ERIC C. LARSON | RESEARCH STATEMENT

My research interests focus on **how signal processing, machine learning, and sensing can support initiatives in environmental sustainability and health management**. My dissertation focuses on how machine learning, both supervised and semi-supervised, can enable fine-grained sensing of water usage and increase awareness of wasteful practices (expected completion date April of 2013). However my research agenda is not limited to sustainability, as is apparent from my diverse research projects and my overarching research questions:

- How can signal processing create and support initiatives in sustainability?
- How can sensing outside the clinic help management and diagnosis of medical ailments?

Figure 1 outlines my research projects as a color-coded map. Figure 1 is intended to illustrate not only the diversity of projects I work on, but also how each project experience contributes to my research agenda in health and sustainability. With this in mind, I will be focusing exclusively on applications in sustainability and health, rather than every project from Figure 1.

My interest in sustainability and sensing naturally bring me to examine ways in which signal processing can help provide insights into efficiency. For instance, I developed single-point sensing systems for monitoring water [8] and gas [4] usage down to the individual fixture or appliance (commonly called disaggregated sensing). For example, providing information like "the shower used a total of 10.7 gallons of hot water today, and the upstairs toilet was used 4 times for a total of 7.3 gallons of cold water." My water disaggregation system, called HydroSense, uses machine learning to infer this information from a single pressure sensor attached anywhere on the plumbing system. I have been able to leverage my experience in signal processing, convex optimization, and evolutionary computation in order to make HydroSense robust and easy to calibrate. This work is highly inter-disciplinary. I often work with water management scientists, policy makers, utilities, different disciplines in computer science (*e.g.*, human computer interaction and machine learning), and different disciplines in electrical engineering (*e.g.*, embedded systems and theoretical signal processing). I believe that this type of collaboration brings focus to the practicality of the research and ensures that diverse aspects of the problem drive the solution.

Outside of my dissertation, I have a growing interest supporting disease management with commonplace sensors, an area which will be a mainstay of my research agenda. I collaborated with a group of pulmonologists to create methods of sensing and counting coughs from the microphone of a mobile phone [16]. This system is currently being used to evaluate patients with chronic lung ailments. Continuing this collaboration, I created a system for performing spirometry (*i.e.*, the most widely accepted medical measure of lung function [20]) using a mobile phone microphone [14]. These systems enable patients to monitor chronic lung ailments outside of the clinic at little or no extra cost, allowing for earlier detection of exacerbations, trending, and ultimately increased survival rates [21,24]. Everything I incorporate into my research agenda follows a common scientific theme: problem driven identification of solutions, coupled with real world deployments and evaluations. Although I give great importance to field deployments, I also understand the need for theoretical understanding and development. As such, I spend a great deal of time optimizing feature extraction and training methods so that they appropriately incorporate physical phenomena, linearity (or lack of), and generalization.

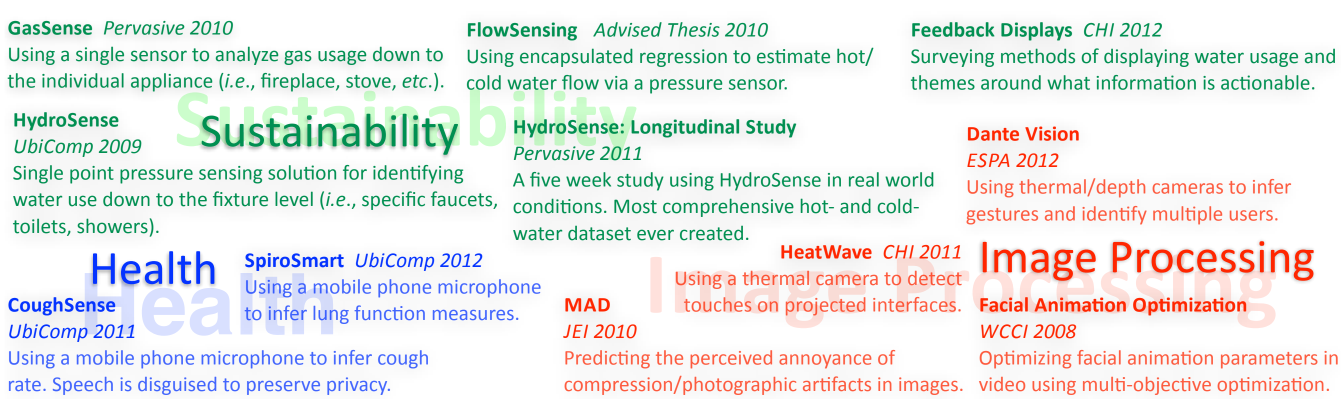

Figure 1. An overview of my previous research studies. My early work focuses on image processing and sustainable water use. My most recent work focuses on sensing markers of respiratory health using mobile phone microphones.

