled container could further improve accuracy. However, as presented in Fig. 13, increasing the number of calibration points reduces the error ratio slightly, but the improvement is not significant. Thus, the constant value of a can be sufficiently optimized with only two-point calibration. Therefore, LiquidListener requires minimal calibration effort (i.e., only two calibration points) for liquid volume sensing, minimizing the effect on usability. Through a field study with real-world users, we observed that it takes less than 1 min to collect data for the two-point calibration.

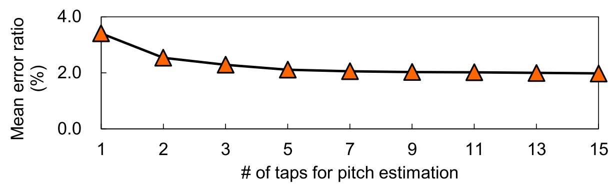

FIGURE 14. Effect of the number of data collection.

2) NUMBER OF TAPS FOR PITCH ESTIMATION

LiquidListener can require a user to collect singing sounds several times for a particular volume, reducing the effect of outliers in pitch estimation as described in Section V-A, and increasing the measurement accuracy of LiquidListener. For example, Fig. 14 demonstrates that the error ratio gradually decreases up to 2.0% as the number of singing sounds (n^T) used for pitch estimation increases. When n^T equals 3, LiquidListener achieves nearly optimal accuracy (an average error ratio of 2.3%) within only a few seconds (less than 2 s) of data collection. Based on this experimental observation, we set n^T to 3 by default, which prevents LiquidListener from losing much usability by requiring only a small number of taps by users.

D. DEPLOYABILITY TEST

The previous sections reveal that LiquidListener could work well with everyday containers with little calibration effort. We further verified the deployability of LiquidListener in more diverse configurations (e.g., with various liquids, tapping tools, and smartphone). Specifically, to observe how well LiquidListener can be deployed in these environments without compromising usability, we compared the following two methods of measuring liquid volume for each configuration:

- `Extra-Cal` measures the liquid volume after additional calibration for the target configuration. The additional calibration takes about 15 s, similar to that required for a normal calibration process.
- `No-Cal` measures the volume without additional calibration, employing the data previously calibrated with the default configuration (i.e., water, a teaspoon, and Pixel4).

In addition, during this experiment, we only used a specific container (G1) to focus solely on the effects of other factors.

1) EFFECT OF LIQUID TYPE

First, we evaluated the accuracy of LiquidListener in measuring liquid volumes for seven types of liquids: coke, milk, vodka, corn syrup, soybean oil, detergent, and ethanol 80%. Fig. 15 reveals that `Extra-Cal` can measure liquid

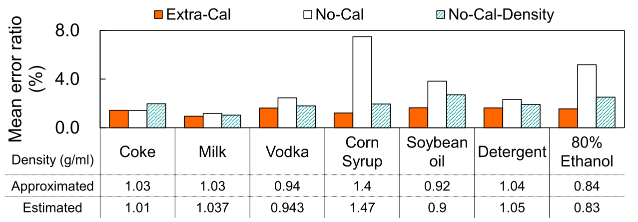

FIGURE 15. Liquid volume sensing errors for various liquid conditions. The approximate density of each liquid was obtained from previous field studies [38], [39], [40], [41], [42], [43]. The ground truth (i.e., the estimated density) was measured using a scale and measuring cup.

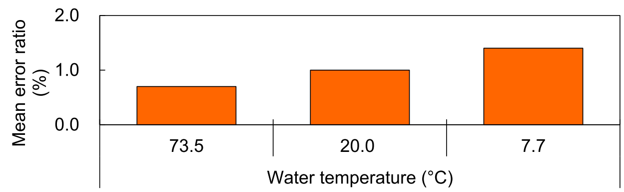

FIGURE 16. Influence of water temperature. (e.g., hot, room temperature, cold).

