SoQr: Sonically Quantifying the Content Level inside Containers

Mingming Fan, Khai N. Truong
Xemxa, Inc.
Toronto, ON Canada
{mingming, khai}@xemxa.com

ABSTRACT
In this paper, we present SoQr, a sensor that can be attached to an external surface of a household item to estimate the amount of content inside it. The sensor consists of a speaker and a microphone. It outputs a short duration sine wave probing sound to excite a container and its content, and then records the container's impulse response. SoQr then extracts Mean Mel-Frequency Cepstral Coefficients from impulse response recordings of a container with different content levels and learns a support vector machine classifier. Results from a 10-fold cross validation of the prediction models on 19 common household items demonstrate that SoQr can correctly estimate the content level for these products with an average overall F-Measure above 0.96. We then further evaluated SoQr's robustness in different usage scenarios to gain an understanding of how the system performs and specific challenges that might arise when users interact with these products and the sensor.

Author Keywords
Content level measurement; container; active probing; impulse response

ACM Classification Keywords
H.5.5. Sound and music computing: Signal analysis, synthesis, and processing.

INTRODUCTION
Sensing the amount of content in a container is an important ubiquitous computing research challenge with a wide variety of practical applications. It can provide an understanding of when an item has been used, how much content has been consumed and whether a product needs to be replenished. Such knowledge, for instance, can help patients manage their medication compliance or remind care-givers when they may need to refill patients' medication bottles. On the other hand, householders can leverage such knowledge to better plan what and when to buy so as to reduce the amount of shopping trips which might help reduce their shopping stresses [16] and potential food wastes by increasing the visibility of the amount of existing foods and not overbuying [17].

Yet, determining the remaining content level in a container remains a challenge for application developers. Many existing approaches may be impractical or expensive. For example, some previous solutions need to directly make contact with the measured liquid [3, 20, 29, 30, 33, 34] and most often are designed for large tanks. The ones that work for portable containers either require users to pay extra efforts in installation and calibration [9,11] or only work in specially designed containers [10,35]. It is not uncommon that many are only able to distinguish very few content levels, for instance, above or below certain point (*e.g.*, [13]), and the high precision solutions may require extra components to wrap [25] or place [26] around the target container.

In this paper, we investigate how to keep track of the content level in portable containers that are typically found in homes regardless of the physical form of the content. We have designed a sensor which can be attached to the products at a single point and uses an active acoustic probing method to determine how much of these products (or content level) exists inside of each item's original container/packaging. The sensor, called SonicQuantifier or SoQr for short (pronounced as "soccer"), emits a probing sound and then records and analyzes the container's impulse response (IR). SoQr uses supervised learning methods to develop a model of the different content level for each product in their containers based on a corresponding set features extracted from the recorded impulse response.

Our evaluation of the system shows that SoQr can be externally installed on the surface of a container to quantify how much content exists inside it. We demonstrate the efficacy of SoQr through tests with 19 common household items in their original packaging. The results of our tests show that SoQr achieves an average precision, recall, and F-Measure result of 0.970, 0.969 and 0.969 respectively with these 19 products. Through preliminary tests which examine the robustness of SoQr to different ways that a user might potentially interact with household products and the sensor, we show that a prediction model built at one location can work reasonably well at another. We demonstrate that the

UbiComp '15, September 7-11, 2015, Osaka, Japan.
© 2015 ACM. 978-1-4503-3574-4/15/09…$15.00.
http://dx.doi.org/10.1145/2750858.2804264

prediction model can be trained in two common storage orientations and used to predict a product's content level in either of those orientations. We also learn that SoQr is not able to perform well with deformable containers because its shape changes each time the user closes it by clipping or tying and mechanisms are needed to help the user install the sensor at a specific position on a product matching where the position where a data for a prediction model was collected. These findings suggest that SoQr has the potential to accurately estimate the content level of common household items but mechanisms to facilitate the training of these prediction models in different context and to help the user install the sensor must be explored.

RELATED WORK

In this section, we provide a review of different sensing methods for estimating the content level in a container. We then discuss in depth how the Impulse Response technique that we adopt in SoQr for analyzing the container, has been used in literature.

Methods for Estimating the Content Level in a Container

Capacitive Sensing

Capacitive sensing is a method that has been widely explored as a mechanism for measuring the amount of liquid in a container [3, 13, 18, 20, 29, 30, 33, 34]. Previous systems have leveraged the fact that the liquid inside a container is a dielectric material that affects the overall capacitance of the container; thus, when the liquid level changes in a container, the capacitance of the container changes accordingly. Many of these previous systems [3, 20, 29, 30, 33, 34] typically require the capacitive sensors to be in direct contact with the liquid being measured. In contrast, Goekler [18] and Dietz *et al.* [13] have proposed capacitive sensor designs that can be externally added to a container. While Goekler's sensor was designed to infer whether the current liquid level is above or below a pre-defined point, Dietz *et al.*'s sensor design [13] is capable of performing continuous liquid level measurement but requires that the sensor wraps around the entire container. To a lesser extent, capacitive sensing has also been used for measuring solid items. For example, AdhereTech [1] has developed a special purpose pill bottle to measure the number of pills inside it.

Although previous work has shown that capacitive sensing can be used to measure the content level of different types of materials inside a container, the physical requirements of this method (*i.e.,* the need to place the sensor in contact with the content, wrap completely around the container, or replace the container itself with a specially designed one) can limit its use in measuring many common household items. In contrast, we explore the design of a small sensor that can be easily attached to an external surface of the container/packaging for common household products.

Load Sensing

A different way of measuring a product's amount is to use load cells. Load cells are typically used in digital scales to weigh objects. A system can track a product's weight to determine how the amount of that item changes over time. Chi *et al.* [5] previously augmented a kitchen counter and a stove with load cells to measure the amount of food being cooked as a part of the system for estimating the total calories in the food that is being cooked. Lo *et al.* [21] created a playful tray to help engage a child in her eating behavior. The tray used a load cell to monitor the weight change of the food being placed on its top. The expensive cost of high precision load cells may be a barrier to attaching this sensor onto individual household products in order to track their content level. Alternatively, when they are used to augment entire surfaces to effectively act as a large scale, such systems may not be able to detect when specific items are being used if multiple products have been placed on it. Less expensive load cells may not be able to detect changes in the content level of very light items.