SIGNAL PROCESSING & MACHINE LEARNING FOR SUSTAINABILITY

The goal of signal processing and sensing in sustainability is twofold: (1) support existing initiatives by making them easier, lower cost, or more effective and (2) enable new initiatives by sensing new phenomena. I have looked at addressing these goals for electricity [9] and natural gas [4], but my dissertation focuses on water [8,12,13]. In particular, it focuses on low-cost, non-intrusive methods for sensing disaggregated water use. There is a diverse interest in this type of technology. For instance, policy makers often conduct large studies to evaluate the end-uses of residential water [2,19], influencing tax programs and plumbing regulations. Computer scientists and psychologists are investigating "eco-feedback" methods which breakdown water usage using different visualizations. This is particularly important because it helps to initiate and sustain behavior change by increasing awareness (*i.e.*, what does it cost to take a hot shower?). For electricity, some studies have shown eco-feedback can increase efficiency up to 30%, but much less is known about water. Disaggregated water studies are often resource constrained because the manpower to install and maintain the sensing is prohibitive [2].

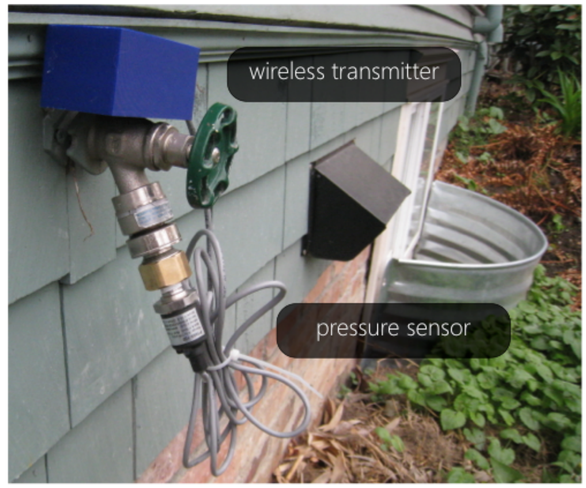

Figure 2. An example installation of HydroSense. Based on the pressure readings, the system infers the amount of water flow from every valve in the home and classifies which fixture is responsible.

My water sensing work, HydroSense, has begun to alleviate that burden. It relies on output from *a single pressure sensor* which can be installed anywhere on the plumbing system and does not require installation by a professional, as shown in Figure 2. HydroSense uses signal processing and machine learning to infer water activity from this pressure stream. This works because the plumbing system forms a closed loop pressure system, with water held at a stable pressure throughout the piping when no water is flowing. Once a fixture is activated two phenomena occur: (1) a pressure gradient in the plumbing system forms and (2) pressure waves (*i.e.*, transients) are excited. An example pressure wave is shown in Figure 3. The properties of the waves are affected by where water is activated in the system and what type of fixture is used. HydroSense uses these properties to classify the fixture that generated the pressure wave. The first part of my dissertation focuses on the *feasibility* of measuring water use from pressure waves—feature extraction and evaluating accuracy in a longitudinal field deployment. The second portion of my dissertation focuses on using semi-supervised learning to make the calibration of the system *practical*.

HydroSense: Non-intrusive Water Fixture Classification

The plumbing system in a home is a linear system or network. Like any linear network, its response to sudden changes is governed by initial conditions and the arrangement of its components (*i.e.*, pipes, in this case). When a fixture is turned on, a sudden impulsive change occurs in the system and it responds linearly—*i.e.*, the response is governed (almost) entirely by linear differential equations. Therefore, there will be a finite number of resonances that exponentially decay over time until the network reaches equilibrium. I employ a number of features to exploit this physical phenomenon including resonance tracking, exponential fitting, and frequency transforms. I also devised a form of Cepstral transformation that leverages the way bandwidths of resonances shift in the plumbing system when water is running. This transformation is more succinct and I showed it to be a more meaningful representation than the FFT [12]. The Cepstra also have the added advantage of partially separating some of the "source excitation effects" from the "filtering effects" of the plumbing system—a form of blind de-convolution. There are also a number of features I employ that are specific to the way a valve is operated. For instance single handle shower valves typically activate cold water before mixing with hot water, and flapper valves on toilets typically shut off gradually—each phenomenon can be observed in the pressure stream.