volume with high accuracy (an average error ratio of 1.4%) for all seven liquids. However, this method has low usability because it requires recalibration whenever the liquid type changes. In contrast, `No-Cal`, which leverages calibration parameters obtained with water, can achieve an accuracy as high as that for `Extra-Cal` only for liquids with densities (0.94 to 1.05 g/mL) similar to water, even without additional calibration (an average error ratio of 1.9%). Nevertheless, its accuracy significantly declines with some liquids, such as corn syrup, soybean oil, and 80% ethanol, because they have significant density differences from water. We can mitigate this problem using the approximate density value of the target liquid (i.e., ρ_l^M in (6)). According to the experiment, this method (denoted as `No-Cal-Density`) supports all seven liquids with an average error ratio of 1.7%, which is not significantly different from the accuracy of `Extra-Cal`. The approximate densities of various liquids can be easily found in many open materials, and they are not significantly different from the densities measured in their experiments, as illustrated in Fig. 17. LiquidListener manages the approximate densities as a database to support a wide range of liquid types.

2) EFFECT OF WATER TEMPERATURE

We evaluated the accuracy of LiquidListener under varying water temperatures, including hot, room temperature, and cold conditions, as illustrated in Fig. 16. The results indicate that the system accuracy is higher in hot than cold water. This phenomenon in hot water can be attributed to the decreased density of water at higher temperatures, which leads to an increased speed of water molecules. Consequently, the speed of sound is enhanced, allowing sound waves to propagate more efficiently. This results in a clearer pitch, which LiquidListener detects with precision. It's notable that the error rate for cold water is 1.4%, demonstrating minimal difference from that of hot water, indicating that the system is less affected by water temperature.

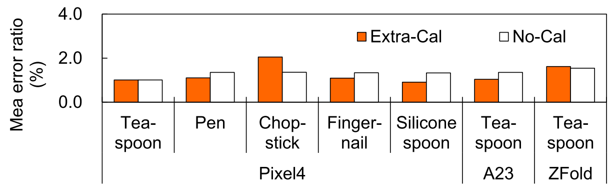

FIGURE 17. Influence of tapping tools and devices.

3) EFFECT OF TAPPING TOOLS AND RECORDING DEVICES

Different tapping tools have non-identical structures, shapes, sizes, and materials. Depending on the smartphone model, microphone response also varies. These differences can cause variations in the characteristics of singing sounds. For example, leveraging a soft object (e.g., the tip of a fingernail) as a tapping tool results in a lower amplitude of singing sound than using more rigid tools (e.g., teaspoons).

To investigate such effects, we measured the amount of water using LiquidListener under various configurations consisting of five tapping tools and three smartphone models. Fig. 17 illustrates that the `Extra-Cal` and `No-Cal` methods achieve low error ratios of less than 2.1% for all configurations. In particular, `No-Cal` has an accuracy value similar to `Extra-Cal` without additional calibration, despite using the default calibration data obtained with the teaspoon and Pixel4. This result is because the significant factors determining the pitch of singing sounds are the physical properties of containers and liquids, as mentioned in Section III. Thus, LiquidListener can capture these core characteristics from received sounds using the proposed algorithms and finally supports high accuracy regardless of the tapping tool and smartphone without requiring recalibration.

E. ROBUSTNESS TEST

Users attempt to measure liquid volumes ubiquitously in diverse and dynamically changing environments, implying that singing sounds could be collected from differing environments at each phase (i.e., the calibration and measurement phases). Therefore, we evaluated how robustly LiquidListener works against such environmental changes. More specifically, during this experiment, we conducted calibration with the dataset collected in ideal environments as described in Section VII-A and made changes in the measurement environment.

1) AGAINST AMBIENT NOISE

The accuracy of acoustic-based algorithms can be degraded significantly due to ambient noise. To observe the influence of this environmental factor, we measured the accuracy of LiquidListener in four different places with varying

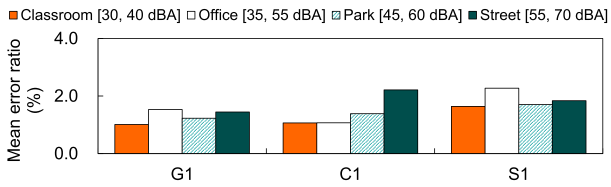

FIGURE 18. Robustness of LiquidListener in diverse places with various noise characteristics.