Camera Sensing

Cameras can also be used to determine how much content is inside a container. For example, Playful Bottle [6] is a camera-based system which used the rear camera of a mobile phone attached to a transparent bottle with pattern bars painted on its outer surface to detect the level of water in a bottle. However, this method only works on a transparent container filled with non-opaque liquid. In this work, we aim to develop a sensor that can estimate the content level of different household products in their original packaging, regardless of whether that is transparent or not.

Electromagnetic Wave based Sensing

By sending a radio wave from the bottom of a liquid container and analyzing the phase angle change in the received reflected wave, Mukherjee [25] presented a method of sensing the liquid level in a container. However, the method required to wrap two pieces of electrodes around the target container and was only able to detect the liquid level within the height of the electrodes. By leveraging the fact that the absorption coefficients of millimeter wave are much higher for liquid than for air, Nakagawa *et al.* [26] used a millimeter Doppler sensor and a piezoelectric vibrator to measure the liquid level in an opaque container. Unfortunately, this method required to place the two components on two opposite sides of the target container. In contrast, SoQr only needs to contact with the target container at a single point yet is able to sense the content level beyond its positon.

Acoustic Sensing

Another widely explored mechanism for measuring the content level in a container is acoustic sensing. Time-of-flight is the most prevalent acoustic sensing technique adopted [2, 11, 19, 27, 32]. The time-of-flight method measures the distance between two points by determining the time elapsed between emitting an acoustic signal and receiving it. This concept has been leveraged previously in both contact and contactless designs. In contact designs, an ultrasonic emitter and a receiver are immersed into a container of liquid to measure its level [2, 19]. In contactless methods, a sensor unit is installed on the outside surface of a container. Different placements of the acoustic sensor unit have been explored. The acoustic sensor can be hung in air above the surface of the liquid inside an open container [32], and used to emit a sonic or ultrasonic signal towards the liquid surface and measure its time-of-flight. Olson and Christensen [27] proposed a method that installed the acoustic sensor unit at the bottom of a container. The sensor emits towards the container an ultrasonic signal which would propagate through the container wall into the liquid inside, and continue until it eventually reaches the surface of the liquid and then bounces back and is captured by the sensor. The time elapsed was used to estimate the liquid level. In contrast, Dam and Austerlitz [11] proposed a sensor design that can be added to the outside of a container along a vertical wall. Their sensor consists of an ultrasound transmitter and an ultrasound receiver, which need to be placed at two different vertical positions on the container. Time of flight between these two parts is used to estimate the internal liquid level between where the transmitter and the receiver units are placed. Dam and Austerlitz's design—which has been implemented as a sensor for measuring the liquid level in plastic containers [9]—requires that the sensor extend the vertical length of the measured container. If built as a single part, it would have a length constraint—making it difficult to reuse the sensor on any container. In contrast, the SoQr sensor unit is designed to be a small part that can be placed at one position of a container to estimate the content level for the entire container.

A major limitation with the time-of-flight method is that a millisecond error in measurement will result in ~34 centimeter error in distance measurement (assuming the speed of sound is 340 meter/second). Given the relatively small sizes of containers/packaging of household items, it would be quite challenging to use time-of-flight as the measuring technique. Use of the time-of-flight method typically requires costly, high precision instruments to accurately measure the time of travel. In this paper, we propose the use of active probing as the means of exciting the container acoustically and then analyzing the impulse response. Although impulse response analysis has been previously used for sensing context, we explore its use for measuring the content level of a container specifically in this paper.

Applications of Impulse Response Analysis

An impulse response is the output of an environment (*e.g.,* a space, a physical object or a container) when presented with a short duration audio signal (*e.g.,* Dirac impulse). Because an impulse response yields a complete description of the changes that a sound signal undergoes when it travels from one point to another [23], it has been widely used to exam the properties of a space or a physical object.

Kunze and Lukowicz [22] previously demonstrated that a mobile phone's location can be inferred by analyzing how its surrounding space acoustically reacts to a vibration and short narrow frequency "beeps." Diaconita *et al.* [12] examined the efficacy of different acoustic probing sequences for determining a mobile phone's location. They discovered that using Gaussian Noise as the probing sound and Mel-Frequency Cepstral Coefficients (MFCCs) and Delta MFCCs as features together produced the best performance. Rossi *et al.* [31] showed that the impulse response analysis approach can be used to localize a user's room level position with 98% accuracy. Fan *et al.* [15] explored the use a sine wave sweep as an active probe sound to excite the environment and then analyzed the impulse response of the space to determine when the user is a restroom or not.

In addition to examining the acoustic fingerprint of a space, where air is the major transmission media for sound, impulse response analysis has also been used to infer the internal physical properties of objects (*e.g.*, food and agriculture), where a solid or liquid acts as the major transmission media. For example, Diezma-Iglesias *et al.* [14] constructed a device that used a mechanical impulse sound generator and a microphone to detect the internal quality of watermelons (*e.g.,* how hollow and ripe a watermelon is). A mechanical impulse sound generator hits a watermelon on one side, and a microphone records the sound that travels through the watermelon on the other side. The device then performs a spectral analysis on the recorded sounds and builds a model for diagnosing the condition of a watermelon. Conde *et al.* [8] have used a similar approach to measure the internal cracks in cheese. Their device uses two impact probes to hit the cheese and a microphone to record the sound that traveled through the cheese. Finally, Lin *et al.* [24] use a light stick to hit the surface of an egg and record the reaction sound through a microphone. They then extract and analyze several frequency features to identify whether an egg's shell has a crack. Instead of using a mechanical system to generate the probing sound, SoQr uses a speaker to output a probing sound to excite a container and its content. We then explore the use of impulse response analysis to quantify the content level inside that container.

The Touch & Activate [28] project showed that a vibration speaker and a piezo-electric microphone can be attached to an ordinary object to make it touch sensitive. By constantly emitting an acoustic signal and measuring the change in frequency spectrum, five hand touch positons and six hand postures can be identified. In comparison, SoQr does not need to emit sounds all the time and focuses specifically on sensing the amount of contents in a container.