Of course, physical phenomenon is only half the story—water use is driven by human activity. There are a number of features that leverage human behavior about how water is used; features such as usage duration, time of day, flow adjustments, and so on. From a machine learning point of view, this transforms the problem from "classifying a fixture based on one transient" into "classifying a sequence of fixtures based on their entire pressure signature"—a time series classification problem. I chose to adopt and manipulate a number of machine learning methods from speech recognition, including hierarchical Markov models like Bayesian networks and

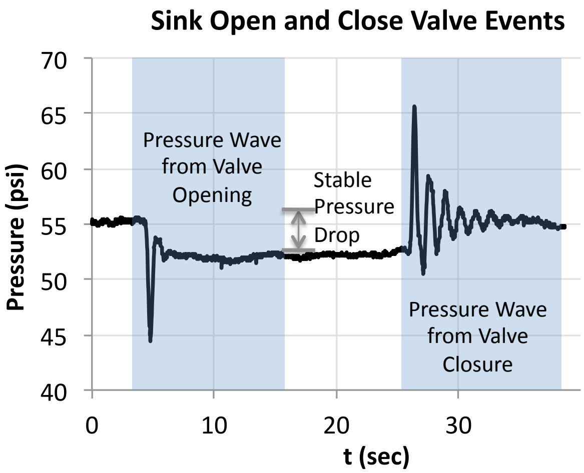

Figure 3. Example pressure waves from turning on and off a kitchen sink. I use the pressure drop to infer the flow rate of the sink and properties of the pressure wave to classify it as a "sink."

conditional random fields (CRFs). I was able to successfully manipulate the decoding algorithms of these Markov models to better incorporate novel characteristics of water use. For example, pairing the open and close transients together and creating a "grammar" for how valves can form a sequence.

To evaluate my methods, I conducted a longitudinal study of real water use over a five-week period. To collect ground truth labels of pressure data, I designed and installed an embedded network of wireless sensors in five homes—a collection of accelerometers, reed switches, hall effect sensors, and power meters that monitored every valve in the home, from bathroom sinks to clothes washers. These sensors logged when a fixture was activated and whether the water draw was hot, cold or mixed. The deployment allowed me to look at real water use in real homes with real families. **This is a common theme of my research—painstakingly crafting deployments to obtain the most realistic evaluation possible**. At the end of data collection, the dataset included almost 15,000 water transients and labels, allowing me to evaluate different feature's importance and methods of inference. To the best of my knowledge, this is the most comprehensive dataset of disaggregated hot- and cold-water use ever collected.

Using this dataset, I showed that it is possible to classify the *type of fixture* being used with 98% accuracy (*e.g.*, a toilet), the exact fixture with 95% accuracy (*e.g.*, master bathroom toilet), and whether that fixture was using hot, cold, or mixed water with 84% accuracy (*e.g.*, hot water from master bathroom sink) [13]. This evaluation showed that a single point sensing approach could be *feasible*. However, it also showed that reliable calibration of such a system required almost 1000 labeled events, a significant adoption barrier in almost all applications. With this in mind, I began researching methods of semi-supervised calibration that mitigate the need for labeled examples.

HydroSense: Unlabeled Big Data

Because of the calibration requirements, the problem fits into a larger, underexplored area of "big data"—where 10,000 homes require 10,000 separately trained models. That is, the vast majority of measureable features do not generalize across different homes and, thus, different models for each installation are required. Some knowledge about how features and classes are distributed, however, can generalize. For example, the bathroom sink is mostly used while the toilet is running and typically has a smaller pressure drop than the toilet. Using this knowledge to update parameters of a machine-learning algorithm, however, is challenging. To leverage this kind of prior knowledge, I turned to Generalized Expectation (GE) criteria. GE allows incorporation of prior knowledge through feature and label constraints. For example, if prior experience tells me that 27% of all compound events (*i.e.*, two water fixtures running simultaneously) come from the bathroom sink, then GE affords a method of updating the parameters of a classifier until this constraint is met. No labels are required—only some prior knowledge about the expectation of feature constraints, given their predicted label.