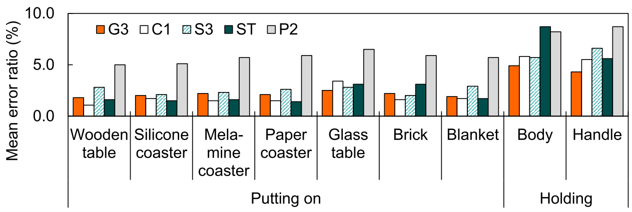

FIGURE 19. Influence of the container placement. Five containers (G3, C1, S3, ST, P2) that have a handle were used to observe the effect of diverse holding types.

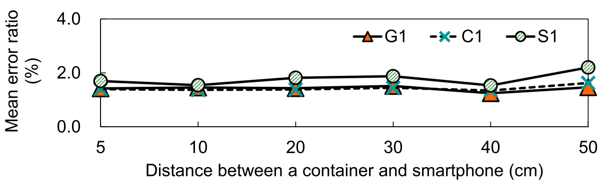

FIGURE 20. Robustness of LiquidListener against changes in measuring distance.

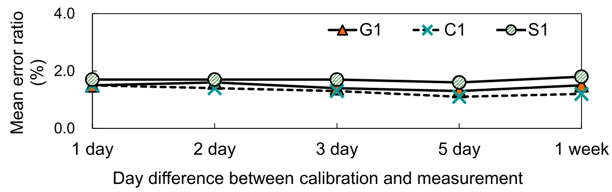

FIGURE 21. Time-invariant characteristics of LiquidListener.

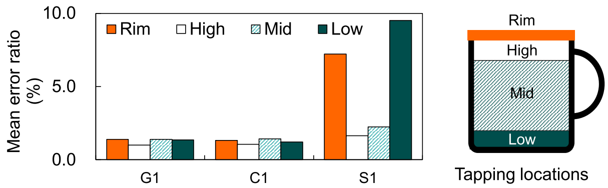

FIGURE 22. Influence of tapping location variations.

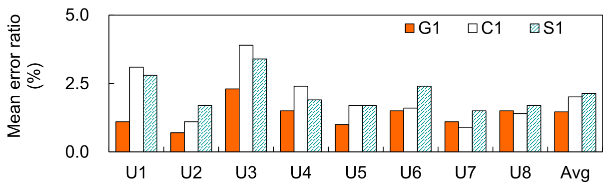

FIGURE 23. Robustness of LiquidListener against user differences.

noise levels: a classroom, office, park, and street (near a construction site). As illustrated in Fig. 18, the error ratios in all containers are within 2.3%, demonstrating that LiquidListener is not significantly affected by the surrounding noise. This outcome is because the proximity between a smartphone and a container (about 10 cm) enables it to capture singing sounds at a louder intensity than the noise. Furthermore, LiquidListener can operate robustly against noise because it filters unnecessary frequency bands containing noise based on prior observations of singing sound characteristics.

2) AGAINST CHANGES IN THE PLACEMENT OF A CONTAINER

In the measurement phase, users could position containers differently than when performing calibration. For example, a user may put a container on a surface of a different material, hold a container by hand, or place a container away from or close to a smartphone. Therefore, these experiments evaluate how much changes in container placement affect the performance of LiquidListener.

Fig. 19 presents the error ratios of LiquidListener measured when five containers are placed on surfaces made of seven materials and when a user is holding the container body or handle. According to the experimental results, LiquidListener provides similar accuracy (e.g., with a standard deviation of 0.3% error ratio for G3) regardless of the surface on which the containers are placed. This result is because the surfaces have little influence on the vibrations of the container body. In other words, singing sounds produced from the oscillating body have similar characteristics regardless of where a container is placed. In contrast, the measurement error ratios drastically increase up to 8.3% when a user holds the container body or handle. The main reason is that the vibrations of the container body are absorbed by the user's hand, generating low-level singing sounds.

