

**Industry Affiliates Annual Research Day
November 15, 2017**

Agenda

Meeting Sessions & Talks

Talks viewed here: <http://tinyurl.com/CSEtalks2017>

Lunchtime Keynote Address

Poster and Demo Session – **Be sure to vote for the People's Choice in EACH LAB! Vote for your favorite poster in EACH LAB THAT YOU VISIT!**

Posters in each lab may be viewed here: <http://tinyurl.com/CSEposters2017>

Building Maps

Internet Access

Meeting Materials:



Posters:



Agenda

The day opens at 10:00 a.m. with registration and coffee in the Gates Commons. The [meeting sessions](#) begin at 10:30 a.m. with a state-of-the-school overview that highlights initiatives and research started over the past year. The remainder of the day consists of research talks (in parallel tracks), lunch with a keynote talk (by robotics manipulator and new Allen School Professor [Sidd Srinivasa](#)), and wraps up with our annual open house reception and lab tours with poster presentations and demonstrations. Please join us at 7:15 p.m. for this year's prize presentation. Seattle's Madrona Venture Partners has awarded the Madrona Prize for the past ten years to students participating in Allen School's Open House whose research work shows the greatest potential for commercialization. The People's Choice Award, for which attendees vote for their favorite projects, will also be awarded.

10:00 – 10:30am	Registration and coffee/breakfast Gates Commons (691)		
10:30 – 11:10am	Welcome and Overview by Ed Lazowska and Hank Levy + various faculty on research areas Gates Commons (691)		
Session I 11:15am – 12:20pm	Robotics CSE 305	Deep Learning for Natural Language Processing CSE 403	Security I CSE 691
12:25 – 1:30pm	Lunch + Keynote Talk Atrium		
Session II 1:30 – 2:35pm	HCI: Accessibility CSE 305	Cross Cutting Systems Research CSE 403	Security II CSE 691
Session III 2:40 – 3:45pm	ML + Data Science CSE 305	HCI: Fabrication CSE 403	Ubicomp I CSE 691
Session IV 3:50 – 4:55pm	Systems CSE 305	Computational Biology CSE 403	Ubicomp II CSE 691
5:00 – 7:00pm	Open House: Reception and Poster Session/Lab Tours Atrium		
7:15 – 7:45pm	Program: Madrona Prize, People's Choice Awards Microsoft Atrium		

Session I

- **Robotics (CSE 305)**

- **11:15-11:20: Introduction and Overview, [Maya Cakmak](#)**
- **11:20-11:32: *Programming of mobile manipulator robots for non-experts*, [Justin Huang](#)**

Mobile manipulator robots could potentially automate tedious jobs and help our aging population live independently. Today, only expert roboticists have the know-how to program robots to do common tasks. Our work investigates how robots can be easily programmed by non-experts. We have developed a combination of solutions, including kinesthetic programming, drag-and-drop coding, and new concepts in perception. In our user studies and in informal workshops, we have shown that non-roboticists can get robots to do useful tasks. Finally, we are also looking at how combining visual perception with human interaction can lead to even easier modes of programming robots.
- **11:32-11:44: *Reinforcement Learning for Dynamic Robot Control*, [Kendall Lowrey](#)**

Reinforcement Learning is a technique and framework to allow systems to do complex tasks on their own. While recently making a lot of news in game playing, we demonstrate our algorithms for using RL in complex robotics control tasks. Robots in factories are designed to be precise and repeatable, but for a future with robots working in a unplanned and changing world, we need robust and dynamic control behaviors that RL can help a robot discover. In this talk, we present our work on generalizing RL techniques to many different morphologies, and demonstrate the the learned dynamic behaviors on real world manipulation tasks.
- **11:44-11:56: *Structured deep dynamics models for visuomotor control*, [Arunkumar Byravan](#)**

The ability to predict how an environment changes based on forces applied to it is fundamental for a robot to achieve specific goals. Traditionally in robotics, this problem is addressed through the use of pre-specified models or physics simulators, taking advantage of prior knowledge of the problem structure. On the other hand, learning based methods such as Predictive State Representations or more recent deep learning approaches have looked at learning these models directly from raw perceptual information in a model-free manner. In this talk, I will present some work that tries to bridge the gap between these two paradigms by proposing a specific class of deep visual dynamics models (SE3-Nets/SE3-Pose-Nets) that explicitly encode strong physical and 3D geometric priors (specifically, rigid body motion) in their structure. I will present results on applying these deep architectures for real-time visuomotor control of a Baxter robot based only on raw depth data.
- **11:56-12:08: *Dexterous and fluent robotic manipulation with and around people*, [Gilwoo Lee](#)**

Robots are extremely effective in environments like factory floors that are structured for them, and currently ineffective in environments like our homes that are structured for humans. The Personal Robotics Lab is developing the fundamental building blocks of perception, navigation, manipulation, and interaction that will enable robots to perform useful tasks in environments

structured for humans. This talk will provide an overview of some of our research achievements, and challenges.

- **12:08-12:20: *Probabilistic Semantic Mapping using Graph-Structured Sum-Product Networks*, [Kaiyu Zheng](#)**

Graph-structured data appears in wide range of domains, from social network analysis, to computer vision and robotics. In this talk, we introduce Graph-Structured Sum-Product Networks (GraphSPNs), a probabilistic approach to modeling structured data, where dependencies between latent variables are expressed in terms of arbitrary, dynamic graphs. While many approaches to structured prediction place strict constraints on the interactions between inferred variables (e.g. grids and sequences), many real-world problems can only be characterized using complex graph structures of varying size, often contaminated with noise when obtained from real data. Here, we focus on one such problem in the domain of robotics. We demonstrate how GraphSPNs can be used to bolster inference about semantic, conceptual place descriptions using noisy topological relations discovered by a robot exploring large-scale office spaces. Through experiments, we show that GraphSPNs consistently outperform the traditional approach based on undirected graphical models, successfully disambiguating information in global semantic maps built from uncertain, noisy local evidence. We further exploit the probabilistic nature of the model to infer marginal distributions over semantic descriptions of as yet unexplored places and detect spatial environment configurations that are novel and incongruent with the known evidence.

• Deep Learning for Natural Language Processing (CSE 403)

- **11:15-11:20: Introduction and Overview, [Yejin Choi](#)**
- **11:20-11:40: *Simulating Action Dynamics with Neural Process Networks*, [Antoine Bosselut](#)**

Understanding procedural language requires anticipating the causal effects of actions, even when they are not explicitly stated. In this work, we introduce Neural Process Networks to understand procedural text through (neural) simulation of action dynamics. Our model complements existing memory architectures with dynamic entity tracking by explicitly modeling actions as state transformers. The model updates the states of the entities by executing learned action operators. Empirical results demonstrate that our proposed model can reason about the unstated causal effects of actions, allowing it to provide more accurate contextual information for understanding and generating procedural text, all while offering more interpretable internal representations than existing alternatives.