Part of my dissertation focuses on using GE with conditional random fields, extending the work of [5]. However, I also created a more flexible training methodology that was able to incorporate different training objectives. This allows the training methodology to include some labeled data for the home, some unlabeled data, and possibly generalizing data from other homes (*i.e.*, features related to water usage behavior). I therefore created a framework for optimizing *many* constraint objectives in GE. In essence, I turned the training of GE into a multi-objective optimization problem. Such problems are many times impractical to solve, so I created an update algorithm that trades off the benefits of evolutionary computation in non-convex spaces with the efficiency of gradient-based optimization for quickly finding local optima [18]. I also eliminated the need for manually setting tuning parameters.

Using this framework, I am currently investigating its utility in calibrating HydroSense. With my current timeline, I should have results in March 2013 that demonstrate the accuracy, training time, and tradeoffs between different constraints. I am also looking at using this framework outside of water disaggregation—for example to calibrate electricity disaggregation systems and on-body sensing algorithms. I also have a number of ideas for using this framework in health and wellness applications (see "research plan").

SIGNAL PROCESSING AND MACHINE LEARNING FOR HEALTH

I have found that the number and type of applications the ubiquitous computing community can address is almost without scope. I find this very comforting because the areas in which I would like to impact are similarly diverse. Aside from my dissertation, I am exploring a number of projects for out-of-clinic respiratory sensing. In particular, sensing respiratory baselines such as objective cough measures and spirometry. These topics highlight a central theme of my research—finding an accepted medical quantity and using already ubiquitous devices to sense it outside of the medical clinic. In respiratory diseases like asthma, chronic obstructive pulmonary disease (COPD), and cystic fibrosis the goal of sensing outside the clinic is typically to manage care, providing baselines and trends that reduce hospital visits and ultimately improve outcomes [21,24]. Using a mobile phone to sense these quantities reduces adoption barriers such as cost while it also increases compliance—the phone is almost always near the patient and automatic uploading to electronic medical records is trivial.

Privacy Preserving Cough Sensing

The first ubiquitous respiratory technology I created for the mobile phone was made for assessing how often a person coughs. Self-report of cough frequency and severity is notoriously unreliable, particularly in patients with chronic respiratory conditions. Therefore, objective monitoring of cough frequency has the potential for substantial clinical benefits. In particular, it allows early detection of respiratory exacerbations in patients with chronic respiratory diseases. Early intervention in exacerbations of these conditions has been shown to decrease hospitalization rates and improve long-term outcomes, including survival [23]. I worked with a team of pulmonologists to characterize the design goals of such a system, organizing the needs as follows: (1) the system should be as accurate as possible with a small number of false positives per hour, (2) it should only require a small amount of calibration, (3) it should be non-invasive, only minimally instrumenting the user, (4) it should be privacy preserving so that users are not worried about private audio or sensor information leaking to an unknown party, (5) and, finally, it should have a long battery life for 24 hour cough monitoring.

The system I created addressed all of the concerns and design requirements. I used the microphone of a mobile phone to record ambient noises and classify cough sounds using the audio spectrogram. This ensured the system was sufficiently non-invasive. I designed features that captured the key characteristics of a general cough sound without capturing the "personal" sounds of a cough that are specific to an individual. In this way, I was able to create a single model that generalized across users and did not require calibration—something that previous work had not been able to achieve. The machine learning approach used carefully constructed principal components and a random forest classifier. I evaluated the system in 19 participants wearing the cough monitor during their normal routines [16]. The system identified coughs from the audio stream with a 92% true positive rate and 0.5% false positive rate—the most accurate published results to date for an automatic system. The approach has the added benefit of being able to reconstruct the cough audio from saved features, allowing physicians to listen to the cough sound. However, because privacy of collected audio is a major concern, the features were designed such that all audio is incomprehensible except the cough sound. I verified this in a subjective experiment where users rated the fidelity of cough sounds and tried to annotate recorded speech. I was able to achieve a rating, on average, of 4 out of 5 (*i.e.,* a "good quality cough sound"), while simultaneously disguising 95% of speech.