Fig. 20 depicts the error ratios of LiquidListener measured while varying the distance between the container and smartphone. We confirmed that the error ratios remain within 2.2% even when the distance is up to 50 cm apart because the intensity of singing sounds is strong enough. If the distance between a container and smartphone extends further than 50 cm, the accuracy of LiquidListener may decrease. However, considering the typical user habit of using a smartphone up close, placing a smartphone more than 50 cm away from a container would be extremely rare.

3) AGAINST TEMPORAL VARIATIONS

Fig. 21 illustrates that LiquidListener maintains high accuracy even when there are temporal differences exist between calibration and measurement. More specifically, this method consistently achieves error ratios of less than 1.7% for all containers. As the pitch of the singing sound is determined by the physical properties of the container and liquid, LiquidListener can provide stable performance over time.

4) AGAINST BEHAVIORAL VARIATIONS

Users might have different behavioral characteristics (e.g., tapping intensities and locations) when using LiquidListener. Worse, the behavior can vary for the same user depending on the situation. We verified how this behavioral variation affects the performance of LiquidListener with the following two experiments. First, we evaluated the accuracy of measuring the liquid volume over different tapping locations as depicted in Fig. 22. We used data collected when tapping the upper side of a container, denoted as `High`, for calibration. The figure illustrates that LiquidListener maintains high accuracy regardless of the tapping locations for G1 and C1. However, the error ratio increases to 7.2% and 9.5% when tapping on the rim and bottom side of S1, respectively. The tumbler (S1) has a unique structure, with a threaded finish and thick support at its top and bottom, respectively (Fig. 9). This generates noisy or weak singing sounds, decreasing the performance of measuring liquid volume. However, except for these cases, LiquidListener can still provide only subtle errors, even for S1, at an average error ratio of 1.9%.

Second, we observe the influence of the behavioral difference between users on measuring liquid volume with eight real-world users recruited from the university. During this experiment, each participant was asked to freely fill water in a container and measure the volume using LiquidListener. More specifically, each of them tapped the container's upper side five times using a stainless teaspoon. This measurement process was repeated ten times for each container (G1, C1, and S1). Then, we measured the accuracy using the data collected from one user for measurement and the others for calibration. As illustrated in Fig. 23, LiquidListener provides a high level of accuracy in measuring liquid volumes (e.g., an average error ratio of 1.7%, 2.3%, and 2.2% for each type of container). Some users (e.g., U3) had significantly different tapping styles from others. However, only a subtle increase in error was observed, indicating that tapping styles have a minimal influence on the measurement.

Through the above experiments, we confirmed that LiquidListener works robustly in most cases. No additional calibration for a specific container is required even in environmental changes. However, a few exceptions exist. One possible solution to avoid such exceptional situations would be to provide users with simple guidelines. For example, we can ask them to tap the container's upper or middle part while not holding it.

VIII. DISCUSSION

A. USE OF HEAVY TAPPING TOOLS

Although LiquidListener requires lightweight tapping tools, such as teaspoons or pens, some may prefer to use heavier objects. For example, a user can directly use a smartphone, which records singing sounds, as a tapping tool, as in Knocker [44], a knock sound-based system that supports object recognition with a smartphone. This usage makes LiquidListener more convenient because the user does not need to prepare an additional tool for tapping. However, in the experiments where the subjects hit containers with their smartphones, we observed that the containers were likely to be pushed back or damaged easily, resulting in difficulties in recording and making singing sounds. As future work on this matter, we consider using a smartphone's vibration motor to generate singing sounds and detect liquid volumes with the smartphone.

B. TYPES OF UNMEASURABLE CONTAINERS

We observed that LiquidListener has difficulty measuring liquid volume for three types of containers. The first type is paper containers. Paper material inherently has a much higher damping ratio, which is an essential condition for vibration. Thus, it cannot generate vibrations caused by external forces, leading to no production of singing sounds. Therefore, the system is unable to measure liquid volume in paper containers.

The second type is thermos bottles that feature a double-walled internal structure. The space between the double walls is maintained in a vacuum state to minimize thermal conductivity. Thus, even if the outer wall of the container is struck, the applied force experiences a significant attenuation while propagating the inner wall, making it impossible for LiquidListener to measure the liquid volume inside thermos bottles.