THEORY OF OPERATION

The acoustical properties of a container depend not only on the container itself, but also what content that it holds and how much of that content is present. This is because sound waves get absorbed and reflected while traveling at and through different medium (*e.g.,* air, liquid and solid). The physical properties of a media dictate their absorption and reflection rates. For instance, viscosity and thermal conduction are examples of two properties that can affect the absorption of sound waves in fluids [7]. Therefore, the same sound traveling towards different physical objects may be affected in different ways because of the differences in those

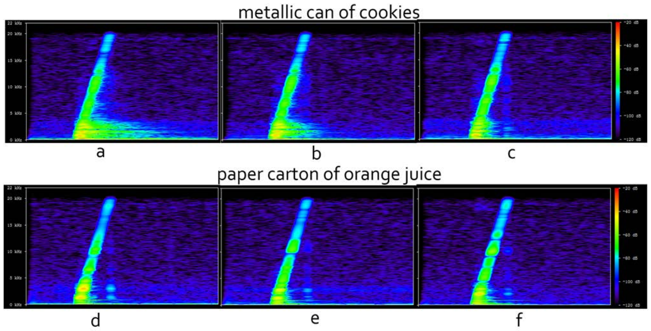

Figure 1. Spectrograms of impulse responses for 2 products (cookies, orange juice) in their original packaging when excited by a sine wave sweep at different content levels: empty (a, d), half (b, e), and full (c, f). The vertical axis is frequency and the horizontal axis is time. Color represents the intensity of sound.

objects. This means that when a sound is emitted at a container, the impulse response to that sound will be different when the amount of content in it changes.

This concept is leveraged by some consumers when they knock on the surface of a watermelon to judge the ripeness of the fruit by listening to the sound. The knocking act produces a probing sound, which is affected by the internal properties of the watermelon (which includes how much sugar and how much water it contains). The final sound is used by the consumer to tell if the fruit has ripened without having to cut open the fruit. Because a ripe watermelon consists of more than 90% water, it sounds different than an unripe one.

In this paper, we explore the translation of the concept and practice described above into a sensor design that outputs a sine wave sweep probing sound to excite a container instead of physically knocking against it, and records and analyzes the impulse response of the container to determine how much content is inside it. Because we do not have prior knowledge of what frequencies are best to excite different household items, we use a wide frequency range from 20 to 20K Hz for the sine wave sweep. Figure 1 shows example spectrograms of the recorded impulse responses to a sine wave sweep probing sound for two different containers (a metallic cookie can and a carton of orange juice) at three different content levels (empty, half, and full). The high intensity slope line area is the impulse response of a container to the outputted probing sound. Note that for the same container, the impulse response differs for the different content levels. In this paper, we explore how to identify and learn these different impulse responses to estimate the content level inside different containers.

SYSTEM IMPLEMENTATION DETAILS

To estimate the content level of a container, we developed a system, called SoQr, which uses and analyzes sound to quantify the amount of content. The goal of this sensor is to acoustically excite a container by outputting a probing signal and at the same time record the impulse response. We then use supervised learning methods to build prediction models, which take the impulse response audio recordings as input and estimate the content level of a container as either discrete levels or continuous levels.

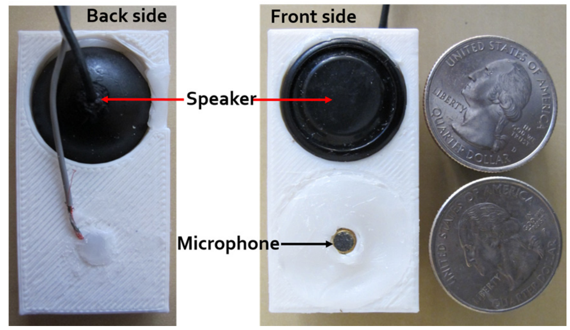

Figure 2. The SoQr sensor prototype which consists of a speaker and a microphone. When the sensor is installed on a container, the front side of the sensor (right) is placed to face and make contact with the container's surface.

Sensor Unit Design

The sensor unit consists of a microphone and a speaker. The speaker is used to output a probing sound, and the microphone records the impulse response of the examined container. To make it possible to install the sensor on containers of many various shapes and sizes, we developed a small 3D printed case (47 cm in length by 24 cm in width) to hold the speaker and microphone together (Figure 2). We position the speaker and microphone to both face outwards with their surfaces aligned flushed with the case's surface. In this way, when the sensor is attached onto a container, the speaker and microphone can come into contact as close as possible to the container's surface.

For prototyping purposes, we currently connect the SoQr sensor to the audio jack of a commodity smartphone (Galaxy Nexus). The smartphone runs a data sampling application that we developed to play a probing sound and record the impulse response at the same time. The wires connecting the SoQr sensor to the smartphone can eventually be replaced by a Bluetooth connection, however, we do not implement and test with that feature in this paper.

Data Collection Procedure

The data sampling application that runs on a commodity smartphone outputs a 20~20K Hz sine wave sweep as the probing sound and starts recording at the same time using 44.1 KHz with 16 bit depth sampling method. The recording lasts 1 second after the sweep stops in order to fully capture the impulse response. We tested three different lengths (1, 0.1 and 0.01 second) for the probing sound and chose 0.01 second as the length because it produced the highest prediction performance and was the least obtrusive.

We used the data collection protocol describe below for all the evaluations described in this paper unless otherwise noted. For a given container X filled with content type Y, we follow the following steps to collect data samples.

1. Measure the weight of a container with full content: W_f.

2. Remove all the content stored in the container being examined into an intermediate container and measure the weight of the empty container being examined: W_e
3. Calculate $\frac{1}{N}$ weight of its content by using the equation: $\frac{W_f - W_e}{N}$. Note that for rigid containers, the $\frac{1}{N}$ weight is equivalent to $\frac{1}{N}$ volume.
4. Install the SoQr sensor on the center of a flattest side surface of the testing container.
5. Use the sensor unit to play a 0.01 second sine sweep probing sound and at the same time record the impulse response for 1 second after the sweep completely stops.
6. Wait for *5* second and then repeat the step 5 to start another round of probing and recording until a total of *K= 100* rounds are finished.
7. Remove $\frac{1}{N}$ content from the examined container, seal the container if it has a lid or leave it open if it does not have one. Then repeat the steps 5, 6 and 7, until the container is completely empty.

The data collection procedure should generate data samples for N + 1 different content levels, and (N+1) * K impulse response recordings in total for each examined container.

