AquaCat: Radar and Machine Learning for Fluid and Powder Identification

David Morrison
SACHI, School of Computer Science, University of St Andrews, Scotland, UK
dm236@st-andrews.ac.uk

David Harris-Birtill
SACHI, School of Computer Science, University of St Andrews, Scotland, UK
dhccb@st-andrews.ac.uk

Alasdair N Houston
SACHI, School of Computer Science, University of St Andrews, Scotland, UK
anh4@st-andrews.ac.uk

Aaron Quigley
SACHI, School of Computer Science, University of St Andrews, Scotland, UK
aquigley@st-andrews.ac.uk

Abstract
AquaCat is a low-cost radar-based system capable of discriminating between a range of liquids and powders. This work builds on techniques developed for RadarCat, a similar radar-based system that has shown high accuracy in the classification of physical objects.

AquaCat records radar signals returned from various substances, building up a dataset of labelled signal traces. These are used as input into a machine learning system that produces a predictive model capable of identifying new examples of these substances. We present a study assessing the accuracy of the system in recognizing water, various sugars and sugar solutions, iron sulphate, and various alcoholic spirits. For this test set, AquaCat achieved a mean accuracy of 78.33% and correctly identified the powders 100% of the time. We see applications of the technology in fields such as chemical manufacturing, public health, and environmental monitoring.

Author Keywords
Radar; Liquid Classification; Powder Classification; Pollutant Detection; Machine Learning.

ACM Classification Keywords
H.5.m. Information interfaces and presentation (e.g., HCI): User interfaces – Input devices and strategies;

Figure 1 – Polystyrene cuvettes with liquid and powder samples.

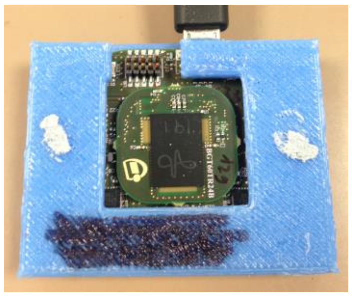

Figure 2 – Infineon BGT60 based radar system in custom housing.

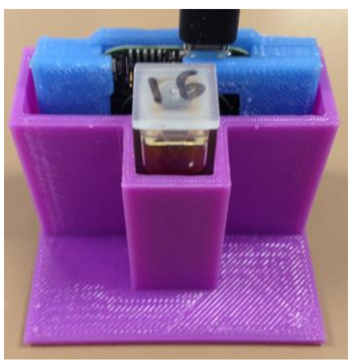

Figure 3 – Custom mount holds sample and radar in place to reduce differences between readings.

Introduction

The AquaCat project makes use of low-cost miniaturized radar technology and machine learning to build a system that can reliably discriminate between different liquids and powders. The project builds on RadarCat, a radar-based object classification system developed in the University of St Andrews School of Computing's Human-Computer Interaction Group (SACHI) [1]. This paper presents an extension of the capacities of such as system to the discrimination of liquids and powders. This extension has a wide range of applications in domains such as; control systems in food and pharmaceutical production, the non-destructive identification of drugs and alcohol (in both law enforcement and public health), the monitoring of water pollution levels in rivers, and the automated assessment of drinking water quality.

In this position paper, we will discuss the initial results of a study designed to test the capacities of a system like RadarCat, repurposed to classify liquids and powders. As far as we are aware this is the first attempt to use radar technology with machine learning to perform such a task. As such;

1. We outline the structure of a system designed to classify liquids and powders.
2. We present the results of a study reporting the level of accuracy achieved by the system when training and classifying different liquids and powders.

We conclude this paper with a brief discussion of the results and the possible applications of the technology.

AquaCat uses a radar systems from Infineon Technologies based on their BGT60 millimeter-wave RF transceiver. During development, we also experimented with Google Soli [3]. Initial testing suggests the Soli exhibits similar levels of accuracy and patterns of confusion to the BGT60 based system, however a detailed comparison is beyond the scope of this study. For the work present here, the BGT60 based system was selected for reasons of availability. Both systems are based on Frequency Modulated Continuous Wave (FMCW) radar principles [4] and use identical frequency bands. We expect the results of this study to be reproducible using the Soli.

Background Work

RadarCat [1] demonstrates the ability of millimeter-wave radar systems when combined with machine learning to discriminate between a wider range of physical objects. Part of that work hinted at the possibility of extending the system to discriminate different kinds of liquids non-destructively. RadarCat is based on Google's Soli [3], a system initially designed for capturing subtitle figure motions for use in gesture recognition. It is expected that Soli, or similar technologies, will be embedded into a wide range of devices in the near future, for example in mobile phones, making miniaturized radar a ubiquitous and cheap technology.