The third type is a container with a closed lid. The lid can block the propagation of singing sounds, significantly attenuating their amplitude. Therefore, the proposed system fails to sense the liquid level. However, besides these cases, we confirmed through intensive experiments that LiquidListener has high deployability, supporting containers comprising various materials, such as glass, ceramic, stainless steel, and even plastic.

C. PRIOR INFORMATION ON CONTAINER CAPACITY

As mentioned, LiquidListener requires the capacity information of a container to calculate liquid volume, which vendors commonly provide through product manuals. The proposed system manages this information as a database, supporting a wide range of containers. However, if the capacity is not given, the proposed system can ask a user to measure the container's height and radius using a ruler or a measuring tool app on the smartphone and calculate the capacity approximately. The error ratio for sensing liquid volume in cylindrical containers with the approximate capacity is within 2.7% on average, which is negligible. In contrast, for noncylindrical containers such as wine glasses, the error ratio increases up to 10.4% because users cannot accurately estimate the capacity due to their complex shape. Thus, if users attempt to measure liquid volume in noncylindrical containers with an unknown capacity, LiquidListener notifies them that the measurement may be inaccurate.

D. PERFORMANCE IN LARGE CONTAINERS

LiquidListener primarily targets cup-like containers used for drinking water or other beverages. The experiment found that LiquidListener could also measure liquid volumes with a similar error ratio of about 2% to 3%, even in larger containers (e.g., pots). This outcome is because the system calculates the liquid volume by considering the ratio of the current liquid level to the total height of the container, according to (5). However, as the container's maximum capacity v^* in (5) increases, the absolute error value might inevitably increase. Thus, the accuracy of liquid volume sensing must be further improved to use LiquidListener practically in large containers. We leave this problem as future work.

IX. CONCLUSION

This study presents LiquidListener, a novel sensing method that enables the ubiquitous measurement of liquid volume in containers while providing a high level of accuracy, deployability, and usability. The key idea of LiquidListener is to predict the liquid volume by capturing and analyzing the singing sound of a container. To realize this method, we devised effective acoustic-based sensing algorithms to detect the pitch of the singing sound robustly despite environmental noise. The method performs calibration with little effort and precisely calculates the liquid volume of the container. The extensive experiments demonstrate that LiquidListener can detect liquid volume in containers of various materials and shapes with high accuracy, even under diverse situations and dynamically changing environments. Currently, LiquidListener is limited to containers without lids; however, extending its capability to measure water levels in sealed containers leaves a goal for future work. This advancement would ensure drinking water safety and management across various applications. We believe that LiquidListener can serve as a promising assistant for ubiquitous liquid volume sensing in daily life with many practical application scenarios, such as in healthcare systems for tracking water intake.

REFERENCES

[1] Y. Ren, S. Tan, L. Zhang, Z. Wang, Z. Wang, and J. Yang, "Liquid level sensing using commodity WiFi in a smart home environment," *Proc. ACM Interact., Mobile, Wearable Ubiquitous Technol.*, vol. 4, no. 1, pp. 1–30, Mar. 2020.

[2] M. Fan and K. N. Truong, "SoQr: Sonically quantifying the content level inside containers," in *Proc. ACM Int. Joint Conf. Pervasive Ubiquitous Comput.*, Sep. 2015, pp. 3–14.

[3] S. Mukherjee, "Non-invasive measurement of liquid content inside a small vial," in *Proc. IEEE Radio Wireless Symp. (RWS)*, Jan. 2010, pp. 527–530.

[4] T. Nakagawa, A. Hyodo, K. Osada, H. Kurata, and S. Oho, "Contactless liquid-level measurement through opaque container using millimeter-wave sensor," in *Proc. IEEE SENSORS*, Oct. 2011, pp. 1421–1424.

[5] P. Mohindru, "Development of liquid level measurement technology: A review," *Flow Meas. Instrum.*, vol. 89, Mar. 2023, Art. no. 102295.

