

Automatic Classification of Daily Fluid Intake

¹Jonathan Lester, ²Desney Tan, ¹Shwetak Patel, and ²A.J. Bernheim Brush

¹Department of Electrical Engineering
Computer Science & Engineering, DUB Group
University of Washington
Seattle, WA, USA
{jlester,shwetak}@u.washington.edu

²Microsoft Research
Redmond, WA, USA
{desney,ajbrush}@microsoft.com

Abstract—Despite the potential health benefits of being able to monitor and log one’s food and drink intake, manually performing this task is notoriously hard. While researchers are still exploring methods of automating this process for food, less work has been done in automatically classifying beverage intake. In this paper, we present a novel method that utilizes optical, ion selective electrical pH, and conductivity sensors in order to sense and classify liquid in a cup in a practical way. We describe two experiments, one that uses a high end commercial off-the-shelf spectrometer, and the other which uses a cheap sensor package that we engineered. Results show both that this method is feasible and relatively accurate (up to 79% classification for 68 different drinks), but also that we would be able to build this in such a way as to make it practical for real-world deployment. We describe the vision for building a sensor rich cup capable of determining the kind of liquid a person is drinking, as well as the opportunities that the success of such sensors may open.

Keywords—*Drink classification, spectrometer, food, caloric intake, weight loss, cup, sensor, pH, conductivity, and health*

I. INTRODUCTION

Even with widespread attention in the news, popular media, and a multitude of commercial products, weight management and proper nutrition remains a very real and pressing problem for Americans and a growing segment of the world population. Obesity creates a range of lifestyle, psychological, and social challenges as well as contributing to long term health problems ranging from diabetes to heart disease. It is not surprising that 31% of the US population is actively trying to lose weight [9].

At its most basic, weight management is a mismatch between the number of calories a person consumes and the number of calories they expend. Research has shown that people often struggle at managing this balance [6], and that they are only able to maintain behavior changes for short periods of time [21]. For dieticians and counselors to help their patients, they use their patient’s exercise routines, their daily caloric intake, and their food choices for the day and why they made those choices. Technology has helped in many regards, pedometers and more advanced calorie counting systems have been developed in recent years to provide greater insight into where individuals expend calories [4]. However, methods for understanding what a person has consumed are still largely manual and self-reported, with the most common method being 24 hour recall surveys, where people are asked to remember the foods they have eaten in the past 24 hours. Using this information, a patient’s energy intake and nutrition can be estimated. However, continuous manual tracking requires a non-trivial

amount of effort and many people find it difficult to keep to a diligent regiment.

One particularly important piece of information about one’s diet is the amount and types of fluids consumed throughout the day. Drinks make up a surprisingly large portion of daily caloric intake with some research suggesting that 21% of a person’s daily caloric intake comes from beverages (458 calories) [5]. These tend to be ‘optional’ calories, not consumed exclusively to satiate hunger, and thus potentially easier to eliminate or replace with healthier alternatives. There is also some evidence to suggest that even drinking diet soft drinks can actually cause an increase in weight gain, making it important to also keep track of these low calorie beverages [7].

In our work, we attempt to design and build a cup capable of non-destructively identifying the liquid inside it. To enable long term monitoring on the scale of months and years, it is important that the technology be unobtrusive and require minimal effort. This is particularly important because people tend to drink much more frequently than they eat, which makes manual logging even harder. Our approach to long term fluid intake monitoring, is a smart cup which combines pH, conductivity, and light spectrum to fingerprint different liquids and allow us to distinguish different beverages.

In this paper, we first describe related work and then our design goals and the sensors we used. We then present an overview of our approach for sensing beverages and the prototype sensing system we built. Next, we discuss the results of classifying a variety of drinks across multiple days. We then present a low cost prototype sensor package and the results of an experiment testing its performance. Finally, we discuss how we might build a deployable version of our approach.

II. RELATED WORK

Detecting food behaviors had been of great interest to the research community, one of the earliest examples is Gellersen et al.’s [8] MediaCup. MediaCup was a coffee cup augmented with an accelerometer and temperature sensor along with infrared communication. Butz et al. [22] created a similar system except using an augmented beer placemat for interacting with drinking cups.

Chang et al. [3] developed a dining table made up of multiple weighing and RFID reader surfaces capable of tracking the movement of tagged food containers and interactions with them. Amft et al. [16] have worked with a number of different sensing modalities such as neck based EMG for swallowing detection, bone conduction microphones for chewing detection, and inertial sensing for detecting intake (eating) gestures. They

also showed that it was possible to distinguish 19 different food items with 80% accuracy using chewing sounds alone. Man-koff et al. [14] used grocery receipts as a proxy for measuring food intake while Patterson et al. [17] used RFID tagged objects to infer cooking and food preparation activities. Kranz et al. [11] have used an augmented kitchen with microphones, a load sensing knife, and smart cutting board to classify six different foods being cut with 85% accuracy.