- **11:40-12:00: *A Neural Framework for Generalized Topic Models*, [Dallas Card](#)**

Topic models for text corpora comprise a popular family of methods that have inspired many extensions to encode properties such as sparsity, interactions with covariates, and the gradual evolution of topics. In this work, we combine certain motivating ideas behind variations on topic models with modern techniques for variational inference to produce a flexible framework for topic modeling that allows for rapid exploration of different models. I will first discuss how our framework relates to existing models, and then demonstrate that it achieves strong performance, with the introduction of sparsity controlling the trade off between perplexity and topic coherence.

- **12:00-12:20: *Deep Learning for Broad Coverage Semantics*, [Luheng He](#)**
Semantic role labeling (SRL) systems aim to recover the predicate-argument structure of a sentence, to determine essentially “who did what to whom”, “when”, and “where”. We introduce a new deep learning model for SRL that significantly improves the state of the art, along with detailed analyses to reveal its strengths and limitations. We use a deep highway BiLSTM architecture with constrained decoding, while observing a number of recent best practices for initialization and regularization. Our 8-layer ensemble model achieves 83.2 F1 on the CoNLL 2005 test set and 83.4 F1 on CoNLL 2012, roughly a 10% relative error reduction over the previous state of the art. Extensive empirical analysis of these gains show that (1) deep models excel at recovering long-distance dependencies but can still make surprisingly obvious errors, and (2) that there is still room for syntactic parsers to improve these results. These findings suggest directions for future improvements on SRL performance.

- **Security I (CSE 691)**

- **11:15-11:20: Introduction and Overview, [Franzi Roesner](#)**
- **11:20-11:35: *Securing Augmented Reality Output*, [Kiron Lebeck](#)**
Augmented reality (AR) technologies, such as Microsoft’s HoloLens head-mounted display and AR-enabled car windshields, are rapidly emerging. AR applications provide users with immersive virtual experiences by capturing input from a user’s surroundings and overlaying virtual output on the user’s perception of the real world. These applications enable users to interact with and perceive virtual content in fundamentally new ways. However, the immersive nature of AR applications raises serious security and privacy concerns. Prior work has focused primarily on input privacy risks stemming from applications with unrestricted access to sensor data. However, the risks associated with malicious or buggy AR output remain largely unexplored. For example, an AR windshield application could intentionally or accidentally obscure oncoming vehicles or safety-critical output of other AR applications. In this work, we address the fundamental challenge of securing AR output in the face of malicious or buggy applications. We design, prototype, and evaluate Arya, an AR platform that controls application output according to policies specified in a constrained yet expressive policy framework. In doing so, we identify and overcome numerous challenges in securing AR output.
- **11:35-11:50: *Computer Security, Privacy, and DNA Sequencing. Compromising Computers with Synthetic DNA, Privacy Leaks, and More*, [Peter Ney](#)**
The rapid improvement in DNA sequencing has sparked a big data revolution in genomic sciences, which has in turn led to a proliferation of bioinformatics tools. To date, these tools have encountered little adversarial pressure. This talk discusses the robustness of such tools if (or when) adversarial attacks manifest. We demonstrate, for the first time, the synthesis of DNA which --- when sequenced and processed --- gives an attacker arbitrary remote code execution. To study the feasibility of creating and synthesizing a DNA-based exploit, we performed our attack on a modified downstream sequencing utility with a deliberately introduced vulnerability. After sequencing, we observed information leakage in our data due to sample bleeding. While this phenomena is known to the sequencing community, we provide the first discussion of how this leakage channel could be used adversarially to inject data or

reveal sensitive information. We then evaluate the general security hygiene of common DNA processing programs, and unfortunately, find concrete evidence of poor security practices used throughout the field. Informed by our experiments and results, we develop a broad framework and guidelines to safeguard security and privacy in DNA synthesis, sequencing, and processing.

- **11:50-12:05: *Privacy in Online Dating*, [Camille Cobb](#)**
Online dating services let users expand their dating pool beyond their social network and specify important characteristics of potential partners. To assess compatibility, users share personal information — e.g., identifying details or sensitive opinions about sexual preferences or worldviews — in profiles or in one-on-one communication. Thus, participating in online dating poses inherent privacy risks. How people reason about these privacy risks in modern online dating ecosystems has not been extensively studied. We present the results of a survey we designed to examine privacy-related risks, practices, and expectations of people who use or have used online dating, then delve deeper using semi-structured interviews. We additionally analyzed 400 Tinder profiles to explore how these issues manifest in practice. Our results reveal tensions between privacy and competing user values and goals, and we demonstrate how these results can inform future designs.
- **12:05-12:20: *End User Security & Privacy Concerns with Smart Homes*, [Eric Zeng](#)**
The Internet of Things is becoming increasingly widespread in home environments. Consumers are transforming their homes into smart homes, with internet-connected sensors, lights, appliances, and locks, controlled by voice or other user-defined automations. Security experts have identified concerns with IoT and smart homes, including privacy risks as well as vulnerable and unreliable devices. These concerns are supported by recent high profile attacks, such as the Mirai DDoS attacks. However, little work has studied the security and privacy concerns of end users who actually set up and interact with today's smart homes. To bridge this gap, we conduct semi-structured interviews with fifteen people living in smart homes (twelve smart home administrators and three other residents) to learn about how they use their smart homes, and to understand their security and privacy related attitudes, expectations, and actions. Among other findings, we identify gaps in threat models arising from limited technical understanding of smart homes, awareness of some security issues but limited concern, ad hoc mitigation strategies, and a mismatch between the concerns and power of the smart home administrator and other people in the home. From these and other findings, we distill recommendations for smart home technology designers and future research.