Current liquid sampling and analysis techniques, such as inductively coupled plasma mass spectrometry (ICPMS) [5], require expensive specialist equipment and are very often destructive in nature. Similar spectroscopic methods, such as near-infrared spectroscopy [6], are often employed to analyses pharmaceutical and food products.

Methodology

Our experiment was designed to test the capacity of a radar and machine learning based system to discriminate between different liquids and powders. In selecting our test substances, we applied the following criteria; the substances must be safe to work with in a non-specialist laboratory setting, the substances must be readily available, and the substances must have relevance to a proposed application domain.

Applying these constraints, for liquids we selected pure water, dextrose solutions, sucrose solutions, and three kinds of alcoholic sprit (vodka, whiskey, and rum). These were selected for their relevance in public health applications. Aqueous solutions of dextrose and sucrose were each tested at two different concentrations: 1-Molar (i.e. 2.5millimoles dissolved in pure water to achieve a total sample volume of 2.5ml) and at their respective saturated concentrations. For powders, we tested dextrose, sucrose, and iron sulphate – a common garden fertilizer, selected as a possible pollutant that was safe to handle.

Standardized polystyrene cuvettes, with a capacity of 2.5ml, were used as containers for the test liquids (figure 1 shows these labelled and arranged into a matrix for testing). To accurately characterize the liquid samples, with a consistent radar background signal, a mount was designed and 3D printed (show in figure 2). This mount holds the radar system and a single cuvette tightly in a relative position that can be easily reproduced as different samples are measured.

Data collection was undertaken by sampling 4-channels from the Infineon radar system. Each channel represents the signal received by one of the systems four receive antenna, labelled Rx1, Rx2, Rx3, and Rx4. The system continuously emits a radar wave, repeatedly ramping its frequency from 57-64GHz. These individual frequency ramps are called Chirps. We configured the system to emit 200 chirps-per-second (providing a good tradeoff between temporal-resolution and system heat). The signal returned from each chirp is convolved with its outputted signal to produce a signal showing the differences at various time points during the frequency ramp. This is known as the intermediate frequency and referred to as the raw signal. A line graph showing a representative trace of this signal can be seen in figure 4.

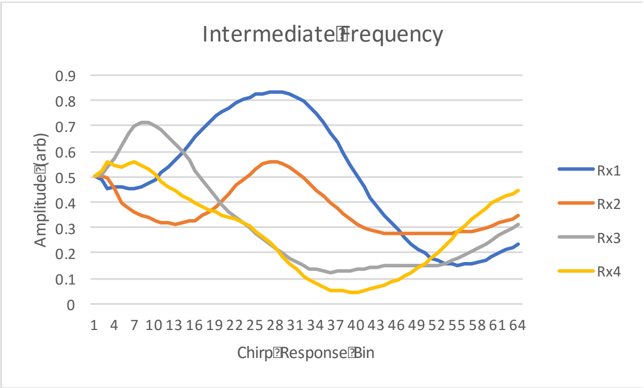

Figure 4 - Example of 4-channel intermediate frequency function from the Infineon radar system.

	Predicted Substance											
Actual Substance	Air	Water	Dextrose SS	Dextrose 1M	Sucrose SS	Sucrose 1M	Dextrose Dry	Sucrose Dry	Iron(II) Sulfate Dry	Vodka	Whiskey	Rum
Air	100.00%											
Water		100.00%										
Dextrose SS			73.33%		13.33%							13.33%
Dextrose 1M		20.00%		80.00%								
Sucrose SS					60.00%	33.33%						6.67%
Sucrose 1M			6.67%			53.33%				13.33%	6.67%	
Dextrose Dry							100%					
Sucrose Dry								100%				
Iron(II) Sulfate									100%			
Vodka			13.33%							73.33%	13.33%	
Whiskey											100.00%	
Rum			13.33%		60.00%	6.67%					20.00%	0.00%

Figure 4 - Confusion matrix showing the results of our study. Rows show which substances were predicted by AquaCat when it was shown actual substances. Empty cells indicate that substance was predicted 0 times.

These signals were recorded 100 times in 0.5s to produce one sample. This process was repeated 30 times for each class of substance. This was done to ensure that enough samples were recorded to compensate for small differences in positioning between sampling for training and sampling for recognition. Each sample was associated the name of the liquid that was sampled and this was used create the set of classes for use during machine learning.