This technology is currently being used in a study to assess cough rate of infants and young children with cystic fibrosis. It is vitally important that these patients cough often to prevent mucus pooling, which can cause infection and death if not well controlled. In the past, physicians relied on parents to report cough frequency but found this to be unreliable, especially for assessing nocturnal coughing. During the recruitment for the current study, the physicians informed me that enrollment was easier because of the privacy preserving features. They could tell parents that any conversations recorded by the monitor were unintelligible, which had been a barrier to recruitment and compliance in previous studies.

SpiroSmart: Spirometry on a Mobile Phone

The second ubiquitous respiratory technology I created for the mobile phone was made for measuring lung function via spirometry. Spirometry is the most widely employed objective measure of lung function [20] and is central to the diagnosis and management of chronic lung diseases. In a normal spirometry test, users breathe in their full lung volume and forcefully exhale through a mouthpiece until their entire lung volume is expelled. A spirometer measures the flow rate and volume of the exhaled air and reports several measurements that pulmonologists use for managing and diagnosing lung ailments. I investigated if it was possible to calculate these same measurements by having patients exhale at the screen of a mobile phone, using the mobile phone microphone to infer the flow rate and volume of air. The system allows patients to perform a spirometry test with their phone in their home without an expensive medical device (spirometers cost upwards of $2000).

Measurement of spirometry outside of the clinic allows patients and physicians to more regularly monitor for trends and detect changes in lung function. In the words of the American Thoracic Society:

> *"Many of us would love to have our patients monitor [spirometric measures such as FEV1] for monitoring a number of lung diseases. Unfortunately, the equipment required is typically too expensive and difficult to use outside the clinical setting."*

Home spirometry has the potential to result in earlier treatment of exacerbations, more rapid recovery, reduced health care costs, and improved outcomes [24]. Moreover, there are important advantages of a smartphone-based spirometry solution: the low-cost and inherent portability of a phone allows much greater uptake of home spirometry, a smartphone spirometer can have built-in coaching and feedback—mechanisms to maximize measurement acceptability—and smartphones provide the capability of easy data uploading, enabling longitudinal trending and instantaneous alerts. With these concerns in mind, I built the system, called SpiroSmart [14], where users exhale at the screen of the phone. The phone's microphone records the exhalation and sends the audio data to a server, which calculates the exhaled flow rate by estimating models of the user's vocal tract and the reverberation of sound around the user's head (see Figure 4).

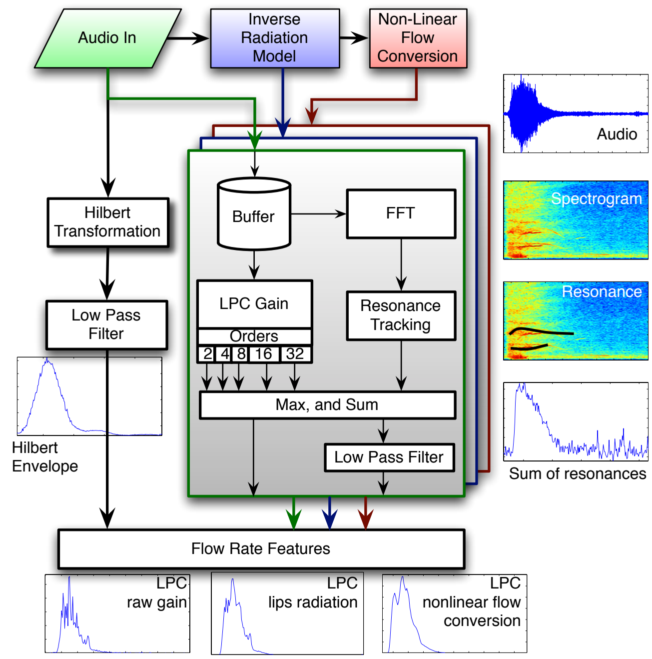

Figure 4. The feature extraction diagram for SpiroSmart. The top blocks account for the way audio radiates around the head and converts the pressure at the lips to an estimate of flow rate. The remaining processing measures different energies to reverse the effects of AC-coupling as well as leveraging effects from the vocal tract.