Session II

- **HCI: Accessibility (CSE 305)**

- **1:30-1:35: Introduction and Overview, [James Fogarty](#)**
- **1:35-1:50: *Interaction Proxies for Runtime Repair and Enhancement of Mobile Application*, [Xiaoyi Zhang](#)**

We introduce interaction proxies as a strategy for runtime repair and enhancement of the accessibility of mobile applications. Conceptually, interaction proxies are inserted between an application's original interface and the manifest interface that a person uses to perceive and manipulate the application. This strategy allows third-party developers and researchers to modify an interaction without an application's source code, without rooting the phone, without otherwise modifying an application, while retaining all capabilities of the system (e.g., Android's full implementation of the TalkBack screen reader). This paper introduces interaction proxies, defines a design space of interaction re-mappings, identifies necessary implementation abstractions, presents details of implementing those abstractions in Android, and demonstrates a set of Android implementations of interaction proxies from throughout our design space. We then present a set of interviews with blind and low-vision people interacting with our prototype interaction proxies, using these interviews to explore the seamlessness of interaction, the perceived usefulness and potential of interaction proxies, and visions of how such enhancements could gain broad usage. By allowing third-party developers and researchers to improve an interaction, interaction proxies offer a new approach to personalizing mobile application accessibility and a new approach to catalyzing development, deployment, and evaluation of mobile accessibility enhancements.

- **1:50-2:05: *Epidemiology as a Framework for Large-Scale Mobile Application Accessibility Assessment*, [Annie Ross](#)**

Mobile accessibility is often a property considered at the level of a single mobile application (app), but rarely on a larger scale of the entire app "ecosystem," such as all apps in an app store, their companies, developers, and user influences. We present a novel conceptual framework for the accessibility of mobile apps inspired by epidemiology. It considers apps within their ecosystems, over time, and at a population level. Under this metaphor, "inaccessibility" is a set of *diseases* that can be viewed through an epidemiological lens. Accordingly, our framework puts forth notions like *risk* and *protective factors*, *prevalence*, and *health indicators* found within a *population* of apps. This new framing offers terminology, motivation, and techniques to reframe how we approach and measure app accessibility. It establishes how app accessibility can benefit from multi-factor, longitudinal, and population-based analyses. Our epidemiology-inspired conceptual framework is the main contribution of this work, intended to provoke thought and inspire new work enhancing app accessibility at a systemic level. In a preliminary exercising of our framework, we perform an analysis of the *prevalence* of common *determinants* or accessibility barriers. We assess the *health* of a stratified sample of 100 popular Android apps using Google's Accessibility Scanner. We find that 100% of apps have at least one of nine accessibility errors and examine which errors are

most common. A preliminary analysis of the frequency of co-occurrences of multiple errors in a single app is also presented. We find 72% of apps have five or six errors, suggesting an interaction among different errors or an underlying influence.

- **2:05-2:20: *Livefonts: Animated Letterforms for Low Vision*, [Danielle Bragg](#)**
The emergence of personal computing devices offers both a challenge and opportunity for displaying text: small screens can be hard to read, but also support higher resolution. To fit content on a small screen, text must be small. This small text size can make computing devices unusable, in particular to low-vision users. Usability is also decreased for sighted users straining to read the small letters, especially without glasses at hand. We propose animated scripts called livefonts for displaying English with improved legibility for all users. Because paper does not support animation, traditional text is static. However, modern screens support animation, and livefonts capitalize on this capability. We evaluate livefont legibility through a controlled lab study with low-vision and sighted participants, and find our animated scripts to be legible across vision types at approximately half the size (area) of traditional letters. We also find it to be comparably learnable to static scripts after two thousand practice sentences.
- **2:20-2:35: *Bridging the Pedestrian Accessibility Informational Gap: User-facing Applications and Large-Scale Virtual Auditing*, [Manaswi Saha](#) and Nick Bolten**
Huge numbers of people don't have access to the information they need to get around: manual wheelchair users need to know where sidewalks are steep, cane users need to know where it's safe to cross the street, and electric wheelchair users need to know where curb ramps are. We present the AccessMap, OpenSidewalks, and Project Sidewalk projects. AccessMap is like Google Maps, but for all pedestrians, providing a web interface for automated trip planning on sidewalks in the Seattle area. The OpenSidewalks project has developed a data schema for describing pedestrian networks (at least) on par with street networks, as well as a diverse suite of software tools for acquiring this data from municipal open data as well as crowdsourcing from ground-level surveys and aerial imagery. Finally, Project Sidewalk is an online tool that allows anyone — from motivated citizens to government workers — to remotely label accessibility problems by virtually walking through city streets. Basic game design principles such as interactive onboarding, mission-based tasks, and stats dashboards are used to train, engage, and sustain users. This data can then be used to build applications for pedestrians with different mobility needs.

• Cross Cutting Systems Research (CSE 403)

- **1:30-1:35: Introduction and Overview, [Luis Ceze](#)**
- **1:35-1:47: *DNA Data Storage and Hybrid Molecular/Electronic Systems*, [Luis Ceze](#)**
DNA data storage is an attractive option for digital data storage because of its extreme density, durability and eternal relevance. This is especially attractive when contrasted with the exponential growth in world-wide digital data production. In this talk we will present our efforts in building an end-to-end system, from the computational component of encoding and decoding to the molecular biology component of random access, sequencing and fluidics automation. We will also discuss some early efforts in building a hybrid electronic/molecular computer system that has the potential to offer more than just data storage.

- **1:47-1:59: *Specializing the Planet's Computation: ASIC Clouds*, [Michael Taylor](#)**
As more and more services are built around the Cloud model, we see the emergence of planet-scale workloads (think Facebook's face recognition of uploaded pictures, or Apple's Siri voice recognition, or the IRS performing tax audits with neural nets) where datacenters are performing the same computation across many users. These scale-out workloads can easily leverage racks of ASIC servers containing arrays of chips that in turn connect arrays of replicated compute accelerators (RCAs) on an on-chip network. The large scale of these workloads creates the economical justification to pay the non-recurring engineering (NRE) costs of ASIC development and deployment. As a workload grows, the ASIC Cloud can be scaled in the datacenter by adding more ASIC servers. Our research examines ASIC Clouds in the context of four key applications that show great potential for ASIC Clouds, including YouTube-style video transcoding, Bitcoin and Litecoin mining, and Deep Learning. We developed tools that consider all aspects of ASIC Cloud design in a bottom-up way, and methodologies that reveal how the designers of these novel systems can optimize TCO in real-world ASIC Clouds. Finally, we proposed a new rule that explains when it makes sense to design and deploy an ASIC Cloud, considering NRE.
- **1:59-2:11: *Designing Systems for Push-Button Verification*, [Luke Nelson](#)**
Formal verification has proven effective in eliminating large classes of bugs in system software. Unfortunately, achieving these strong guarantees often requires significant programmer effort in the form of manual proofs, which can be orders of magnitude larger than the implementations of the systems themselves. Our research is in co-designing systems and verification techniques that can provide these strong guarantees with less programmer effort, using automated "push-button" techniques that leverage SMT solvers. So far, we have demonstrated the feasibility of this approach with a verified file system, Yggdrasil (OSDI '16), and a verified operating systems kernel, Hyperkernel (SOSP '17). Moving forward, we are interested in making these techniques more general and more deployable in real-world scenarios.
- **2:11-2:23: *LightDB: Data Management for VR/AR*, [Brandon Haynes](#)**
In this talk, we introduce the data model and architecture of LightDB, a database management system designed to efficiently manage virtual reality (VR) video content at scale. VR video differs from its two dimensional counterpart in that it is spherical, non-uniformly sampled, and only a small portion of the sphere is within a user's field of view. To address these differences, LightDB treats VR video data as a logically-continuous six-dimensional light field. LightDB supports a rich set of operations over these fields, which are automatically transformed into efficient physical execution plan at runtime. This approach allows developers to declaratively express queries over VR video data and avoids the need to manually optimize workloads. In our experiments, LightDB automatically offers up to 4x improvement over manually-tuned solutions.
- **2:23-2:35: *The UW Sandcat Project: Synthesis and Verification Across the System Stack*, [Rastislav Bodik](#)**
Despite the end of Moore's Law, we will continue building machines for groundbreaking new applications. These new computers may, however, overwhelm human programmers because new capabilities will come at the cost of increasing the complexity across the system stack. For