All of these systems use sensors on the user or the environment to infer the foods people are eating. Another approach has been to use camera images to track what foods people have been eating. Wu and Yang [20] developed a system which uses a wearable web camera to recognize the fast food items people are eating. They are able to recognize food from 9 different restaurants with 73% accuracy. In the nutritional field the remote food photography method (RFPM) has been used to help supplant food surveys and recall studies. Participants take pictures of the food before and after they eat, and these pictures are shown to trained clinicians who analyze the pictures by matching them up to a library of stored foods and estimating their size. Martin et al. [15] tested their system by providing a cooler with pre-weighed foods to 52 subjects, who consumed the food, and showed -5.1% and -6.6% accuracy in laboratory and free living conditions (consistent underestimates).

While these approaches share our goal of monitoring consumption, the most similar technical work comes from the fields of food chemistry and materials analysis. Because food production is highly automated and controlled there are many advanced systems targeted at very specific applications. For example, using microwave transmission to measure the moisture in salted-butter [19]. In contrast to the majority of these systems, which measure particular features of a known food, we seek to identify liquids. In this realm, Liu et al. [13] developed a system which used visible and near infra-red spectroscopy to discriminate between 5 different types of instant tea using a neural network with 98.7% accuracy. Ketola et al. [10] used a mass spectrometer to identify 13 different cola beverages with 93.8% accuracy. Sádecká et al. [18] used fluorescent spectroscopy to differentiate 13 different brandies and 30 different wines. While most of the approaches rely on expensive hardware in a controlled laboratory environment, we sought to develop an approach that could detect what liquid a person is drinking simply by pouring it into our cup

Another area of similar research interest is the development of chemical or electronic nose sensors. One of the main disadvantages of electronic noses sensors is that they are very expensive. Zhang et al. [24] use a different approach to create inexpensive electronic nose sensors to classify 14 different soft drinks with <2% misclassification rate. Rather than using a MEMS sensor, Zhang uses an array of 25 chemically responsive dyes in a ~1cm square. When the dyes are exposed to a liquid they change color, depending upon the signature of the color changes in the 25 dyes the liquids can be distinguished from one another.

III. SENSING BEVERAGES

We are usually very good at determining what liquid is in our glass, because we almost always know, a priori, what we poured into it. When we are presented with an unknown drink

| | | |
|------------------------|-------------|-------------------------------------|
| Milks | 1 | 1% Chocolate Milk |
| | 2 | 2% Reduced Fat Milk |
| | 3 | Fat Free Milk |
| Talking Rain | 4 | Talking Rain - Black Raspberry |
| | 5 | Talking Rain - Ice orange mango |
| | 6 | Talking Rain - Lemon Lime |
| | 7 | Talking Rain - Pink Grapefruit |
| | 8 | Talking Rain Natural |
| | 9 | Talking Rain Peach Nectarine |
| | 10 | Talking rain - kiwi strawberry |
| Gatorade | 11 | Gatorade Berry Rain |
| | 12 | Gatorade Blue Berry Low Calorie |
| | 13 | Gatorade Fruit Punch |
| | 14 | Gatorade Lime Rain |
| Coke | 15 | Cherry Coke |
| | 16 | Coke |
| | 17 | Coke Zero |
| | 18 | Coke Zero Vanilla |
| | 19 | Diet Chery Coke |
| | 20 | Diet Coke |
| | 21 | Diet Coke Caffeine Free |
| | 22 | Diet Coke Lime |
| | 23 | Diet Coke with Splenda |
| Additional Soft Drinks | 24 | AW RootBeer |
| | 25 | Diet AW Root Beer |
| | 26 | Dr Pepper |
| | 27 | Diet Dr Pepper |
| | 28 | Diet Barqs Root Beet |
| | 29 | Pepsi |
| | 30 | Diet Pepsi |
| | 31 | Caffeine Free Pepsi |
| | 32 | Mountain Dew |
| | 33 | Diet Mountain Dew |
| | 34 | Diet Mug Root Beer |
| | 35 | Mug Root Beer |
| | 36 | Fresca - Original Citrus |
| | 37 | Ginger Ale |
| | 38 | Orange Crush |
| | 39 | Orange Juice |
| | 40 | Sprite |
| 41 | Sprite Zero | |
| Beers | 42 | Alaskan Amber |
| | 43 | Fat Tire |
| | 44 | Heineken |
| | 45 | Pilsner Urquell |
| | 46 | Red Bull |
| Wines | 47 | Redhook ESB |
| | 48 | Samuel Adams Octoberfest |
| | 49 | Stella Artois |
| | 50 | Chardonnay |
| Energy Drinks | 51 | Merlot |
| | 52 | Sauvignon Blanc |
| | 53 | Shiraz |
| Juices | 54 | Monster Lo-Carb |
| | 55 | Rockstar Energy Drink |
| | 56 | Apple Juice |
| | 57 | Cranberry Juice |
| | 58 | Grape Juice |
| | 59 | Grapefruit Juice |
| Coffee | 60 | V8 |
| | 61 | Coffee Decaf |
| | 62 | Coffee Decaf - with half and half |
| | 63 | Coffee Regular |
| Other | 64 | Coffee.Regular - with half and half |
| | 65 | Coco |
| Water | 66 | Nestea Lemon Iced Tea |
| | 67 | Distilled Water |
| | 68 | Water |