Feature Extraction

Before extracting features, we analyze the portion of the audio before the sweep starts to see if there is environment noise and perform noise removal on the recordings accordingly. To learn the prediction models, we then need to process the raw impulse response audio recordings into representative features. Based on the existing literature [12,15,22,31], we have identified the Mel-Frequency Cepstral Coefficients (MFCCs) related features as the most promising ones for our tasks. Thus, we extract features as follows:

1. Convert the time domain data samples into frequency domain features by first dividing the audio recording of impulse response into frames using a sliding window of size *W = 512* samples, with 50% overlap between two adjacent windows. Then we apply a Hamming window function to each frame. Lastly, we perform a Fast Fourier Transform (FFT) on each frame j *(j = 1.. N. N: the number of frames in the impulse response recording)* to get the magnitude of each frequency bin *k (k = 1,.., W)*: $FFT_{j,k}$.
2. Locate the portion of the recording that the probing sound impacts the most. Because the length of a recording is longer than the length of the impulse response itself, the recording might potentially include unwanted environment sound. To estimate the start *i* and the end of the portion that the sweep impacts the most, apply the optimization: $arg\ max_{i \in \{1 \dots N\}} \sum_{j=i}^{\min\{i+S,N\}} \sum_{K=0}^{W-1} FFT_{j,k}$ (*S: the number of frame-for the length of the sweep itself;* $S = \frac{44100 * 0.01}{W * 50\%}$).
3. For the portion of the recording extracted in step 2, calculate MFCCs for each frame *i* and keep the first 13 coefficients: $MFCC_{i,1}, \dots, MFCC_{i,13}$.
4. To capture the temporal change over adjacent frames, compute the mean of each of the 13 MFCC over all frames: $MMFCC_k = \frac{\sum_{i=1}^{N} MFCC_{i,k}}{N}$ (*i = 1, …, N; k = 1,…, 13*).

Prediction Model Learning

Our goal is to learn a prediction model that can tell the content level in a container. We can treat this goal as a classification problem, in which the model predicts the content level from the *N+1* levels that we have gathered data samples and trained on. To explore the efficacy of different machine learning models on our prediction task, we explored three machine learning techniques: K-Nearest Neighbor, Support Vector Machines (SVM) and Random Forest. In this paper, we only reported the results of SVM for its superior performance. We leveraged libSVM [4] for our implementation and used Radial Basic Function (RBF) as the kernel function for its superior performance over linear, polynomial and sigmod kernels on our prediction tasks.

At the same time, because the household products that we tested would be consumed or used in continuous manner, we can also treat our prediction task as a regression problem, in which our model predicts the ratio of empty space in a container as a continuous value ranged from 0 to 1. We leveraged Support Vector Regression (SVR) model and used libSVM for implementation again for this task.

EVALUATION

We first evaluate the efficacy the SoQr system for predicting the content levels of different household items in their original container or packaging. Then we investigate how the potential ways that the user interacts with a container and the sensor can potentially influence the system's performance. Note that we use the term container and packaging interchangeably here to mean the object which holds the household products being quantified.

Estimating the Content Level of Common Household Items

To evaluate the SoQr system's ability to predict the content level in a container, we tested the sensor on 19 common household items in their natural packaging. We selected these products to have coverage of many different container properties (packaging material, its shape, whether or not it's deformable, and if it can be reclosed) and content properties (its type, density/thickness if it is a liquid or gel, and granularity if it is a solid). A detailed description of the 19 items is shown in Table 1.

During data collection, we placed each item in Table 1 at the center of a table and placed the SoQr sensor at the center of a flattest side surface for that container. We wrapped a small cling film around the SoQr sensor to help secure it on a container's surface. For any container that has a curved surface primarily (*e.g.,* wine bottle), we chose a vertical spot at the center of the container to install the sensor. For deformable items that could not stand on their own (*e.g.,* bags of chips and rice), we placed them at the edge of the table so that they could lean against the wall to keep them standing upright, mirroring how they might be placed in a pantry or on a shelf. During the data collection, we reclosed the lid on any containers that has a lid and left open those that did not have a lid (*e.g.,* bags of chips and rice). Figure 3 shows the household items tested in this part of the evaluation with the SoQr sensor installed on each.

Item	Product/Content		Container/Packaging			
	Type (Liquid, Solid, Gel, Other)	Density/thickness (if item is a liquid or gel) and Granularity (if item is a solid)	Material	Shape	Deformable	Closable (Has lid?)
Cereal box	Solid	Coarse granularity	Paper (with an additional plastic bag inside)	Rectangular box	Yes	Yes
(Potato) chips	Solid	Coarse granularity	Plastic	Rectangular bag	Yes	No
Coke	Liquid	Thin liquid	Plastic	Cylindrical bottle	No	Yes
Cookie can	Solid	Coarse granularity	Metallic (with the cookies themselves packaged in some plastic bags)	Rectangular box	No	Yes
Cooking oil	Liquid	Thick liquid	Plastic	Roughly rectangular bottle	No	Yes
Flour	Solid	Fine granularity	Paper	Rectangular bag	Yes	No
Grape jelly	Gel	Thick gel	Plastic	Roughly rectangular bottle	No	Yes
Ice cream	Other (semi-solid)	Thick semi-solid	Paper	Roughly rectangular box	No	Yes
Laundry detergent	Liquid	Thick liquid	Plastic	Jug	No	Yes
Laundry stain remover	Solid	Fine granularity	Plastic	Roughly rectangular box	No	Yes
Milk	Liquid	Thin liquid	Plastic	Jug	No	Yes
Orange juice	Liquid	Thin liquid	Paper	Rectangular box	No	Yes
Penut butter	Gel	Thick gel	Plastic	Cylindrical jar	No	Yes
Red wine	Liquid	Thin liquid	Glass	Cylindrical bottle	No	Yes
Rice	Solid	Medium granularity	Nylon (with an additional plastic bag inside)	Rectangular bag	Yes	No
Salsa sauce	Other (semi-solid/liquid)	Thick semi-solid/liquid	Glass	Cylindrical jar	No	Yes
Salt	Solid	Fine granularity	Paper	Cylindrical jar	No	Yes
Shampoo	Liquid	Thick liquid	Plastic	Rectangular bottle	No	Yes
Egg	Solid	Coarse granularity	Styrofoam	Rectangular box	No	Yes

Table 1. The 19 common household items used to test SoQr and details about each product and its container.

Figure 3. The household items shown with SoQr installed on them.