example, legacy data-analytics programs may need to be ported to new big-data frameworks that deliver good scalability by severely restricting the programming model. Analogously, energy-harvesting computers will limit their compute power and thus will require sophisticated tradeoffs between quality and responsiveness. The good news is that future machines will program themselves. The UW SandCat project explores such automatic code synthesis and verification across the entire system stack. This talk will overview our vision and highlight recent results.

- **Security II (CSE 691)**

- **1:30-1:35: Introduction and Overview, [Tadayoshi Kohno](#)**
- **1:35-1:50: *Robust Physical-World Attacks on Deep Learning Models*, [Ivan Evtimov](#)**

Although deep neural networks (DNNs) perform well in a variety of applications, they are vulnerable to adversarial examples resulting from small-magnitude perturbations added to the input data. Attackers can cause inputs perturbed in this way to be mislabeled into a class of their choosing. However, recent studies have demonstrated that such adversarial examples have limited effectiveness in the physical world due to changing physical conditions -- they either completely fail to cause misclassification or only work in restricted cases. We propose a general attack algorithm -- Robust Physical Perturbations (RP2)-- that takes into account the numerous physical conditions and produces robust adversarial perturbations. Using a real-world example of road sign recognition, we show that adversarial examples generated using RP2 achieve high attack success rates in the physical world under a variety of conditions, including different viewpoints. For instance, we demonstrate that stop signs can be misclassified as speed limit signs by adding black and white sticker patches and that right turn signs can be misclassified as stop signs with similar perturbations. Given the lack of standardized evaluation methodology, we also propose a two-stage attack test process that includes tests in lab conditions and in driving conditions. We demonstrate that our attacks achieve high success rates in both stages.
- **1:50-2:05: *Recognizing and Imitating Programmer Style*, [Lucy Simko](#)**

Source code attribution classifiers have recently become powerful. We consider the possibility that an adversary could craft code with the intention of causing a misclassification, i.e., creating a forgery of another author's programming style in order to hide the forger's own identity or blame the other author. We find that it is possible for a non-expert adversary to defeat such a system. In order to inform the design of adversarially resistant source code attribution classifiers, we conduct two studies with C/C++ programmers to explore the potential tactics and capabilities both of such adversaries and, conversely, of human analysts doing source code authorship attribution. Through the quantitative and qualitative analysis of these studies, we (1) evaluate a state of the art machine classifier against forgeries, (2) evaluate programmers as human analysts/forgery detectors, and (3) compile a set of modifications made to create forgeries. Based on our analyses, we then suggest features that future source code attribution systems might incorporate in order to be adversarially resistant.

- **2:05-2:20: *Moving to New Devices in the FIDO Ecosystem*, [Alex Takakuwa](#)**

The FIDO Alliance envisions a world without passwords, providing the tools to revolutionize the way users authenticate on the web. The current ecosystem provides secure standards that promise to improve online account security and simplify the experience for internet connected users. This ecosystem allows users to sign into web services through authenticators (for example, a smartphone or dedicated token) that perform user authentication using an asymmetric cryptographic signature that is resistant to phishing attacks and provides two-factor authentication. Similar to the iPhone's TouchID, users on many platforms will have devices, such as phones, that can serve as FIDO authenticators. In this talk, we show how to solve the issues that arise when a user upgrades a device serving as a FIDO authenticator. We propose the Transfer Access Protocol, which details minimal additions to the authenticator and the relying party server code that can allow users to transition seamlessly to new devices.

- **2:20-2:35: *SeaGlass: Enabling City-Wide IMSI-Catcher Detection*, [Peter Ney](#)**

Cell-site simulators, also known as IMSI-catchers and stingrays, are used around the world by governments and criminals to track and eavesdrop on cell phones. Despite extensive public debate surrounding their use, few hard facts about them are available. For example, the richest sources of information on U.S. government cell-site simulator usage are from anonymous leaks, public records requests, and court proceedings. This lack of concrete information and the difficulty of independently obtaining such information hampers the public discussion. To address this deficiency, we build, deploy, and evaluate SeaGlass, a city-wide cell-site simulator detection network. SeaGlass consists of sensors that measure and upload data on the cellular environment to find the signatures of portable cell-site simulators. SeaGlass sensors are designed to be robust, low-maintenance, and deployable in vehicles for long durations. The data they generate is used to learn a city's network properties to find anomalies consistent with cell-site simulators. We installed SeaGlass sensors into 15 ridesharing vehicles across two cities, collecting two months of data in each city. Using this data, we evaluate the system and show how SeaGlass can be used to detect signatures of portable cell-site simulators. Finally, we evaluate our signature detection methods and discuss anomalies discovered in the data.