Table 1. Drinks sampled during data collection using the spectrometer, pH sensor, and conductivity probe.

we primarily use our sense of vision, smell, and taste. Visually we are limited to categorizing drinks based on colors, for example brown soft drinks versus green sports drinks. Flavor, a combination of our sense of smell and taste, is our primary tool for identifying the specific drink. Our sense of smell is a reusable chemical sensor capable of detecting thousands of different compounds; combined with our sense of taste we have a very accurate chemical sensing system to discriminate liquids with. Drinks are also tailored to human tastes, they are designed to be appealing and recognizable, so our sense of flavor is especially well suited to this task.

Unfortunately, there are no readily available chemical sensors which can replicate our sense of flavor. While there have been attempts to develop electronic nose sensors there are none that are cheap, compact, readily available, and easily used. Instead of focusing on replicating our human senses we have to go beyond these senses and use technologies which can help bridge the gap in our technological shortcomings. In doing so we can take advantages of several properties of beverages that make automatic classification possible. First, there are a relatively small number of beverages with reduced variance. The commercial nutritional database, Axxya [23] contains 35,000 items; only 2,200 items (6%) are beverages, a small fraction. Most beverages are mass produced at bottling plants, which have a vested interest in producing similar drinks over long periods of time. Liquids also tend to be homogeneous, meaning

you only need to sample a small portion to determine the makeup of the entire glass. Finally, people tend to drink similar sets of drinks for long periods of time, so most often; we only need to recognize a smaller subset of drinks

To find technologies to assist in our task we looked to the medical science fields to identify technologies which were being used to analyze unknown substances. The most promising technologies were spectroscopy, ion selective electrode (ISE) sensing, and radio frequency (RF) sensing. Our goal of developing a technology which could be integrated into a glass requires sensors that would:

- Be sanitary to use over long periods of time;
- Not require replacement reagents or user cleaning of sampling tubes or chambers;
- Not destroy or make the liquid undrinkable; and
- Have the potential of being deployed cheaply to a large segment of the population

While we investigated RF sensing as well, we found that optical sensing and ISE were sufficient for the number of drinks people usually consume.

IV. PROTOTYPE BEVERAGE SENSING SYSTEM

Our prototype beverage sensing system has two components: the optical spectrometer and the pH/Conductivity probes, illustrated in Fig. 1.

A. Optical Sensors

Spectroscopy is the measurement of radiation (such as light energy) as a function of wavelength. A spectrometer contains a 1D camera which uses optics to disperse incoming light across the surface of the sensor. This allows us to measure the composition of the light, in our spectrometer, the OceanOptics USB4000, our camera is a 3648x1 CCD where each pixel is approximately 0.2nm wide, and our spectral resolution is about 1.5nm. Fig. 2 shows an example spectrum received showing the ~300nm (ultraviolet) to 1100nm (near infrared, NIR) range of our spectrometer. The spectrometer has a variable integration time (time during which light accumulates across the sensor before a reading), from 3.8ms to 65s. Longer times allow more light inside. In our work, we use a 10mS integration time, except for very opaque liquids, such as chocolate milk, where a longer integration time is used to let more light enter the spectrometer. The advantage of using a small integration time is that the system is more responsive and is more immune to ambient lighting. For very short integration times, there is almost no signal present even from very bright ambient room lighting.

While mass spectrometry, x-ray, and nuclear magnetic resonance are all forms of spectroscopy and would provide useful data, only visible and NIR spectroscopy are readily available and appropriate for low-cost portable solutions.

B. Acquiring Spectra from Liquid Samples

Our optical setup mirrors standard spectrometer setups, where liquid is poured into a 1cm square plastic container called a cuvette. One side is illuminated with a controlled light source, light travels through the liquid, and depending upon the chemical composition of the liquid, different parts of the light




Figure 1. (top) Experimental data collection setup using a OceanOptics LS1 tungsten light source, a laser cut plastic sample cuvette holder, and a fiber optic probe connected to our spectrometer. (bottom) pH and conductivity probes collected readings from samples in glass containers.

source’s spectrum are absorbed at various rates. The spectrometer reads the light which passes through the liquid. Our system is constructed using off the shelf parts in a fairly simple configuration. Rather than enclosing the system in a box or using very specific alignment of parts we wanted to make sure our system was robust and did not require precision alignment of the various components. This gives us a better confidence that our system will function outside of laboratory conditions and is robust enough to be used by non-experts.