I extended models of the vocal tract that were developed at AT&T in the 1960's [6] that allowed me to convert the audible reverberation of the lips into a measure of the flow from the mouth. I then used machine-learning regression to calibrate the flow estimate. I employed a number of regression techniques, ultimately settling in on an ensemble of conditional random fields and random forests, and then filtering and clustering the outputs. My findings with 52 participants yielded that, SpiroSmart [14] has a median error that is comparable to the typical variation for common measures of lung function using a clinical spirometer [20]. This finding means the technology can be useful for tracking trends in these lung function measures—something especially useful for managing asthma and COPD. For diagnosis, however, we are really interested in one question: can a pulmonologist look at the results of SpiroSmart and assess a patient's lung function? To investigate this, I created a survey and recruited five pulmonologists to look at results generated from a clinical spirometer and SpiroSmart, and then enter a diagnosis. They were unaware of which results came from which device. Using these surveys, I showed that the about three-fourths of the diagnoses were the same, and the remaining quarter were highly similar, with SpiroSmart having a few false positives. The conclusion was that pulmonologists could use SpiroSmart not only to diagnose *if* a lung ailment was present, but also the degree of severity.

The American Thoracic Society[1] has called the application "complex but innovative" and that the results "appear to be fairly reliable." I am currently running additional trials of SpiroSmart in larger populations of patients with asthma, COPD, and cystic fibrosis that will be used to verify the device through the FDA. This is a two-pronged process that validates the accuracy of SpiroSmart as well as the usability of the system in home settings—two aspects the FDA has historically been concerned with for approving devices in home spirometry. Getting the trials and software to a point that is ready for FDA inspection and working with the needs of the hospitals has been a new, eye-opening experience for me. I believe this experience will be wholly beneficial as I push more of my research in mobile health towards applications requiring FDA approval.

RESEARCH PLAN

As my career progresses, I have several areas to explore in my research agenda. The first area has to do with turning my framework for semi-supervised learning into a general-purpose toolkit for non-experts in machine learning. The remaining ideas pertain to sensing health markers in high impact areas. I note also that I have a number of publications in sensing for novel interaction and image processing. I still am highly interested in these areas, but do not consider them part of my core research agenda. I would certainly be open to collaborations, or if a future students shows particular interest. My main research agenda stays with sustainability and health. Nearly all projects relating to health would fall under a number of grants from the NIH, NIA, NDA, and NCI looking at translational research. Research for health workers in the developing world may also fall under some of these grants but are more appropriate for organizations such as the Gates Foundation, WHO, or RWJF. Similar to my past research, many projects have commercial appeal and could be licensed by interested companies or startups.

Semi-Supervised Machine Learning for Non-experts

While working on my dissertation I found that the algorithms that I was creating lowered many adoption barriers for semi-supervised learning approaches. As such, I have interest in further developing these algorithms into a toolkit for non-experts to use for their own unlabeled data applications. The algorithms I created have the added benefit of not requiring tuned parameters, which helps many non-experts to try out the algorithms without knowing much about the methodology—something that can reduce access barriers significantly [22].

Opportunistic Health Markers

Many markers for health require the patient to perform a specific test (*e.g.*, spirometry, blood pressure using a cuff). Because these require active participation, low compliance and even fabricated data are common occurrences. Opportunistic sensing, on the other hand seeks to provide markers of health from everyday activities or already ubiquitous sensors. For example, I have a growing interest in using video baby-monitors to monitor breathing and pulse rate—and analyze movements, crying habits, and apnea. For adults, the phone is often a convenient bedside microphone for analyzing occurrences of apnea or even something as benign as snoring patterns. I intend to investigate ways of leveraging these technologies to sense markers of health and wellness continuously. These applications will be a mainstay of my research agenda and I look forward to working with medical professionals in these diverse areas.

Assistive Care and Elder Care

The technologies for water and gas sensing I developed during my thesis also have a number of ramifications in the field of elder care. For example, detecting the onset of dementia or potential problems of an aging parent, "*the stove has been on for 4 hours*" or "*The bathroom faucet has been running for 20 minutes*" or "*The toilet was used ten times from midnight to 6AM*." These technologies could also have great utility in creating daily activity logs (DALs) for a patient—something vitally important in assessing mobility and independence. Aside from this, I am also interested in looking at ways to address pain management and passively detect pain in the home. For instance, a thermal camera may be able to locate pain points (and severity) and evaluate blood circulation [1]. This type of sensing could have ramifications for managing many ailments—even early detection of stroke and diabetes management. I look forward to working with psychologists and computer scientists in creating such systems.