Session III

- **ML + Data Science (CSE 305)**

- **2:40-2:45: Introduction and Overview, [Kevin Jamieson](#)**
- **2:45-2:55: *TVM: End to End IR stack for AI Frameworks*, [Tianqi Chen](#)**

Deep learning and AI has become ubiquitous and indispensable. Part of this revolution has been fueled by scalable AI systems. In this talk, I am going to talk about TVM: a unified intermediate representation (IR) stack that will close the gap between the productivity-focused deep learning frameworks, and the performance- or efficiency-oriented hardware backends. TVM is a novel framework that can: Represent and optimize the common deep learning computation workloads for CPUs, GPUs, and other specialized hardware; Automatically transform the computation graph to minimize memory utilization, optimize data layout and fuse computation patterns; Provide an end-to-end compilation from existing front-end frameworks down to bare-metal hardware.
- **2:55-3:05: *High-Precision Model-Agnostic Explanations*, [Marco Ribeiro](#)**

Understanding why ML models make their predictions is paramount in many application areas, and useful for all practitioners. In this talk, we will introduce anchors: a novel model-agnostic system that explains the behavior of complex models with high-precision rules, representing local, “sufficient” conditions for their predictions. We demonstrate the flexibility of anchors by explaining a myriad of different models for different domains and tasks, and compare anchors to alternative explanation methods in a user study.
- **3:05-3:15: *Learning brain connectivity networks from neuroimaging data*, [Rahul Nadkarni](#)**

A fundamental problem in neuroscience is inferring the functional connectivity networks between brain regions that underlie cognitive behaviors such as vision, speech, and audition. Magnetoencephalography (MEG) has become a popular neuroimaging technique for studying these networks, having enough temporal resolution that we can treat the MEG signals as time series rather than independent observations. We represent the functional connectivity network as a graphical model of time series, extending previous techniques for graphs of time series by accounting for latent signals that would otherwise lead to inferring spurious connections. We apply our model to real MEG data collected from an auditory attention task.
- **3:15-3:25: *Comparative Evaluation of Big-Data Systems on Scientific Image Analytics Workloads*, [Parmita Mehta](#)**

Scientific discoveries are increasingly driven by analyzing large volumes of image data. Many new libraries and specialized database management systems (DBMSs) have emerged to support such tasks. It is unclear how well these systems support real-world image analysis use cases, and how performant the image analytics tasks implemented on top of such systems are. In this paper, we present the first comprehensive evaluation of large-scale image analysis systems using two real-world scientific image data processing use cases. We evaluate five representative systems (SciDB, Myria, Spark, Dask, and TensorFlow) and find that each of them

has shortcomings that complicate implementation or hurt performance. Such shortcomings lead to new research opportunities in making large-scale image analysis both efficient and easy to use.

- **3:25-3:35: *Supervising Music Transcription*, [John Thickstun](#)**

Music transcription can be viewed as a multi-label classification problem. We will introduce a new large-scale dataset, MusicNet, consisting of music recordings and labels suitable to supervising transcription and other learning tasks. We will then explore neural network architectures that leverage this dataset to achieve state-of-the-art performance for music transcription.

- **3:35-3:45: *Deep Learning as a Mixed Convex-Combinatorial Optimization Problem*, [Abram Friesen](#)**

As neural networks grow deeper and wider, learning networks with hard-threshold activations is becoming increasingly important, both for network quantization, which can drastically reduce time and energy requirements, and for creating large integrated systems of deep networks, which may have non-differentiable components and must avoid vanishing and exploding gradients for effective learning. However, since gradient descent is not applicable to hard-threshold functions, it is not clear how to learn them in a principled way. We address this problem by observing that setting targets for hard-threshold hidden units in order to minimize loss is a discrete optimization problem, and can be solved as such. The discrete optimization goal is to find a set of targets such that each unit, including the output, has a linearly separable problem to solve. Given these targets, the network decomposes into individual perceptrons, which can then be learned with standard convex approaches. Based on this, we develop a recursive mini-batch algorithm for learning deep hard threshold networks that includes the popular but poorly justified straight-through estimator as a special case. Empirically, we show that our algorithm improves classification accuracy in a number of settings, including for AlexNet and ResNet-18 on ImageNet, when compared to the straight-through estimator.

• **HCI: Fabrication (CSE 403)**

- **2:40-2:45: *Introduction and Overview*, [Jennifer Mankoff](#)**

- **2:45-3:00: *Fabricating Accessibility*, [Jennifer Mankoff](#)**

With the increasing power and flexibility of technologies available to consumers, we are seeing a revolution in what is being created. This talk will highlight the challenges that end users face in leveraging them effectively to address AT issues. I will touch on issues around process, materials, design and followup. In each case, there are a host of open problems that can benefit from computational methods and better tools.

- **3:00-3:15: *Designing 3D-Printed Deformation Behaviors Using Spring-Based Structures: An Initial Investigation*, [Liang He](#)**

Recent work in 3D printing has focused on tools and techniques to design deformation behaviors using mechanical structures such as joints and metamaterials. In our work, we explore how to embed and control mechanical springs to create deformable 3D-printed objects. We propose an initial design space of 3D-printable spring-based structures to support a wide range of expressive behaviors, including *stretch and compress*, *bend*, *twist*, and all

possible combinations. The talk will provide an overview of our basic approach and design tool and a set of applications uniquely enabled by 3D-printable embedded springs.

- **3:15-3:30: PL Techniques for 3D printing, [Chandrakana Nandi](#)**
This talk is about our work on using programming language techniques for improving desktop-class manufacturing such as 3D printing. Our goal is to help make these processes more accurate, fast, reliable, and accessible to end-users. We focus on three major areas where 3D printing can benefit from programming language tools: design synthesis, verified compilation of CAD to G-code, and runtime monitoring of the printing process. We present preliminary results on synthesizing editable CAD models from difficult-to-edit surface meshes, give some insights about verifying CAD compilers using proof assistants and propose runtime monitoring techniques. We conclude by discussing additional near-future directions we intend to pursue.
- **3:30-3:45: 3D Printing Wireless Connected Objects, [Vikram Iyer](#)**
Our goal is to 3D print wireless sensors, input widgets and objects that can communicate with smartphones and other Wi-Fi devices, without the need for batteries or electronics. To this end, we present a novel toolkit for wireless connectivity that can be integrated with 3D digital models and fabricated using commodity desktop 3D printers and commercially available plastic filament materials. Specifically, we introduce the first computational designs that 1) send data to commercial RF receivers including Wi-Fi, enabling 3D printed wireless sensors and input widgets, and 2) embed data within objects using magnetic fields and decode the data using magnetometers on commodity smartphones. To demonstrate the potential of our techniques, we design the first fully 3D printed wireless sensors including a weight scale, flow sensor and anemometer that can transmit sensor data. Furthermore, we 3D print eyeglass frames, armbands as well as artistic models with embedded magnetic data. Finally, we present various 3D printed application prototypes including buttons, smart sliders and physical knobs that wirelessly control music volume and lights as well as smart bottles that can sense liquid flow and send data to nearby RF devices, without batteries or electronics.