[6] E. Slawek. (2017). *I Stayed Hydrated for Two Weeks and it Changed My Life*. [Online]. Available: https://www.nbcnews.com/better/diet-fitness/i-stayed-hydrated-two-weeks-it-changed-my-life-n731131

[7] S. K. Schwartz. (2012). *How to Help Your Patients Manage Their Medications*. [Online]. Available: https://www.physicianspractice.com/view/how-help-your-patients-manage-their-medications

[8] S. Purkayastha. (2016). *Iot Application That Monitors Your Food Pantry*. [Online]. Available: https://www.ibm.com/cloud/blog/iot-inventory-monitor-part1

[9] S. C. Bera, H. Mandal, S. Saha, and A. Dutta, "Study of a modified capacitance-type level transducer for any type of liquid," *IEEE Trans. Instrum. Meas.*, vol. 63, no. 3, pp. 641–649, Mar. 2014.

[10] C. A. Gong, H. K. Chiu, L. R. Huang, C. H. Lin, Z. D. Hsu, and P.-H. Tu, "Low-cost comb-electrode capacitive sensing device for liquid-level measurement," *IEEE Sensors J.*, vol. 16, no. 9, pp. 2896–2897, May 2016.

[11] K. Loizou, E. Koutroulis, D. Zalikas, and G. Liontas, "A low-cost capacitive sensor for water level monitoring in large-scale storage tanks," in *Proc. IEEE Int. Conf. Ind. Technol. (ICIT)*, Mar. 2015, pp. 1416–1421.

[12] K.-W. Park and H. C. Kim, "High accuracy pressure type liquid level measurement system capable of measuring density," in *Proc. IEEE Region 10 Conf.*, Nov. 2015, pp. 1–5.

[13] P. Esmaili, F. Cavedo, and M. Norgia, "Differential pressure based liquid level measurement in sloshing condition," in *Proc. IEEE Int. Instrum. Meas. Technol. Conf. (I2MTC)*, May 2018, pp. 1–6.

[14] P. Li, Y. Cai, X. Shen, S. Nabuzaale, J. Yin, and J. Li, "An accurate detection for dynamic liquid level based on MIMO ultrasonic transducer array," *IEEE Trans. Instrum. Meas.*, vol. 64, no. 3, pp. 582–595, Mar. 2015.

[15] S. L. Mohammed, A. Al-Naji, M. M. Farjo, and J. Chahl, "Highly accurate water level measurement system using a microcontroller and an ultrasonic sensor," *IOP Conf. Ser., Mater. Sci. Eng.*, vol. 518, no. 4, May 2019, Art. no. 042025.

[16] H. Liang, S. Li, X. Ma, N. Hendrich, T. Gerkmann, F. Sun, and J. Zhang, "Making sense of audio vibration for liquid height estimation in robotic pouring," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Nov. 2019, pp. 5333–5339.

[17] C. Gao, Y. Li, and X. Zhang, "LiveTag: Sensing human-object interaction through passive chipless Wi-Fi tags," *GetMobile: Mobile Comput. Commun.*, vol. 22, no. 3, pp. 32–35, Jan. 2019.

[18] T.-Y. Wu, L. Tan, Y. Zhang, T. Seyed, and X.-D. Yang, "Capacitivo: Contact-based object recognition on interactive fabrics using capacitive sensing," in *Proc. 33rd Annu. ACM Symp. User Interface Softw. Technol.*, Oct. 2020, pp. 649–661.

[19] W. Sun, Y. Chen, S. Zhan, T. Han, F. Tian, H. Wang, and X.-D. Yang, "RElectrode: A reconfigurable electrode for multi-purpose sensing based on microfluidics," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, May 2021, pp. 1–12.

[20] M.-C. Chiu, S.-P. Chang, Y.-C. Chang, H.-H. Chu, C. C.-H. Chen, F.-H. Hsiao, and J.-C. Ko, "Playful bottle: A mobile social persuasion system to motivate healthy water intake," in *Proc. 11th Int. Conf. Ubiquitous Comput.*, Sep. 2009, pp. 185–194.