C. Ion Selective Electrode Sensors

Ion Selective Electrodes (ISE) are sensors which react to specific ions in a solution and allow their concentrations to be measured. For example, pH is a measurement of H⁺ ions in a solution. The pH range is from 0 (very acidic) to 14 (very alkaline) with 7 being a neutral pH. pH is measured by the amount ion interactions with the electrode, the reading is calibrated against a known sample inside the probe.

One drawback with ISE probes is that they do require ions to be exchanged across a membrane. For pH probes this membrane is usually glass which is simple to keep clean, long last-




Figure 2. (left) Example spectrum of a tungsten light source (top curve) from a Vis-NIR spectrometer. The bottom curve is the output after the light source passes through a sample of water. The received spectrum ranges from 345nm to 1040nm. (right) USB4000 spectrometer used to capture the spectrum and measures 89mm x 63 mm x 34mm.

ing, and durable. For other ISE probes a crystal disk or plastic membrane, is used. These probes tend to have limited lifetimes - they can only be exposed to a certain number of solutions before they must be replaced - and are slightly less sanitary. Because of the limited life spans of crystal and polymer membrane ISE sensors we only use pH measurements in our system. The one short coming of glass pH probes is that they usually must be stored in a buffer solution to avoid drying out. However, it is feasible that probes could be made robustly enough to function without being stored this way. In addition there are other optical based methods for measuring pH which could reasonably last long enough [12] and estimates suggest these pH sensors can be manufactured for pennies.

The other electrode sensor used in our system is not necessarily an ion selective sensor, but instead a probe which measures the ability of the solution to conduct an electric current. Conductivity provides some information on the salinity is in the drink; for example, Coke is much more conductive than milk (due to its high sodium content, not its sugar as is often believed. Sugar is a non-electrolyte, so it does not conduct electricity very well in water; salt does however).

The pH and conductivity probes are simply immersed in liquid samples to measure their respective properties. Separate containers are used for convenience because the pH and conductivity probes cannot be both sampling at the same time in the same liquid, otherwise the two would emit signals that would confuse one another. When integrated into a glass the two sensors would take turns, one would sample while the other would sit idle.

D. Data Processing

Our light source does not illuminate the entire range of our spectrometer, we remove portions of the spectrum that do not contain useful information. For the Ocean Optics LS1 light source, this ends up being between 433nm to 988nm. This range is found by examining the emission spectrum of the light source, on its own, and finding the range which contains 95% of the energy. Depending on the emission spectrum of different light sources, a different range might be more appropriate. Using this subset of wavelength data we correct for intensity variations by aligning each signal to a stored free air reference signal from our light source. This alignment helps to counteract problems with slight misalignments in the optical system and changes in sample concentrations (for thicker liquids like V8 or coffee). The normalization helps make the system more immune to variations in samples; however, it does throw out some useful information about how much light actually was absorbed by the liquid, rather than the shape of the absorption curve.

We select 974 bins from the spectral data and use them as features for our classifier. We also use feature weighting to store an equal number of pH/conductivity features. The pH/conductivity data is discretized into bins of increasing size to avoid over fitting to specific pH/conductivity readings.

Data from the primary data collection experiment and our prototype sensor package was organized and classified using the Naïve Bayes classifier in the Weka machine learning software. Although Weka offers several classifiers, each produces roughly similar results and rather than picking the classifier

with highest performance, we selected Naïve Bayes as it provides an effective, simple, and common classifier which is well suited to modeling the distributions of large sets of features and which requires relatively small amounts of training data.

V. PRIMARY DATA COLLECTION

To test if our system could distinguish different drinks from one another we collected the spectrum, pH, conductivity, and temperature of 68 different beverages, listed in Table 1. We collected two sets of data on different days, all beverages except the Gatorade, Energy Drink, Wine, and Beer drinks were collected from new cans of each drink. For the other groups, we resealed the containers to see if beverages left overnight would change significantly enough to cause our system to fail, no such drastic change was observed.

A. Experimental Protocol

We refer to our two datasets as Run 1 and Run 2. The drinks were randomly ordered and sampled for Run 1. Run 2 used the same order but in reverse. Drinks were processed in groups of approximately 6 drinks. As each sample was processed it was opened, and ~1.5mL of each drink was placed into a plastic cuvette. The process for collecting the spectra data from the spectrometer was as follows:

1. selected the sample from a menu
2. recorded the temperature of the liquid in the cuvette
3. before placing the liquid sample cuvette into the holder the spectra of light emitted from the light source was recorded
4. the sample was placed into the sample holder as show in Fig. 1
5. the user indicated that the sample was inserted and ready and the system began an automated data collection process. 10 blocks of samples were recorded of the light spectrum of the sample. Each block consisted of 500 spectra using a 10 millisecond exposure
6. each block took approximately 5 seconds and in all around 1 minutes worth of data was collected taking approximately 3 minutes with processing and data saving in between each block
7. the data collector then removed the sample and a post empty spectrum was recorded from the light source

After completing this procedure the data collection for that sample was complete. The data collector was presented with a view of the raw spectrum collected and asked to review any spectra which were quite dissimilar from the rest. Samples varied for primarily two reasons 1) bubbles from carbonation and 2) subtle variations in liquids like V8, grapefruit juice, and very dense liquids like cocoa. Opaque samples like coffee, cocoa, and V8 had an additional data collection run where an integration time longer, varying from 500ms to 1s was used. This additional data was collected because the 10ms exposure had such a small amplitude.