• Ubicomp I (CSE 691)

- **2:40-2:45: Introduction and Overview, [Hanchuan Li](#)**
- **2:45-3:05: DigiTouch: Reconfigurable Thumb-to-Finger Input and Text Entry on Head-mounted Displays, [Eric Whitmire](#)**
Input is a significant problem for wearable systems, particularly for head mounted virtual and augmented reality displays. Existing input techniques either lack expressive power or may not be socially acceptable. As an alternative, thumb-to-finger touches present a promising input mechanism that is subtle yet capable of complex interactions. We present DigiTouch, a reconfigurable glove-based input device that enables thumb-to-finger touch interaction by sensing continuous touch position and pressure. Our novel sensing technique improves the reliability of continuous touch tracking and estimating pressure on resistive fabric interfaces. We demonstrate DigiTouch's utility by enabling a set of easily reachable and reconfigurable widgets such as buttons and sliders. Since DigiTouch senses continuous touch position, widget layouts can be customized according to user preferences and application needs. As an example of a real-world application of this reconfigurable input device, we examine a split-

QWERTY keyboard layout mapped to the user's fingers. We evaluate DigiTouch for text entry using a multi-session study. With our continuous sensing method, users reliably learned to type and achieved a mean typing speed of 16.0 words per minute at the end of ten 20-minute sessions, an improvement over similar wearable touch systems.

- **3:05-3:25: *BiliScreen: Smartphone-Based Jaundice Monitoring for Liver and Pancreatic Disorders*, [Alex Mariakakis](#)**

Pancreatic cancer has one of the worst survival rates amongst all forms of cancer because its symptoms manifest later into the progression of the disease. One of those symptoms is jaundice, the yellow discoloration of the skin and sclera due to the buildup of bilirubin in the blood. Jaundice is only recognizable to the naked eye in severe stages, but a ubiquitous test using computer vision and machine learning can detect milder forms of jaundice. Ubiquitous monitoring can also be used as a disease management tool, allowing a person to conveniently track their condition over time after receiving treatment. We are developing BiliScreen, a smartphone app that captures pictures of the eye and produces an estimate of a person's bilirubin level, even at levels normally undetectable by the human eye. We tested two low-cost accessories that reduce the effects of external lighting: (1) a 3D-printed box that controls the eyes' exposure to light and (2) paper glasses with colored squares for calibration. In a 70-person clinical study, we found that BiliScreen with the box achieves a Pearson correlation coefficient of 0.89 and a mean error of -0.09 ± 2.76 mg/dl in predicting a person's bilirubin level. As a screening tool, BiliScreen identifies cases of concern with a sensitivity of 89.7% and a specificity of 96.8% with the box accessory.

- **3:25-3:45: *Heterogeneous Bitwidth Binarization in Convolutional Neural Networks*, [Josh Fromm](#)**

Recent work has shown that performing inference with fast, very-low-bitwidth (e.g., 1 to 2 bits) representations of values in models can yield surprisingly accurate results. However, although 2-bit approximated networks have been shown to be quite accurate, 1 bit approximations, which are twice as fast, have restrictively low accuracy. We propose a method to train models whose weights are a mixture of bitwidths, that allows us to more finely tune the accuracy/speed trade-off. We present the "middle-out" criterion for determining the bitwidth for each value, and show how to integrate it into training models with a desired mixture of bitwidths. We evaluate several architectures and binarization techniques on the ImageNet dataset. We show that our heterogeneous bitwidth approximation achieves superlinear scaling of accuracy with bitwidth. Using an average of only 1.4 bits, we are able to outperform state-of-the-art 2-bit architectures.

Session IV

- **Systems (CSE 305)**

- **3:50-3:55: Introduction and Overview, [Xi Wang](#)**
- **3:55-4:10: *Hyperkernel: Push-Button Verification of an OS Kernel*, [Helgi Sigurbjarnarson](#)**

This talk describes an approach to designing, implementing, and formally verifying the functional correctness of an OS kernel, named Hyperkernel, with a high degree of proof automation and low proof burden. We base the design of Hyperkernel's interface on xv6, a Unix-like teaching operating system. Hyperkernel introduces three key ideas to achieve proof automation: it finitizes the kernel interface to avoid unbounded loops or recursion; it separates kernel and user address spaces to simplify reasoning about virtual memory; and it performs verification at the LLVM intermediate representation level to avoid modeling complicated C semantics. We have verified the implementation of Hyperkernel with the Z3 SMT solver, checking a total of 50 system calls and other trap handlers. Experience shows that Hyperkernel can avoid bugs similar to those found in xv6, and that the verification of Hyperkernel can be achieved with a low proof burden.
- **4:10-4:25: *Eris: Coordination-Free Consistent Transactions Using In-Network Concurrency Control*, [Ellis Michael](#)**

Distributed storage systems aim to provide strong consistency and isolation guarantees on an architecture that is partitioned across multiple shards for scalability and replicated for fault tolerance. Traditionally, achieving all of these goals has required an expensive combination of atomic commitment and replication protocols – introducing extensive coordination overhead. Our system, Eris, takes a different approach. It moves a core piece of concurrency control functionality, which we term multi-sequencing, into the datacenter network itself. This network primitive takes on the responsibility for consistently ordering transactions, and a new lightweight transaction protocol ensures atomicity. The end result is that Eris avoids both replication and transaction coordination overhead: we show that it can process a large class of distributed transactions in a single round-trip from the client to the storage system without any explicit coordination between shards or replicas in the normal case. It provides atomicity, consistency, and fault tolerance with less than 10% overhead – achieving throughput 3.6–35x higher and latency 72–80% lower than a conventional design on standard benchmarks.
- **4:25-4:40: *MultiNyx: A multi-level abstraction framework for systematic analysis of hypervisors*, [Pedro Fonseca](#)**

Modern virtualization systems rely on complex processor extensions to provide efficient virtualization. Unfortunately, the reliance on such extensions adds another dimension to the already challenging task of implementing correct virtualization systems and further complicates the automatic analysis of such systems. This work proposes MultiNyx, a system that systematically analyzes modern virtual machine monitors (VMM). MultiNyx uses symbolic execution techniques to simultaneously analyze the VMM implementation and the specification of the processor extension in a scalable manner. This analysis is achieved through a selective multi-level approach, wherein most of the VMM instructions are analyzed at a high semantic level but, in contrast, the complex processor instructions are analyzed at a low

semantic level by leveraging an executable specification. Importantly, MultiNyx is able to seamlessly transition between the different semantic levels of analysis by converting their state. Furthermore, this work proposes a methodology to break down the execution of virtual machine monitors into small units of execution that are amenable to systematic analysis. Our experiments demonstrate that MultiNyx is practical and effective at analyzing VMMs. In particular, we applied MultiNyx to KVM and generated 206,628 tests. Our results show that many of the automatically generated test cases revealed inconsistencies in the results of the KVM implementation that may have security implications. In particular, 98 test cases revealed different results across KVM configurations on Intel and 641 produced different results across different architectures. We reported some of the inconsistencies found to the KVM developers, one of which already has a patch proposal.