[21] Y. Jiang, E. Schenck, S. Kranz, S. Banerjee, and N. K. Banerjee, "CNN-based non-contact detection of food level in bottles from RGB images," in *Proc. 25th Int. Conf. MMM*, Thessaloniki, Greece. Cham, Switzerland: Springer, 2019, pp. 202–213.

[22] K.-W. Chen, C.-K. Wang, C.-L. Lu, and Y.-Y. Chen, "Variations on a theme by a singing wineglass," *Europhys. Lett. (EPL)*, vol. 70, no. 3, pp. 334–340, May 2005.

[23] A. P. French, "In vino veritas: A study of wineglass acoustics," *Amer. J. Phys.*, vol. 51, no. 8, pp. 688–694, Aug. 1983.

[24] Y. Zhao, S. Yao, S. Li, S. Hu, H. Shao, and T. F. Abdelzaher, "VibeBin: A vibration-based waste bin level detection system," *Proc. ACM Interact., Mobile, Wearable Ubiquitous Technol.*, vol. 1, no. 3, pp. 1–22, Sep. 2017.

[25] J. Ryu, H. Jeong, Y. Chen, C. Oh, J. Kim, H. Kim, S. Cho, K. No, Y.-H. Park, S. Park, and S. Hong, "Flexible piezoelectric liquid volume sensor," *Sens. Actuators A, Phys.*, vol. 276, pp. 219–225, Jun. 2018.

[26] S. Yue and D. Katabi, "Liquid testing with your smartphone," in *Proc. 17th Annu. Int. Conf. Mobile Syst., Appl., Services*, Jun. 2019, pp. 275–286.

[27] B. Sun, S. R. X. Tan, Z. Ren, M. C. Chan, and J. Han, "Detecting counterfeit liquid food products in a sealed bottle using a smartphone camera," in *Proc. 20th Annu. Int. Conf. Mobile Syst., Appl. Services*, Jun. 2022, pp. 42–55.

[28] J. Chan, A. Raghunath, K. E. Michaelsen, and S. Gollakota, "Testing a drop of liquid using smartphone LiDAR," *Proc. ACM Interact., Mobile, Wearable Ubiquitous Technol.*, vol. 6, no. 1, pp. 1–27, Mar. 2022.

[29] Y. Huang, K. Chen, L. Wang, Y. Dong, Q. Huang, and K. Wu, "Lili: Liquor quality monitoring based on light signals," in *Proc. 27th Annu. Int. Conf. Mobile Comput. Netw.*, Oct. 2021, pp. 256–268.

[30] K. An, Q. Zhang, and E. Kwong, "ViscoCam: Smartphone-based drink viscosity control assistant for dysphagia patients," *Proc. ACM Interact., Mobile, Wearable Ubiquitous Technol.*, vol. 5, no. 1, pp. 1–25, Mar. 2021.