Bubbles caused by carbonation appeared as variations in the intensity of the sample, for example if a bubble floated through the section of the sample being analyzed there would




Figure 3. Conductivity and pH values for drinks collected in Run 1 and Run 2. The difference between recorded pH values was 3.1% (std. dev. 3.1%). The difference between recorded conductivity values was 8.8% (std. dev. 17.8%).

be dip in the intensity of light received (as the bubble typically scattered the light). Carbonation bubbles were not considered a challenge for the system to overcome because they are easily detected by variations in the signal. In a deployed system they can be detected and avoided if necessary and can be physically eliminated using an ultrasonic actuator to automatically ‘tap’ the wall of the glass to remove bubbles in front of the sampling window, if necessary. For samples where variations were found, we avoided manually picking measurements to get a consistent sample. Instead we ran through our data collection procedure a second or third time to ensure we were getting consistent readings of the current sample and used them as is. Fig. 4 contains an overlaid plot of the various light spectrums recorded for the first data set.

For pH and conductivity data, we placed approximately 25mL of each liquid in a small glass beaker and placed the pH and conductivity probes inside. We waited a few seconds for the reading to stabilize and then began collecting data from pH sensor, and then the conductivity sensor. Conductivity readings tended to vary over time, with the reading increasing slightly if the probe was stirred. pH readings were usually very consistent and stabilized within a few seconds of being placed inside the liquid.

The only sample for which the pH reading was inconsistent was distilled water. pH values are actually not valid for substances which are too pure, substances with very few free ions to measure result in nonsensical readings from the pH sensor. Samples with conductivities below 10µSiemens¹ (µS) were considered distilled/pure water and were set to a pH of 7 (neutral). From our data set, only distilled water samples required this modification, with a measured conductivity of 6µS and 2µS. The next highest value was tap water with a conductivity

of 56µS and 886µS. Fig. 3 shows the pH and Conductivity data recorded from Run 1 and Run 2.

VI. RESULTS

Table 2 shows the confusion matrix for our drink classification scheme. We achieved an overall accuracy of 79.4% with a precision and recall of 77.5% and 79.4%, respectively. These results are especially encouraging considering that random chance is 1.5% for a 68-class problem.

Upon closer inspection, many of the misclassifications occurred with similar drinks of that class, such as the root beers. Interestingly, we were still able to classify among different types of drinks. For example, we could see the difference between coke and diet coke, which both appear visually similar, but have very different nutritional content. Moreover, there was strong predictive power in differentiating between the juices and milks.




Figure 4. Plot of all 68 beverages’ spectrum recorded for Run 1.

¹ Water with conductivity <10µS is considered distilled water by the National Committee for Clinical Laboratory Standards




Table 2. (left) Confusion matrix for drink classification using leave-one out training with training data from Run 1 and test data from Run 2. Overall accuracy is 69.1% with precision of 68.7% and recall of 69.1%. (right) Confusion matrix using leave-one out training with training data from Run 2 and testing data from Run 1. Each run contains 10 separate measurements of the same sample, the coloring indicates the number of classified measurements in a particular cell. Overall accuracy is 79.4% with precision of 77.5% and recall of 79.4%. Random chance is 1 out of 68 or 1.5%. The ~10% accuracy between the left and right confusion matrix is caused by additional confusion with ‘Sauvignon Blanc’, ‘distilled water’, and slightly poorer performance with ‘Diet Coke with Splenda’.

From our initial results, we observed we were not able to differentiate between caffeinated (diet Coke) and non-caffeinated drinks (caffeine free diet Coke), because they typically have very similar chemical compositions. Caffeine is also a tasteless, colorless, and odorless substance that does not manifest itself in our optical or conductive features space. More work is needed in identifying a low-cost and practical approach to automatically detecting caffeine. The closest commercial system capable of performing caffeine detection requires a spectrometer with range much further into the infrared (up to 2500nm as opposed to our current instruments 1040nm range) [2]. Miniature spectrometers with capabilities in this range are not common because there is relatively little demand for such systems. Also at higher infrared wavelengths more and more light is absorbed by water, making it more difficult to analyze liquid samples. This problem can be overcome by using a narrower interaction field (decreasing the amount of water light passes through before reaching the sensor) or increasing IR

light output using high intensity LEDs or laser diodes.