- **4:40-4:55: *Lightweight Data Center TCP Packet Processing*, [Antoine Kaufmann](#)**

TCP is widely used for client-server communication in modern data centers. Despite its popularity, TCP packet handling is notoriously CPU intensive, accounting for an increasing fraction of data center processing time for many applications. Known techniques such as TCP segment offload and kernel bypass are of limited benefit for the small, frequent interactions typical of most data center communication patterns. We show that TCP packet handling can be made efficient, scalable, and flexible for typical data center workloads. Further, we show these goals are compatible with secure resource isolation, a necessary requirement for modern data centers. We propose a unique refactoring of TCP functionality between the application library, the OS kernel, and reconfigurable network interface hardware based on match-action tables and a limited amount of per-flow state. Common case processing is done in hardware with software assist that runs out-of-band and less frequently. Data packets are delivered directly from application to application, while congestion control is enforced by kernel software. Using an emulation-based methodology, we show that our RMT TCP can increase per-core packet throughput by 8.2x compared to the Linux kernel TCP implementation and 3.2x compared to kernel bypass.

• Computational Biology (CSE 403)

- **3:50-3:55: Introduction and Overview, [Jacob Schreiber](#)**
- **3:55-4:15: *Deep matrix factorization for the imputation of missing biological experiments*, [Jacob Schreiber](#)**

The ENCODE project is a nation-wide consortium of researchers that seek to produce high quality, uniformly produced, data sets. One aspect of this data is running various assays to identify several epigenetic marks along the genome for over a hundred cell types. This yields a giant data "cube", where one axis is cell types, one axis is assay, and one axis spans the whole genome, and is filled in with measurements for how enriched a region of the genome is for some epigenetic mark in a specific cell type. Each assay has been run for a few cell types, and each cell type has a few assays run for it, but the vast majority of experiments have not been run. A natural desire is to use the known interactions between these epigenetic marks to use the values for experiments that have been run in order to impute values for the missing experiments. Drawing inspiration from recommendation systems, like the Netflix challenge, we adopt a deep matrix factorization approach to leverage the advances from deep learning to



learn complex interactions between cell types and assays to impute the thousands of missing experiments.

- **4:15-4:35: *Explainable machine learning predictions to help anesthesiologists prevent hypoxemia during surgery*, [Scott Lundberg](#)**

Hypoxemia causes serious patient harm, and while anesthesiologists strive to avoid hypoxemia during surgery, anesthesiologists are not reliably able to predict which patients will have intraoperative hypoxemia. Using minute by minute EMR data from fifty thousand surgeries we developed and tested a machine learning based system called Prescience that predicts real-time hypoxemia risk and presents an explanation of factors contributing to that risk during general anesthesia. Prescience improved anesthesiologists' performance when providing interpretable hypoxemia risks with contributing factors. The results suggest that if anesthesiologists currently anticipate 15% of events, then with Prescience assistance they could anticipate 30% of events or an estimated additional 2.4 million annually in the US, a large portion of which may be preventable because they are attributable to modifiable factors. The prediction explanations are broadly consistent with the literature and anesthesiologists' prior knowledge. Prescience can also improve clinical understanding of hypoxemia risk during anesthesia by providing general insights into the exact changes in risk induced by certain patient or procedure characteristics. Making predictions of complex medical machine learning models (such as Prescience) interpretable has broad applicability to other data-driven prediction tasks in medicine.

- **4:35-4:55: *Reading DNA barcodes in electric current time series data from nanopore sequencing*, [Katie Doroschak](#)**

Measuring output from a biological reaction is crucial for some computational and synthetic biology applications. For example, in DNA circuits, the output of an AND gate is typically measured using fluorescence, which is severely limited by imprecise measurements and a small encoding space with few colors. Instead, we propose using DNA barcodes and nanopore sequencing machines (tiny, card deck-sized devices that produce electric current signals as DNA passes through a nano-scale pore), combined with machine learning techniques to identify and quantify these barcodes. Each barcode produces a unique signature in the electric current time series data, and by using random forests to capture the subtle differences between them, we can greatly expand the encoding space and capture more complex events like sequential barcodes and changes over time. We are developing both the biological and computational sides of this problem in-house, and I will discuss the computational aspects in this presentation.

• Ubicomp II (CSE 691)

- **3:50-3:55: Introduction and Overview, [Eric Whitmire](#)**
- **3:55-4:15: *IDCam: Precise Item Identification for AR Enhanced Object Interactions*, [Hanchuan Li](#)**

Augmented reality (AR) promises to revolutionize the way people interact with their surroundings by seamlessly overlaying virtual information onto the physical world. To achieve this goal, AR systems need to know what objects are present in the ambient environment and

where they are located. AR systems today heavily rely on computer vision for object identification; however, state-of-the-art computer vision systems can only identify the general category of an object given sufficient training, rather than its precise identity, limiting the scope of AR applications. In this work, we propose IDCam, a system designed for precise item identification for AR object interactions. IDCam fuses computer vision and radio frequency identification (RFID) to match item identities with user interactions. To validate our system, we conducted a lab study in which 5 participants interacted with a rack of clothing simultaneously. Our result demonstrated that IDCam could identify item interactions with an accuracy of 82.0% within 2 seconds.

- **4:15-4:35: *CoughSense: Cough sound analysis for pulmonary health sensing*, [Elliot Saba](#)**

Coughing is a common symptom for many pulmonary ailments. Its ubiquity makes it difficult to use for clinical and diagnostic purposes as it is a symptom with an incredibly wide set of source ailments. We propose the design and construction of a system for cough monitoring and analysis, fitting form factors across a variety of applications from ambulatory mobile cough detectors to clinical cough sound analysis. With proper training, medical professionals are able to discern important characteristics of pulmonary health from cough sounds. This motivates the construction of machine learning models to analyze cough sounds for these same important characteristics, such as the presence of tuberculosis granulomas within the lungs of a patient. To validate our system, we collected multiple datasets of cough sounds, both in ambulatory settings as well as clinical settings, with both healthy patients and patients with specific pulmonary ailments. We present and compare novel methods of detecting cough sounds in streams of audio, as well as classification of cough sounds into various categories.