Item	Classification (SVM)			Regression(SVR)	
	Overall Precision	Overall Recall	Overall F-measure	Correlation Coefficient	Mean Absolute Error
Cereal box	0.950	0.949	0.949	0.9857	0.0528
Chips	0.980	0.980	0.980	0.9648	0.0616
Coke	0.956	0.955	0.955	0.9761	0.0551
Cookie can	0.999	0.999	0.999	0.986	0.0489
Cooking oil	0.996	0.996	0.996	0.9851	0.0439
Flour	0.997	0.997	0.997	0.9855	0.0495
Grape jelly	0.995	0.995	0.995	0.978	0.0513
Ice cream	0.963	0.960	0.960	0.9578	0.0671
Laundry detergent	0.980	0.980	0.980	0.9790	0.0522
Laundry stain remover	0.983	0.983	0.983	0.9824	0.0535
Milk	0.994	0.994	0.994	0.9852	0.0573
Orange juice	0.967	0.965	0.966	0.9832	0.0487
Peanut butter	0.972	0.972	0.972	0.9811	0.0516
Red wine	0.967	0.965	0.966	0.9609	0.0665
Rice	0.970	0.968	0.969	0.9736	0.0583
Salsa sauce	0.941	0.939	0.939	0.979	0.0487
Salt	0.864	0.856	0.855	0.9707	0.0572
Shampoo	0.987	0.986	0.986	0.9818	0.0484
Average	0.970	0.969	0.969	0.978	0.054

Table 2. The performance of classification and regression prediction models.

We then followed the protocol described in *Data Collection Procedure* section to collect data samples with each item. Starting from a full package, we removed 10% of the content from its container each time *(N = 10)* until the package was empty. There are 11 content levels in total, including the full and empty levels. For each level, we collected 100 samples. Therefore, there were 1100 samples for each item.

We then extracted features and trained an SVM classifier and an SVR regression model for each household item according to the procedure of *Feature Extraction* and *Prediction Model Learning* sections. To evaluate our models' performances, we adopted the 10-fold cross validation strategy. The SVM and SVR models' performances on the first 18 household products in their original packaging are shown in Table 2. For classification, the overall F-Measure was above 0.85 for all items (with the overall F-Measure being ~0.969). For regression, the mean absolute error was ~0.05 for all items, which was about the half of the two adjacent levels' interval (0.1). The results demonstrated that SoQr can reliably predict the content level of common household items in their original packaging.

In addition to the first 18 items, whose contents always settle at the bottom of the packaging and extend vertically based on how much of the content remains, we also conducted data collection and evaluation on a carton of 12 eggs that were arranged horizontally within the container. We installed the

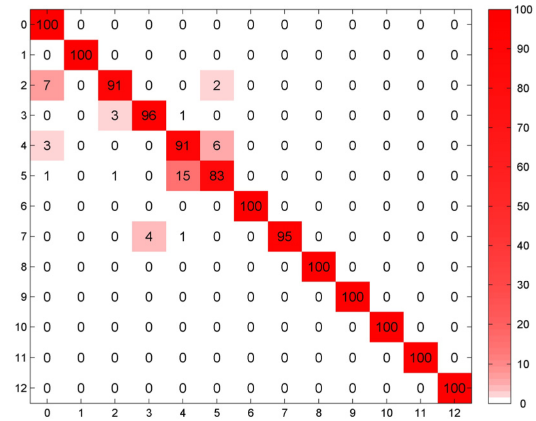

Figure 4. Confusion matrix showing the result of 10-fold cross validation of the SVM model for a carton of 12 eggs (Row: true class; Column: predicted class).

sensor unit on a vertical side surface of the carton (Figure 3 right). Starting from a full carton, we removed one egg each time from one end to the other following a zigzag sequence, collecting again 100 data samples at each "level." The overall precision, recall and F-Measure of the 10-fold cross validation were all over 0.96. The confusion matrix is shown in Figure 4. 65% of the misclassifications were happened in adjacent classes.

Robustness of SoQr in Different Usage Scenarios

The results from above suggest that a prediction model can be trained to predict the content level for each of the 19 tested household products fairly accurately. However, the models were trained and tested in a controlled manner—data was collected in only a single location, the product was always in an upright orientation, the lids were always opened or always closed, and the sensor was installed at only one location on each product. In this section, we remove those constraints and examine how SoQR might perform in response to some potential ways that the user might interact with the household products and the sensor itself.

Different Locations Where a Container Might Be Stored & Used

Many household items can be stored in, moved to, and used at different locations in the home. This raises the question of whether a prediction model trained on data samples collected at one location would still work at a different location.

To explore this problem, we chose a quarter gallon of milk as an example and collected data with this item at two different locations where this item is commonly found in the home: inside a refrigerator and on a kitchen countertop. These two locations are quite different in the following ways. The fridge is a relative small and enclosed space. It also has periodic low frequency noise created by the vapor compression cycle. In contrast, the kitchen countertop is a relatively open space. We followed the same protocol to collect data samples at these two locations and then extracted features. We then built and tested the SVM model in two different ways. First, we trained the model on the data collected in the fridge, and tested on the data collected on the kitchen countertop. Second, we performed the training and testing in the reverse order. The overall precision, recall, and F-Measure of the model trained on the fridge data set and tested on the kitchen data set were 0.976, 0.975, and 0.975 respectively. The overall precision, recall and F-Measure of the reverse case were: 0.935, 0.923 and 0.916. The results demonstrate that the model built on the data samples collected at one location can still work reasonably well at very different location.

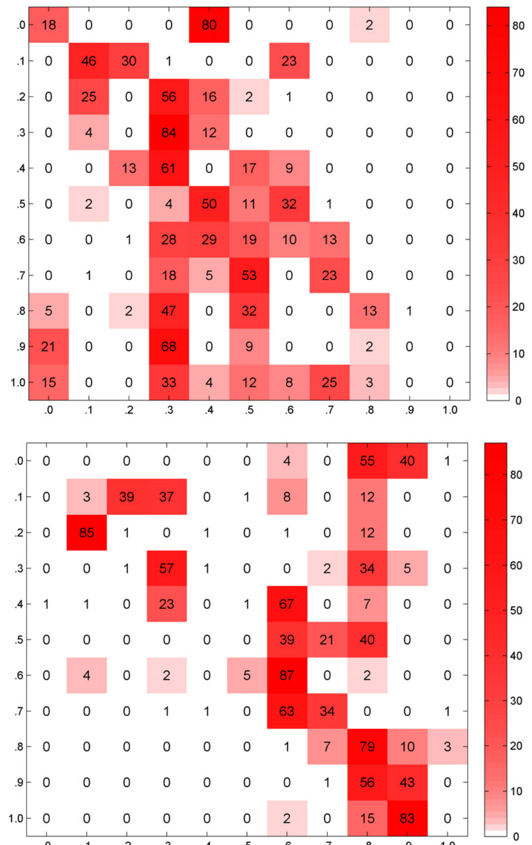

Figure 5. (Top) Confusion matrix of the classification results for the SVM model trained on data collected in standing upright orientation and tested on data collected in laid down orientation. (Bottom) Confusion matrix for the reverse case.