We also found that all of the features were important in classifying drinks. The use of pH only and resistivity only features yielded an accuracy of 11.62% and 31.02%, respectively. The optical data definitely appeared to be the most critical in the classification scheme by being able to discern over half of the drinks when used alone.

A. Clustering Drinks

If we are only interested in what logical group a drink belongs to we can use a simpler classifier to cluster our drinks together into groups. Instead of recognizing a particular brand of soft drink we simply care about soft drinks as a group. Using the data from our data collection run we grouped the drinks into a dozen categories and classified them as shown in Table 3. Average accuracy is 77.2% with a great deal of confusion between diet and regular soft drinks. If we combine these two groups together our accuracy increases from 77.2% to 81.9%.

VII. PROTOTYPE SENSOR PACKAGE

While the primarily data collection using the OceanOptics spectrometer showed that we could achieve promising results from a wide variety of liquids, we would eventually like this system to be integrated into something which costs a few dollars or ideally less than a dollar. The cheaper the components of the system the more likely it can be widely used and accepted by people throughout the world.

The spectrometer that we currently use costs about \$3,000 USD. The majority of these costs are from engineering and support and not actual parts costs. The basic components of a spectrometer are a linear imager (like those found in flatbed scanners), a few mirrors, and a diffraction grating. These parts

| | Beer | Coffee | Diet Soft Drink | Energy Drink | Gatorade | Juice | Milks | Other | Soft Drink | Talking rain | Water | Wines |
|-----------------|------|--------|-----------------|--------------|----------|-------|-------|-------|------------|--------------|-------|-------|
| Beer | 80 | | | | | | | | | | | |
| Coffee | | 40 | | | | | | | | | | |
| Diet Soft Drink | | | 90 | 10 | | | | | 10 | | | |
| Energy Drink | | | | 20 | | | | | | | | |
| Gatorade | | | | | 28 | | | | | | | 12 |
| Juice | 9 | | | | | 30 | | 5 | 5 | | | 11 |
| Milks | | | | | | | 30 | | | | | |
| Other | | 3 | | | | | | 17 | | | | |
| Soft Drink | | | 28 | 20 | | | | | 80 | 22 | | |
| Talking rain | | | | 3 | | | | | 7 | 60 | | |
| Water | | | | | | | | | | | 10 | |
| Wines | | | | | | | | | | | | 40 |

Table 3. Clustering results from our primary data collection. The average accuracy is 77.2% with a precision of 81.7% and recall of 77.2%.

already cost under \$50 in large quantities and could be made for much less. However, one alternative to taking existing spectrometers and making them less expensive, is to make a similar sensor using cheap discrete components. We constructed a prototype sensor package made up of a color sensor (Avago ADJD-S71) and 8 different discrete LEDs that cost less than US\$10, shown in Fig. 5 along with the pH and conductivity probes. Our sensor package uses UV, IR, and visible LEDs to get a characteristic fingerprint of the liquids absorption and transmission. Whereas the spectrometer used light intensity for different wavelengths, we measured the intensity output from the color sensor (Red, Green, Blue, and Clear) with different combinations of our 8 LEDs illuminating the sample. Using our prototype sensor package we collected an example data of set of 10 liquids, shown in Table 4. The 10 liquids were chosen to represent a cross section of our original data set and what someone might drink during the week, with a few difficult to classify variations (i.e. Coke vs. Diet Coke) chosen to ensure the subset contained some more difficult comparisons.

We performed 4 data collection runs for each liquid, each time recording data from the prototype sensor, pH, and conductivity sensor. The average accuracy of the sensor was 60% with a precision of 52% and recall of 60%. While reduced in scale and complexity, this is comparable to our performance using the full spectrometer data alone to classify our 68 drink data set, which achieves an average accuracy of 52%. Adding in pH and conductivity data to our prototype data set boosts the accuracy tremendously to 97.5% due to the limited number of possible liquids. This effect would disappear as the number of liquids increased; which would make the pH and conductivity ranges more crowded as illustrated in Fig. 3.

VIII. DISCUSSION

We have described our approach of coupling an optical spectrometer together with pH and conductivity probes in order to sense and classify different fluids that a user may consume. Results from our first experiment demonstrate the feasibility of using relatively high end commercial off-the-shelf sensors and employing a fingerprinting methodology. Encouraged by these results, we built a prototype sensor package to investigate whether this would hold true with a cheaper and less well calibrated implementation of the spectrometer. Results from that experiment suggest that this is indeed possible as well.

Although the aim of this work was to show the feasibility of automatically classifying drinks among a collection of known fingerprints, we have carefully considered how one might build a fully deployable system in the future. We started by creating design sketches of different objects in which one may embed such sensors. Ideas ranged from stirrers to artificial ice cubes to cups and mugs themselves. For various pragmatic reasons, we settled on the cup being the most practical for everyday use.