[31] F. Shang, P. Yang, Y. Yan, and X.-Y. Li, "PackquID: In-packet liquid identification using RF signals," *Proc. ACM Interact., Mobile, Wearable Ubiquitous Technol.*, vol. 6, no. 4, pp. 1–27, Dec. 2022.
[32] Y. Liang, A. Zhou, H. Zhang, X. Wen, and H. Ma, "FG-LiquID: A contactless fine-grained liquid identifier by pushing the limits of millimeter-wave sensing," *Proc. ACM Interact., Mobile, Wearable Ubiquitous Technol.*, vol. 5, no. 3, pp. 1–27, Sep. 2021.
[33] X. Sun and C. Feng, "Ls-liquid: Towards container-irrelevant liquid sensing on smartphones," in *Proc. SIGCOMM Poster Demo Sessions*, Aug. 2022, pp. 25–27.
[34] Y. Yang, Y. Wang, J. Cao, and J. Chen, "HearLiquid: Nonintrusive liquid fraud detection using commodity acoustic devices," *IEEE Internet Things J.*, vol. 9, no. 15, pp. 13582–13597, Aug. 2022.
[35] Y. Huang, K. Chen, Y. Huang, L. Wang, and K. Wu, "Vi-liquid: Unknown liquid identification with your smartphone vibration," in *Proc. 27th Annu. Int. Conf. Mobile Comput. Netw.*, Sep. 2021, pp. 174–187.
[36] V. Hromadová, P. Kasák, R. Jarina, and P. Brída, "Frequency response of smartphones at the upper limit of the audible range," in *Proc. ELEKTRO (ELEKTRO)*, May 2022, pp. 1–5.
[37] H. J. Landau, "Sampling, data transmission, and the Nyquist rate," *Proc. IEEE*, vol. 55, no. 10, pp. 1701–1706, Oct. 1967.
[38] R. S. Herrick, L. P. Nestor, and D. A. Benedetto, "Using data pooling to measure the density of sodas: An introductory discovery experiment," *J. Chem. Educ.*, vol. 76, no. 10, p. 1411, Oct. 1999.
[39] H. M. Musallam, H. M. Almozogai, S. S. Amkabis, M. A. Aoag, T. M. Hassan, E. A. Elhefian, and F. M. Asseid, "Physicochemical characteristics of various milk samples," *Nova J. Med. Biol. Sci.*, vol. 6, no. 2, pp. 1–3, 2017.
[40] N. Hu, D. Wu, K. Cross, S. Burikov, T. Dolenko, S. Patsaeva, and D. W. Schaefer, "Structurability: A collective measure of the structural differences in vodkas," *J. Agricult. Food Chem.*, vol. 58, no. 12, pp. 7394–7401, Jun. 2010.
[41] J. L. Maxwell, F. A. Kurtz, and B. J. Strelka, "Specific volume (density) of saccharide solutions (corn syrups and blends) and partial specific volumes of saccharide-water mixtures," *J. Agricult. Food Chem.*, vol. 32, no. 5, pp. 974–979, 1984.
[42] T. A. Scott, "Refractive index of ethanol–water mixtures and density and refractive index of ethanol-water-ethyl ether mixtures," *J. Phys. Chem.*, vol. 50, no. 5, pp. 406–412, May 1946.
[43] P. Rice and W. Hamm, "Densities of soybean oil/solvent mixtures," *J. Amer. Oil Chemists' Soc.*, vol. 65, no. 7, pp. 1177–1179, Jul. 1988.
[44] T. Gong, H. Cho, B. Lee, and S.-J. Lee, "Knocker: Vibroacoustic-based object recognition with smartphones," *Proc. ACM Interact., Mobile, Wearable Ubiquitous Technol.*, vol. 3, no. 3, pp. 1–21, Sep. 2019.

TAEWON KIM received the B.S. degree from the Department of Computer Software Engineering, Wonkwang University, in 2022. He is currently pursuing the M.S. degree with the School of Computer Science and Engineering, Chung-Ang University, Seoul, South Korea. His research interests include the IoT, cyber-physical systems, and mobile and ubiquitous systems.

SEOHYEON PARK received the B.S. degree from the Department of Computer Science and Engineering, Chung-Ang University, Seoul, South Korea, in 2023, where she is currently pursuing the M.S. degree with the School of Computer Science and Engineering. Her research interests include cyber-physical systems, signal processing, and mobile and ubiquitous systems.

SANGEUN OH received the B.S. degree in computer and communication engineering from Korea University, Seoul, South Korea, in 2012, and the M.S. and Ph.D. degrees in computer science from Korean Advanced Institute of Science and Technology (KAIST), Daejeon, South Korea, in 2014 and 2020, respectively. He is currently an Assistant Professor with the Department of Software and Computer Engineering, Ajou University, Suwon, South Korea. His research interests include mobile/the IoT systems and real-time embedded systems.

HYOSU KIM (Member, IEEE) received the B.S. degree from the Department of Computer Engineering, Sungkyunkwan University, in 2010, and the M.S. and Ph.D. degrees from the School of Computing, Korean Advanced Institute of Science and Technology (KAIST), Republic of Korea, in 2012 and 2018, respectively. He is currently an Assistant Professor with the School of Computer Science and Engineering, Chung-Ang University, Seoul, South Korea. His research interests include cyber-physical systems, mobile and ubiquitous systems, and privacy.