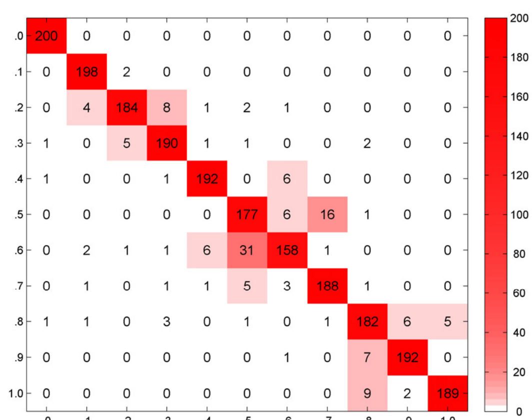

Figure 6. Confusion matrix showing the results of a 10-fold cross validation for the SVM model trained on the combined data set of the two orientations.

Different Orientations in Which a Container Might Be Placed

Although people typically store their products in an upright orientation, there are some items that can be stored in other ways. This raises the question of whether a prediction model trained on data collected in one orientation would still work in a different orientation. We explored this problem using a cereal box as an example product and followed the aforementioned protocol to collect data samples in two different orientations that is common for a cereal box: standing upright and laid down on a table (Figure 7). We then trained an SVM model on data samples from one orientation and tested on the other orientation. The overall precision, recall and F-Measure were (0.211, 0.186, and 0.161) and (0.170, 0.276 and 0.194) respectively. The confusion matrices of the classification results are shown in Figure 5. A significant amount of misclassifications occurred in both cases. This suggests that the models trained in one orientation do not directly apply to other.

A natural follow-up question is: how well would a model trained with a combined data set of the two orientations classify data samples in these two orientations? We trained an SVM model on the combined data samples for these two orientation and then performed 10-fold cross validation. The confusion matrix is shown in Figure 6. The overall precision, recall and F-Measure are: 0.933, 0.932 and 0.932. This indicates that the model can classify the content level in either orientation if it is trained on samples from both orientations.

Different Ways to Store a Product in Its Original Packaging

In all of the tests described so far, we always reclosed a container if it has a lid or left a product opened if it did not

Figure 7. Common orientations of a cereal box: (left) standing upright and (right) laid down on a table.

Figure 8. A bag of chips in various opened/closed states: (left) opened, (middle) clipped, and (right) tied.

have a lid. In practice, a user may not always reclose a container even if it has a lid. The user may also use a clip or a rubber band to reclose a packaging even if it does not have a lid. This raises the question of whether a prediction model trained on data collected for a container in one opened/closed state would still work with the container in a different opened/closed state. To explore this problem, we first chose a bag of potato chips as the test product. We opened the bag by cutting its top end. We then tested the models for three common opened/closed states for the product: when it is left opened, reclosed at the top with a clip, and tied at the top with a rubber band (Figure 8).

We again followed the same data collection protocol to collect data samples for these three opened/closed states. Then we trained an SVM model using the data samples collected while the bag was left open, and tested on the data set collected when the bag was "clipped." The overall precision, recall and F-Measure were: 0.550, 0.491, and 0.431. We then tested the same model with the data set collected when the bag was "tied". The overall precision, recall and F-Measure 0.040, 0.045, and 0.037. The confusion matrices of the classification results are shown in Figure 9. The relatively poor performance of both cases suggest that our prediction model is not able to work on deformable packages, because the degree of deformation will likely always differ each time it is reclosed regardless of the method. This means that although the performance results reported in Table 1 were generally positive for potato chips and rice, a model trained for these products will likely not be useful for predicting the content level of those products again in real practice.

Although it is clear that a prediction model trained for products in a deformable container would not work when the container is closed and the shape of the container is affected, it is important to then confirm that this is because of the container/packaging is deformable. We then performed this test again with a rigid container and has a reclosable lid. We chose for this purpose a bottle of cooking oil and collected data samples with the lid on and without the lid. We then trained an SVM model on the data set with the lid on and tested against the data set without the lid. The overall precision, recall and F-Measure are 0.842, 0.802 and 0.786. We reversed the training and testing data sets and did the evaluation again. The overall precision, recall and F-Measure are 0.834, 0.829 and 0.812. The confusion matrices of the classification results are shown in Figure 9. The overall performance was reasonable and the majority of the misclassifications happened to adjacent classes.

Different Positions Where a Sensor Can Be Installed

In all of the tests described so far, we have only installed SoQr at one location on each product—at the center of the flattest surface on the side of its packaging. Although the classification results shown in Table 2 are promising, an important question to answer is whether the user would need to install the sensor at the same location on for each product?