Interestingly, the vision of an instrumented cup is not too far reaching. Manufactures such as Hamamatsu, Inc. [1] are already building spectrometers in integrated circuit form. These mini-spectrometers are only a few centimeters on each side, which could be integrated into the bottom portion of standard 200-300 ml cup. The integration of the pH sensing is also possible through new advances in IR-based detection schemes. This contactless approach is critical for ensuring safety and durability.




Figure 5. (left) Prototype sensor package. The sensor is submerged inside a drink along with the pH and conductivity sensors to record a fingerprint of the liquids response to various illumination sources. (right - top) Avago ADJD-S71 color sensor (right - bottom) LED color array

| | Apple Juice | Coke | Diet Cherry Coke | Diet Coke | Diet Mug Rootbeer | Gatorade - Berry Rain | Gatorade - Blue Berry | Mug Rootbeer | Talking Rain - Lemon Line | Talking Rain - Peach Nectarine |
|--------------------------------|-------------|------|------------------|-----------|-------------------|-----------------------|-----------------------|--------------|---------------------------|--------------------------------|
| Apple Juice | 4 | | | | | | | | | |
| Coke | | 4 | | | | | | | | |
| Diet Cherry Coke | | 1 | 2 | | | | 1 | | | |
| Diet Coke | | | 2 | 1 | | | 1 | | | |
| Diet Mug Rootbeer | | | | 2 | | | 2 | | | |
| Gatorade - Berry Rain | | | | | | 4 | | | | |
| Gatorade - Blue Berry | | | | | | | 4 | | | |
| Mug Rootbeer | | 1 | | | | | | 3 | | |
| Talking Rain - Lemon Line | | | | | | | | | 3 | 1 |
| Talking Rain - Peach Nectarine | | | | | | | | | | 2 |

Table 4. Combined confusion matrix of prototype sensor package. We performed four iterations of leave one out classification (train on 3 test on 1), the above confusion matrix is the sum of those 4 confusion matrices, average classification accuracy was 60% using the prototype sensor data only. Adding in pH and conductivity data increased accuracy to 97.5%

Like most technology, the cost of these sensors will likely drop over time, especially as spectrometers are integrated into mainstream consumer-oriented markets. Cheaper sensors would enable users to have more than one instrumented cup, for example a coffee cup for hot drinks, wine glass, mug, or regular glass. Specialized glasses would also help to make the classification problem simpler (users are much more likely to have wine in a wine glass).

One benefit of instrumenting the cup is that we would also be able to measure volume and weight in order to infer not only the actual amount of the beverage consumed, but also the particular rate and patterns with which it was consumed. We could also add a cheap output modality to the cup itself in order to present information to the user.

That said, we should caution that the current work does have one major limitation, which we hope to resolve in future work. Although our system shows promise for identifying drinks, we do not believe our sensors are sensitive enough to isolate and sense specific chemical compounds. Improved sensors and further analysis might make this possible.

With sensors that can sense specific chemical compounds, we could extend the applications we have motivated this work with, namely the desire to measure beverage consumption for the purpose of health metrics such as calorie management. For example, cups in bars and nightclubs in the future might contain sensors like these that are able to sense certain chemical compounds, such as gamma-hydroxybutyric acid or benzodiazepines, better known as date rape drugs. It would also be useful to be able to sense the alcoholic content of fluids in a cup to know how much one has been drinking. One could also envision building such devices to sense water quality in distribution wells and pipes, or at infrastructure endpoints (e.g. the faucet in your home).

IX. CONCLUSION

We have described a vision for building a sensor rich cup capable of determining the kind of liquid a person is drinking, as well as the opportunities that the success of such sensors may open. Our feasibility experiments show that a novel combination of optical and ion selective electrical pH, and conductivity sensors can enable the sensing and classification of liquids. We described two feasibility experiments that suggest that it is possible to reliably classify specific drinks with up to 79% classification accuracy for 68 different drinks. Although more verification and larger datasets are needed these results are very promising. We hope that researchers will begin to explore this kind of approach for automatically logging daily fluid intake and eventually spur interest in technology to support automatic food journaling. Although our work is based on the use of commercial test equipment, we feel that building a practical version of this approach is not too far reaching, especially when considering many of the necessary sensors needed for this technology are already becoming available in small form factor. We have already shown that is possible to build a low-cost version of an instrumented cup using readily available electronics costing under \$10 that can identify up to 10 drinks.

ACKNOWLEDGMENT

The authors would like to thank Neil Emerton, John Lewis, John Lutian, and Tim Large in the Microsoft Applied Sciences group for their generous assistance and the use of their equipment. We would also like to thank Gerald DeJean, William Dangel, and Mike Sinclair for their assistance.