When training AquaCat, we spilt our dataset into training data and validation data. 90% of samples (324) were used to train a neural network with 3 hidden layers, each containing 128 artificial neurons. The system was implemented using the Scikit Learn [7] machine learning toolkit for Python. The remaining 10% (36) were used as a validation set – to automatically assess the performance of the network. This process was repeated four times with different training and validation sets and the best model selected.

Once the network had been trained, it’s accuracy was assessed manually by having it predict the labels for the samples in on an independently produced testing set. The test set consisted of 15 samples for each class of substance (180 samples total). These were generated using new samples of the substances in different cuvettes than the training data. These results are shown in the next section.

Results

Figure 4 is a confusion matrix [8] showing the results of this preliminary study. As can been seen from the diagonal axis, the system performed with high accuracy when discriminating between the powders and most of the liquids. The system achieved a mean accuracy of 78.33%. Note that the powders were correctly discriminated 100% of the time.

Discussion and Conclusion

These initial results suggest some interesting capacities and limitations. The empty cuvette (labelled air) and pure water were correctly identified every time. This suggests that detecting the presence of any liquid or impurities in water might be possible with high accuracy. However, the sugar solutions and alcoholic sprits were often confused suggesting that identifying specific impurities or pollutants might be challenging.

The system showed high accuracy in discriminating the powdered forms of dextrose, sucrose, and iron sulfate. This may be due to the powders having different grain sizes and thus having different substance specific scattering and absorption properties. Further physical and chemical testing is required to establish this.

High accuracy non-destructive powder classification could have important applications in chemical and drug production to control manufacturing systems or as part of a quality control process. The non-destructive identification of illicit drugs has applications in public health and law enforcement. We might envision a deployment of the AquaCat technology being used to perform rapid testing on unknown substances as part of a customs process. Another possibility would be to use the system for the rapid testing of urine samples. This could be expanded to a network of low-cost sensors in a city's wastewater system, allowing public health agencies to monitor drug use and map it geographically.

One possible long-term application is to allow environmental researchers to make cheaper, faster, and thus more numerous measurements of water pollution levels in the field and therefore create a richer data set. When miniaturized radar technologies are deployed into smartphones and other devices, AquaCat could be available to a much larger section of the population. This will enable the collection of crowd-sourced pollution data over a large geographic area and over a long time. By having a higher number of samples over a longer period, policymakers will gain a better understanding of the distribution of pollutants and their sources, allowing them to better target regulatory interventions and infrastructure improvements.

In the future, we intend to test the system with a wider range of chemicals, both in liquid solutions and in their powdered form. We will be looking to measure both the accuracies with which different chemicals can be detected, and at what concentrations accurate detection is possible. With a wider map of the possible substances and concentrations, we will be able to narrow down the possible use-cases and move onto creating specific deployable solutions.

References

1. Yeo H.S., Flamich G., Schrempf P., Harris-Birtill D., and Quigley A. 2016. RadarCat: Radar Categorization for Input & Interaction. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (UIST

'16). ACM, New York, NY, USA, 833-841. DOI: https://doi.org/10.1145/2984511.2984515

2. Infineon Technologies. 2017. BGT60. Retrieved on 16 July 2017 from https://www.infineon.com/cms/en/product/rf-and-wireless-control/mm-wave-mmic/mmw-backhaul-and-fronthaul/BGT60/productType.html?productType=5546d46247342c630147918436016a3b.
3. Lien, J., Gillian, N., Karagozler, M.E., Amihood, P., Schwesig, C., Olson, E., Raja, H. and Poupyrev, I., 2016. Soli: Ubiquitous gesture sensing with millimeter wave radar. *ACM Transactions on Graphics (TOG)*, *35*(4), p.142.
4. Lesurf, J. 2017. Improved forms of radar.. Retrievied on 16 July 2017 from https://www.st-andrews.ac.uk/~www_pa/Scots_Guide/RadCom/part15/page2.html.
5. Garbarino, J.R. and Taylor, H.E., 1996. Inductively coupled plasma-mass spectrometric method for the determination of dissolved trace elements in natural water (No. 94-358). US Geological Survey.
6. Williams, P. and Norris, K., 1987. Near-infrared technology in the agricultural and food industries. American Association of Cereal Chemists, Inc.
7. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V. and Vanderplas, J., 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, *12*(Oct), pp.2825-2830.
8. Sammut, C. and Webb, G. 2011. Encyclopedia of machine learning. New York, NY: Springer.