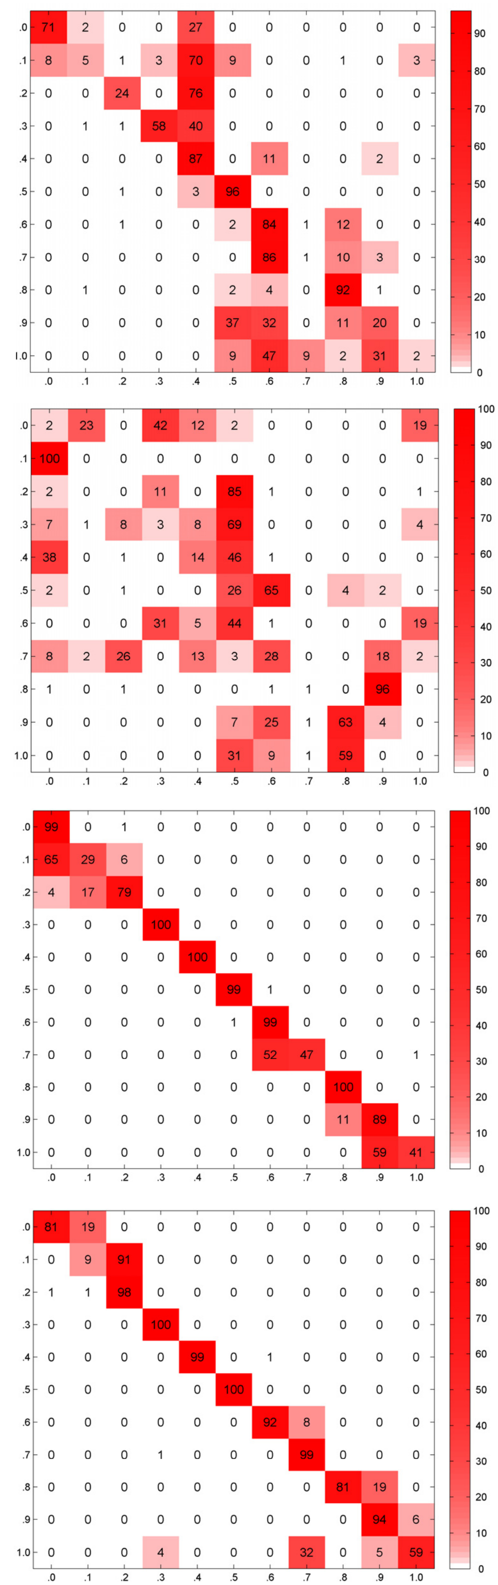

Figure 9. (Top) Confusion matrix of the classification results of an SVM model trained on data collected with the "opened" bag and tested on data collected with bag "clipped." (Second) Results of the same SV model tested on data collected with bag "tied." (Third) Confusion matrix showing the results of an SVM model trained on data collected for a bottle of oil with lid on and tested with the lid off. (Bottom) Confusion matrix of the reverse case.

That is, would installing the sensor unit at a different position of the same product affect a prediction model's ability to determine the content level inside that container? To explore this problem, we used a carton of orange juice as the test product and collected data with the sensor unit placed at 3 different spots: top-center, middle-center and bottom-center of a side surface. We trained SVM models with the data set for one placement and tested on data sets for the other two placements. There were 6 combinations in total. The evaluation results showed that, the overall F-Measure for all the combinations were <0.02. This implies that a model trained with the sensor installed at one spot would not work if the sensor is moved to a different spot. Because there are many different positions where a sensor can be placed on a container, it would not be possible to develop a model that combines data collected from all possible positions.

DISCUSSION

The preliminary tests of how robust the sensor is to the potential ways in which the user might interact with these products and the sensor in practice point out some important practical challenges that must be considered.

Building Customized Prediction Models for Each Product

The unique properties of different types of containers and different types of contents mean that the SoQr sensor would require that customized model is built for each unique product and packaging combination. Although collecting the necessary data samples to build the prediction models requires a large amount of time and effort at the moment, this burden does not have to fall completely on the end-users. Instead, a company that manufactures or produces a product can potentially perform this data collection to create a prediction model for all of their consumers (*e.g.,* a soda company can create a prediction model that can be used for all their 3 liter plastic bottled products). Even if this does not happen, potentially a prediction model for a product would only need to be created once by the first user who wishes to collect the necessary data for that item and share that model with other consumers. However, it might be necessary to examine how to minimize the amount of time and effort needed to collect the data to build a prediction model so that it is feasible to imagine a community of users participating in the collection and sharing this information. One question related to this goal is how many different content levels would need to be sampled to build an accurate and precise prediction model, and which levels are the best ones to ask users to help collect as they use their products?

Installing the Sensor on a Container

There are many potential spots on a container's surface where SoQr can be placed. Our test results revealed that when the model is trained with the SoQr sensor installed at one position on the container, the performance of the prediction model degrades when the sensor is installed in a completely different spot. Although there might be consistent spot (*e.g.,* the barcode area) on different containers that can be used as the default place where users can be instructed to install SoQr, the size and position of such spots may make it hard for the user to accurately place the sensor onto the container in a way that matches how it was originally positioned when training a model. Therefore, two important questions that must be explored are 1) how can the sensor be designed to help guide the end-users with the process for installing the sensor at the desired spot, and 2) is it possible to develop a prediction model that relaxes the need for the SoQr sensor to be installed at a specific spot without comprising too much precision and accuracy?

Quantifying Content Level in Untrained Context

We tested the prediction models ability to estimate the content level of a container when it is placed in an entirely different place than where the training data was collected as well as in a second common orientation in which it might be stored. While the system was already able to accurately predict the content-level of a product in different locations, it needed to be trained with a combined data set collected for a product in the two different orientations before it was able to accurately predict the content level in either orientation. Although this result is promising, this test was limited to only a small number of products, in a small number of locations, and a small number of orientations. Whether the prediction models would perform similarly at more locations in the home and across homes is an important research question that must be explored. The answer to these questions may dictate how the prediction models should be created and by whom. At the same time, it might also be useful to determine whether location independent features can be identified and used, or if additional hardware and software is needed to help isolate the probe sound and the impulse response from the environment sounds. Finally, it is important to note that sometimes a product may be found in a completely arbitrary orientation (*e.g.,* when a product is lies partially atop other items or is leaning against a wall). Thus, it is also necessary to explore how to train the system in a limited number of orientations (*e.g.,* the three orthogonal orientations) but yet can still work in an arbitrary orientations. This requires an examination into whether there is a relationship that exists between how a container is oriented and its impulse response to the probing sound.

CONCLUSION

In this paper, we present SoQr, a portable sensor that can be installed on a surface of a container to estimate its content level. We demonstrated that SoQr can predict the content level with the average overall F-Measure above 0.96 for 19 common household items in their original packaging. We explored how SoQr might perform in response to some potential ways that the user might interact with different household products and the sensor itself. Overall, we show that SoQr can be trained to accurately estimate the content level of common household items found in rigid containers, even when the product placed in different locations in the home and stored different orientations, but mechanisms are needed to facilitate the training of these prediction models in different context and to help the user install the sensor.