REFERENCES

- [1] Hamamatsu Photonics Mini-spectrometer <http://jp.hamamatsu.com/products/sensor-ssd/pd186/4029/C10988MA/>.
- [2] CAMO Software. Determination of caffeine in decaffeinated coffee by NIR spectroscopy. http://www.camo.com/downloads/resources/application_notes/decaffeinated_coffee_NIR_spectroscopy.pdf.
- [3] Chang, K., Liu, S., Chu, H., et al. The Diet-Aware Dining Table: Observing Dietary Behaviors over a Tabletop Surface. In *Pervasive Computing*. 2006, 366-382.
- [4] Consolvo, S., McDonald, D.W., Toscos, T., et al. Activity sensing in the wild: a field trial of ubifit garden. *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, ACM (2008), 1797-1806.

- [5] Duffey, K.J. and Popkin, B.M. Shifts in Patterns and Consumption of Beverages Between 1965 and 2002[ast][ast]. *Obesity* 15, 11 (2007), 2739-2747.
- [6] Duncan, G.E., Sydemann, S.J., Perri, M.G., Limacher, M.C., and Martin, A.D. Can sedentary adults accurately recall the intensity of their physical activity? *Preventive Medicine* 33, 1 (2001), 18-26.
- [7] Fowler, S.P., Williams, K., Resendez, R.G., Hunt, K.J., Hazuda, H.P., and Stern, M.P. Fueling the Obesity Epidemic[quest] Artificially Sweetened Beverage Use and Long-term Weight Gain. *Obesity* 16, 8 (2008), 1894-1900.
- [8] Gellersen, H., Beigl, M., and Krull, H. The MediaCup: Awareness Technology Embedded in an Everyday Object. In *Handheld and Ubiquitous Computing*. 1999, 308-310.
- [9] Kruger, J., Galuska, D.A., Serdula, M.K., and Jones, D.A. Attempting to lose weight: specific practices among U.S. adults. *American Journal of Preventive Medicine* 26, 5 (2004), 402-406.
- [10] Ketola, R.A., Heikkinen, J., Piepponen, S., Lauritsen, F.R., and Kotiaho, T. Classification of cola beverages on the basis of mass spectra measured by membrane inlet mass spectrometry. *Rapid Communications in Mass Spectrometry* 12, 15 (1998), 1011 - 1018.
- [11] Kranz, M., Maldonado, A., Rusu, R.B., et al. Sensing technologies and the player-middleware for context-awareness in kitchen environments. (2007).
- [12] Lin, J. and Liu, D. An optical pH sensor with a linear response over a broad range. *Analytica Chimica Acta* 408, 1-2 (2000), 49-55.
- [13] Liu, F., Ye, X., He, Y., and Wang, L. Application of visible/near infrared spectroscopy and chemometric calibrations for variety discrimination of instant milk teas. *Journal of Food Engineering* 93, 2 (2009), 127-133.
- [14] Mankoff, J., Hsieh, G., Hung, H.C., Lee, S., and Nitao, E. Using Low-Cost Sensing to Support Nutritional Awareness. *Proceedings of the 4th international conference on Ubiquitous Computing*, Springer-Verlag (2002), 371-376.
- [15] Martin, C.K., Han, H., Coulon, S.M., Allen, H.R., Champagne, C.M., and Anton, S.D. A novel method to remotely measure food intake of free-living individuals in real time: the remote food photography method. *The British Journal of Nutrition* 101, 3 (2009), 446-456.
- [16] Oliver Amft, G.T. On-Body Sensing Solutions for Automatic Dietary Monitoring. 2009. <http://www2.computer.org/portal/web/csdl/doi/10.1109/MPRV.2009.32>.
- [17] Patterson, D.J., Fox, D., Kautz, H., and Philipose, M. Fine-Grained Activity Recognition by Aggregating Abstract Object Usage. *Proceedings of the Ninth IEEE International Symposium on Wearable Computers*, IEEE Computer Society (2005), 44-51.
- [18] Sádecká, J., Tóthová, J., and Májek, P. Classification of brandies and wine distillates using front face fluorescence spectroscopy. *Food Chemistry* 117, 3 (2009), 491-498.
- [19] Shiinoki, Y., Motouri, Y., and Ito, K. On-line monitoring of moisture and salt contents by the microwave transmission method in a continuous salted butter-making process. *Journal of Food Engineering* 38, 2 (1998), 153-167.
- [20] Wen Wu and Jie Yang. Fast food recognition from videos of eating for calorie estimation. *IEEE International Conference on Multimedia and Expo, 2009. ICME 2009.*, (2009), 1210-1213.
- [21] Wing, R.R. and Phelan, S. Long-term weight loss maintenance. *Am J Clin Nutr* 82, 1 (2005), 222S-22G.
- [22] Butz, A. and Schmitz, M. Poster: Design and Applications of a Beer Mat for Pub Interaction. *Proceedings of the Seventh International Conference on Ubiquitous Computing*.
- [23] Nutritionist Pro TM Nutrition Database. <http://www.nutritionistpro.com/knowledgebase.php>.
- [24] Zhang, C. and Suslick, K. S. Colorimetric Sensor Arrays for Soft Drink Analysis. In *Proceedings of the Journal of Agricultural and Food Chemistry*, 2007, 55, 237-242.