Ubiquitous Sensing and Testing for Active Disease Management

My recent research has focused on using ubiquitous sensors to carry out common tests for respiratory disease management and diagnosis. These technologies help the user to perform a medical test, like spirometry, without

---

[1] http://www.thoracic.org/education/mobile-musings/articles/november-2012.php

the need for additional sensors. But these tests are not limited to respiratory sensing. Other quantities can also be measured using already ubiquitous sensors—quantities like blood pressure for cardiac diseases, bilirubin levels for infants, and pulse-oximetry. I have a number of ideas for how to lower the cost of such tests and allow individuals to sense baselines with common equipment (*i.e.*, with nothing more than a mobile phone). I also want to pursue uses for these technologies outside of their intention. For example, using spirometry measures to increase awareness of air quality. I look forward to working with collaborators in image and video processing as well as those with experience in building medical sensing devices.

Health Sensing for the Developing World

While the kinds of technologies I have described have a number of impacts in reducing healthcare costs and expanding options for patients, they also have the potential to increase access and quality of life in the developing world. I want to continue to innovate on these technologies so that they are more appropriate for use in developing regions. For example, I am investigating how to adapt SpiroSmart so that it no longer requires an Internet connected smartphone, but instead uses a call-in-service, employing the voice channel to send audio. This makes the system agnostic to the phone type and more appropriate for use in the developing world. From an application level this seems a trivial change, but from a signal-processing viewpoint, the problem is completely changed—compression and channel characteristics become important factors for feature selection.

CONCLUSION

My research combats a common theme in sustainability and healthcare: waste. My research in sustainability encourages efficiency and awareness. In health, I focus on access and reducing the cost of care and medical testing. My work in single point sensing has created new ways of monitoring and encouraging sustainable water use. My research also makes the calibration practical, creating new semi-supervised training procedures. My work has ramifications not only for signal processing and machine learning experts, but also for water management scientists and psychologists in eco-feedback application areas. My work has been disseminated through many different cross-disciplinary venues: ICIP [17], WCCI [18], SPIE [10], UbiComp [3,8,14,16], CHI [7,11], DEV [15], and Pervasive [4,13].

As I explore areas outside my dissertation, I am making a number of cross-disciplinary collaborations with medical professionals. My work in health is enabling larger out-of-clinic studies and increasing access for individuals with chronic respiratory conditions. Such methods are critically important as medical sensing moves from a lab into our homes, where the administrator of the test is not a trained professional, but a family member or a digital coach. I am passionate about finding solutions and collaborating with psychologists, policy makers, computer scientists, and medical professionals as I continue to develop sensing and inference systems for sustainability and health.

BIBLIOGRAPHY

1. Bagavathiappan, S. and Saravanan, T. Infrared Thermal Imaging for Detection of Peripheral Vascular Disorders. *Journal of Medical Physics*, (2009).
2. Beal, C., Stewart, R., and Huang, A. *South East Queensland Residential End Use Study: Baseline Results*. 2010.
3. Campbell, T., **Larson, E.C.**, Cohn, G., Alcaide, R., and Patel, S.N. WATTR: a Method for Self-powered Wireless Sensing of Water Activity in the Home. *Proceedings of the 12th ACM International Conference on Ubiquitous Computing (UbiComp 2010)*, (2010), 169–172.
4. Cohn, G., Gupta, S., Froehlich, J., **Larson, E.C.**, and Patel, S. GasSense: Appliance-level, Single-point Sensing of Gas Activity in the Home. *Proceedings of the 8th International Conference on Pervasive Computing (Pervasive 2010)*, Springer (2010), 265–282.
5. Druck, G. *Generalized Expectation Criteria for Lightly Supervised Learning*. University of Massachusetts Amherst, 2011.
6. Flanagan, J. *Speech Analysis, Synthesis, and Perception*. Springer-Verlag, Berlin - Heidelberg - New York, 1972.
7. Froehlich, J., Findlater, L., Ostergren, M., **Larson, E.C.,** *et al*. The Design and Evaluation of Prototype Eco-Feedback Displays for Fixture-Level Water Usage Data. *Proceedings of the 2012 Annual Conference on Human Factors in Computing Systems (CHI 2012)*, (2012).