REFERENCES

1. AdhereTech. http://adheretech.com/

2. Caldwell, J., Slobodnik, (2008). M. Fluid level measuring system. *U.S. Patent No. 7,421,895*. 9 Sep. 2008.
3. Canbolat, H. (2009). A novel level measurement technique using three capacitive sensors for liquids. *Instrumentation and Measurement, IEEE Transactions on*, 58(10), 3762-3768.
4. Chang, C., Lin, C. (2011). LibSvm: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, vol. 2(3), 2011, 188-205.
5. Chi, P. Y. P., Chen, J. H., Chu, H. H., & Lo, J. L. (2008). Enabling calorie-aware cooking in a smart kitchen. In *Persuasive Technology* (pp. 116-127). Springer Berlin Heidelberg.
6. Chiu, M. C., Chang, S. P., Chang, Y. C., Chu, H. H., Chen, C. C. H., Hsiao, F. H., & Ko, J. C. (2009). Playful bottle: a mobile social persuasion system to motivate healthy water intake. In *Proceedings of the 11th international conference on Ubiquitous computing (UbiComp'09)* (pp. 185-194). ACM.
7. Christensen, J., Romero-Garcia, V., Picó, R., Cebrecos, A., de Abajo, F. G., Mortensen, N. A. & Sánchez-Morcillo, V. J. (2014). Extraordinary absorption of sound in porous lamella-crystals. *Scientific reports, 4*.
8. Conde, T., Mulet, A., Clemente, G., & Benedito, J. (2008). Detection of internal cracks in Manchego cheese using the acoustic impulse-response technique and ultrasounds. *Journal of dairy science, 91(3)*, 918-927.
9. Cosense-Ultrasonic Fluid Detection Sensors. http://www.medicaldevice-network.com/contractors/mechanics/cosense-sensors/cosense-sensors4.html
10. Cuptime. htttp://cuptime.org
11. Dam, N., and Austerlitz, H. (2003). System and method of non-invasive discreet, continuous and multi-point level liquid sensing using flexural waves. *U.S. Patent No. 6,631,639.* 14 Oct. 2003.
12. Diaconita, I., Reinhardt, A., Englert, F., Christin, D., & Steinmetz, R. (2014). Do you hear what i hear? Using acoustic probing to detect smartphone locations. In *Pervasive Computing and Communications Workshops (PERCOM Workshops), 2014 IEEE International Conference on* (pp. 1-9). IEEE.
13. Dietz, P. H., Leigh, D., & Yerazunis, W. S. (2002). Wireless liquid level sensing for restaurant applications. In *Sensors, 2002. Proceedings of IEEE* (Vol. 1, pp. 715-720). IEEE.
14. Diezma-Iglesias, B., Ruiz-Altisent, M., & Barreiro, P. (2004). Detection of internal quality in seedless watermelon by acoustic impulse response. *Biosystems engineering*, 88(2), 221-230.
15. Fan, M., Adams, A. T., & Truong, K. N. (2014). Public restroom detection on mobile phone via active probing. In *Proceedings of the 2014 ACM International Symposium on Wearable Computers (ISWC'14)* (pp. 27-34). ACM.
16. Fram, E. H., & Axelrod, J. (1990). The distressed shopper. *American Demographics*, 12(10), 44-45.
17. Ganglbauer, E., Fitzpatrick, G., Molzer, G. (2012). Creating visibility: understanding the design space for food waste. In *Proceedings of the Proceedings of the 11th International Conference on Mobile and Ubiquitous Multimedia*, (pp1-10). ACM.
18. Goekler, L. E. (1991). Capacitive liquid level sensor. *U.S. Patent No. 5,017,909*. 21 May 1991.
19. Hinz, William R., and Dennis M. Kotz. (2010). Ultrasonic liquid level detector. *U.S. Patent No. 7,802,470*. 28 Sep. 2010.
20. Lin, Y. (2007). Capacitive liquid level sensor. *U.S. Patent No. 20,070,240,506*. 18 Oct. 2007.
21. Lo, J. L., Lin, T. Y., Chu, H. H., Chou, H. C., Chen, J. H., Hsu, J. Y. J., & Huang, P. (2007). Playful tray: adopting Ubicomp and persuasive techniques into play-based occupational therapy for reducing poor eating behavior in young children. In *Proceedings of the 9th international conference on Ubiquitous computing (UbiComp '07)* (pp. 38-55). Springer Berlin Heidelberg.
22. Kunze, K., Lukowicz, P. (2007) Symbolic object localization through active sampling of acceleration and sound signatures. *In Proceedings of the 9th international conference on Ubiquitous computing (UbiComp '07)* (pp. 163-180). Springer Berlin Heidelberg.
23. Kuttruff, H. (2009). *Room Acoustics*. CRC Press.
24. Lin, H., Zhao, J. W., Chen, Q. S., Cai, J. R., & Zhou, P. (2009). Eggshell crack detection based on acoustic impulse response and supervised pattern recognition. Czech J. Food Sci. Vol, 27(6), 393-402.
25. Mukherjee, S. (2010). Non-invasive measurement of liquid content inside a small vial. In *Radio and Wireless Symposium (RWS), 2010 IEEE*, pp. 527-530. IEEE.
26. Nakagawa, T., Hyodo, A., Osada, K., Kurata, H., Oho, S. (2011). Contactless liquid-level measurement through opaque container using millimeter-wave sensor. In *Sensors, 2011 IEEE*, pp. 1421-1424. IEEE.
27. Olson, R. S., & Christensen, E. L. (1990). Nonintrusive acoustic liquid level sensor. *U.S. Patent No. 4,901,245.* 13 Feb. 1990.
28. Ono, M., Shizuki, B., and Tanaka, J. (2013) Touch & activate: adding interactivity to existing objects using active acoustic sensing. In Proceedings of the 26th annual ACM symposium on User interface software and technology *(UIST '13)* (pp. 31-40). ACM.

29. Park, K. (1990). Liquid level sensor. *U.S. Patent No. 4,924,702*. 15 May 1990.

30. Reverter, F.; Li, X.; Meijer, G. (2007). Liquid-level measurement system based on a remote grounded capacitive sensor, *Sensors and Actuators A: Physical*, Volume 138, Issue 1, 20 July 2007, 1-8.

31. Rossi, M., Seiter, J., Amft, O., Buchmeier, S. and Tröster, G. (2013). RoomSense: an indoor positioning system for smartphones using active sound probing. In *Proceedings of the 4th Augmented Human International Conference* (pp.89-95), ACM.

32. Soltz, D. J. (1984). Ultrasonic liquid level meter. *U.S. Patent No. 4,470,299*. Sep. 1984.

33. Terzic, E., Nagarajah, C. R., & Alamgir, M. (2010). Capacitive sensor-based fluid level measurement in a dynamic environment using neural network. *Engineering Applications of Artificial Intelligence*, 23(4), 614-619.

34. Tward, E., & Junkins, P. (1983). Multi-capacitor fluid level sensor. *U.S. Patent No. 4,417,473*. 29 Nov. 1983

35. Vessyl. https://www.myvessyl.com