8. Froehlich, J., **Larson, E.C.**, Campbell, T., Haggerty, C., Fogarty, J., and Patel, S.N. HydroSense: Infrastructure-mediated Single-point Sensing of Whole-home Water Activity. *Proceedings of the 11th ACM International Conference on Ubiquitous Computing (UbiComp 2009)*, (2009), 235–244.

9. Froehlich, J., **Larson, E.C.**, Gupta, S., Cohn, G., Reynolds, M.S., and Patel, S.N. Disaggregated End-use Energy Sensing for the Smart Grid. *IEEE Pervasive Computing*, (2010), 28–39.

10. **Larson, E.C**. and Chandler, D.M. Most Apparent Distortion: a Dual Strategy for Full-reference Image Quality Assessment. *Proceedings of SPIE Image Quality and System Performance (SPIE 2009) 7242*, 1 (2009), 72420S.

11. **Larson, E.C.,** Cohn, G., Gupta, S., et al. HeatWave: Thermal Imaging for Surface User Interaction. *Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems (CHI 2011)*, (2011), 2565–2574.

12. **Larson, E.C.**, Froehlich, J., Campbell, T., *et al*. Disaggregated Water Sensing from a Single, Pressure-based Sensor: An Extended Analysis of HydroSense Using Staged Experiments. *Pervasive and Mobile Computing*, (2010).

13. **Larson, E.C.**, Froehlich, J., Saba, E., et al. A Longitudinal Study of Pressure Sensing to Infer Real-World Water Usage Events in the Home. *Proceedings of the 9th International Conference on Pervasive Computing (Pervasive 2011)*, (2011), 50–69.

14. **Larson, E.C.**, Goel, M., Boriello, G., Heltshe, S., Rosenfeld, M., and Patel, S.N. SpiroSmart: Using a Microphone to Measure Lung Function on a Mobile Phone. *Proceedings of the 14th ACM International Conference on Ubiquitous Computing (UbiComp 2012)*, (2012).

15. **Larson, E.C.**, Goel, M., Redfield, M., Boriello, G., and Patel, S.N. Tracking Lung Function on any Phone. *Third Annual Symposium on Computing for Development (DEV 2013) to appear*, (2013).

16. **Larson, E.C.**, Lee, T., Liu, S., Rosenfeld, M., and Patel, S.N. Accurate and Privacy Preserving Cough Sensing using a Low-Cost Microphone. *Proceedings of the 13th ACM International Conference on Ubiquitous Computing (UbiComp 2011)*, (2011).

17. **Larson, E.C.**, Vu, C., and Chandler, D. Can Visual Fixation Patterns Improve Image Fidelity Assessment? *15th IEEE International Conference on Image Processing (ICIP 2008)*, (2008), 2572–2575.

18. **Larson, E.C.** and Yen, G. Facial Feature Analysis in Dynamic Bandwidth Environments: a Genetic Approach. *Proceedings of the 2008 IEEE World Congress on Computational Intelligence (WCCI 2008)*, IEEE (2008), 2819–2826.

19. Mayer, P., DeOreo, W., Kiefer, J., Opitz, E., Dziegieliewski, B., and Nelson, J. Residential End Uses of Water. *American Water Works Association*, (1999).

20. Miller, M.R., Hankinson, J., Brusasco, V., *et al*. Standardisation of Spirometry. *The European Respiratory Journal 26*, 2 (2005).

21. Morice, A. and Fontana, G. ERS Guidelines on the Assessment of Cough. *European Respiratory Journal*, (2007).

22. Patel, K., Fogarty, J., Landay, J., and Harrison, B. Examining Difficulties Software Developers Encounter in the Adoption of Statistical Machine Learning. *Proceedings of the Association for the Advancement of Artificial Intelligence (AAAI 2008)*, (2008), 1998–2001.

23. Seemungal, T., Donaldson, G., Bhowmik, A., Jeffries, D., and Wedzicha, J. Time Course and Recovery of Exacerbations in Patients with Chronic Obstructive Pulmonary Disease. *American Journal of Respiratory and Critical Care Medicine 161*, 5 (2000).

24. Sevick, M., Trauth, J., Ling, B., *et al*. Patients with Complex Chronic Diseases: Perspectives on Supporting Self-Management. *Journal of General Internal Medicine 22*, (2007).