- **4:35-4:55: *Carpacio: Repurposing Capacitive Sensors to Distinguish Driver and Passenger Touches on In-Vehicle Screens*, [Edward Wang](#)**

Standard vehicle infotainment systems often include touch screens that allow the driver to control their mobile phone, navigation, audio, and vehicle configurations. For the driver's safety, these interfaces are often disabled or simplified while the car is in motion. Although this reduced functionality aids in reducing distraction for the driver, it also disrupts the usability of infotainment systems for passengers. Current infotainment systems are unaware of the seating position of their user and hence, cannot adapt. We present Carpacio, a system that takes advantage of the capacitive coupling created between the touchscreen and the electrode present in the seat when the user touches the capacitive screen. Using this capacitive coupling phenomenon, a car infotainment system can intelligently distinguish who is interacting with the screen seamlessly, and adjust its user interface accordingly. Manufacturers can easily incorporate Carpacio into vehicles since the included seat occupancy detection sensor or seat heating coils can be used as the seat electrode. We evaluated Carpacio in eight different cars and five mobile devices and found that it correctly detected over 2600 touches with an accuracy of 99.4%.

Keynote Address

Wednesday, November 15, 2017: 1:00pm
Paul G. Allen Center, Microsoft Atrium

**Siddhartha Srinivasa, Boeing Endowed Professor
Paul G. Allen School of Computer Science & Engineering**

Physics-based Manipulation with and Around People

Abstract

Robots manipulate with super-human speed and dexterity on factory floors. But yet they fail even under moderate amounts of clutter or uncertainty. However, human teleoperators perform remarkable acts of manipulation with the same hardware. My research goal is to bridge the gap between what robotic manipulators can do now and what they are capable of doing.

What human operators intuitively possess that robots lack are models of interaction between the manipulator and the world that go beyond pick-and-place. I will describe our work on nonprehensile physics-based manipulation that has produced simple but effective models, integrated with proprioception and perception, that has enabled robots to fearlessly push, pull, and slide objects, and reconfigure clutter that comes in the way of their primary task.



But human environments are also filled with humans. Collaborative manipulation is a dance, demanding the sharing of intentions, inferences, and forces between the robot and the human. I will also describe our work on the mathematics of human-robot interaction that has produced a framework for collaboration using Bayesian inference to model the human collaborator, and trajectory optimization to generate fluent collaborative plans.

Finally, I will talk about our new initiative on assistive care that focuses on marrying physics, human-robot collaboration, control theory, and rehabilitation engineering to build and deploy caregiving systems.

Speaker Bio

Siddhartha Srinivasa is the Boeing Endowed Professor at the Paul G. Allen School of Computer Science & Engineering at the University of Washington.

He works on robotic manipulation, with the goal of enabling robots to perform complex manipulation tasks under uncertainty and clutter, with and around people. To this end, he founded the Personal Robotics Lab in 2005. He has been a PI on the Quality of Life Technologies NSF ERC, DARPA ARM-S and the CMU CHIMP team on the DARPA DRC. Sidd is also passionate about building end-to-end systems (HERB, ADA, HRP3, CHIMP, Andy, among others) that integrate perception, planning, and control in the real world. Understanding the interplay between system components has helped produce state of the art algorithms for object recognition and pose estimation (MOPED), and dense 3D modeling (CHISEL, now used by Google Project Tango).

Sidd received a B.Tech in Mechanical Engineering from the Indian Institute of Technology Madras in 1999, an MS in 2001 and a PhD in 2005 from the Robotics Institute at Carnegie Mellon University.

He played badminton and tennis for IIT Madras, captained the CMU squash team, and lately runs competitively.

Poster and Demo Session

Wednesday, November 15: 5:45-7:00pm
Various locations throughout the Paul G. Allen Center
Madrona Prize and People's Choice Award presentation, 7:15-7:45 p.m.
Microsoft Atrium

(Last updated 15 November 2017)

Seattle's Madrona Venture Partners has awarded the Madrona Prize for the past twelve years to students participating in the Allen School's Open House whose research work shows the greatest potential for commercialization. The People's Choice Award, for which attendees vote for their favorite projects, will also be awarded. Please join us at 7:15 p.m. for this year's prize presentation.

Please check the following link for the most up-to-date list of posters:

<http://tinyurl.com/CSEposters2017>

Robotics + Computer Vision (2nd floor landing)

- 5 Characterizing the Design Space of Rendered Robot Faces**
Alisa Kalegina and Maya Cakmak
- 6 Learning Deep Generative Spatial Models for Mobile Robots**
Andrzej Pronobis and Rajesh P. N. Rao
- 1 Learning Stylized Character Expressions from Humans**
Deepali Aneja, Bindita Chaudhuri, Alex Colburn, Gary Faigin, Linda G Shapiro, Barbara Mones
- 7 Probabilistic Semantic Mapping using Graph-Structured Sum-Product Networks**
Kaiyu Zheng, Andrzej Pronobis and Rajesh P. N. Rao
- 8 Rapid programming of a mobile manipulation robot**
Justin Huang and Maya Cakmak
- 9 Semantic segmentation of hands for robot task demonstrations**
Justin Huang, Yu Xiang, Dieter Fox, and Maya Cakmak
- 2 YOLO9000: Better, Faster, Stronger**
Joseph Redmon and Ali Farhadi
- 3 Vega-Lite: a concise specification language for interactive visualizations**
Arvind Satyanarayan, Dominik Moritz, Kanit Wongsuphasawat, Jeffrey Heer
- 4 Voyager: Accelerating Data Exploration with Visualization Recommendations**
K Wongsuphasawat, D Moritz, A Anand, J Mackinlay, B Howe, J Heer

HCI + Accessibility + Learning (4th floor landing)

- 10 A Pilot Deployment of an Online Tool for Large-Scale Virtual Auditing of Urban Accessibility**
Manaswi Saha, Kotaro Hara, Soheil Behnezhad, Anthony Li, Michael Saugstad, Hanuma Maddali, Sage Chen, Jon E. Froehlich

- 11 A Predictive Algebra Assistant for Handwritten Mathematics**
Ben Jones, Steve Tanimoto
- 55 Augury: Predicting visual appeal of website design**
Manuel Nordhoff and Katharina Reinecke
- 12 Comparing Interaction Modalities for Productivity Applications in Virtual Reality**
Ying-Chao Tung, Stephen DiVerdi, Aaron Hertzmann, Jacob O. Wobbrock
- 57 Designing 3D-Printed Deformation Behaviors Using Spring-Based Structures: An Initial Investigation**
Liang He, Huaishu Peng, Joshua Land, Mark D. Fuge, Jon E. Froehlich
- 13 Enhancing the Accessibility of Mobile Applications**
Xiaoyi Zhang, Anne Spencer Ross, Anat Caspi, James Fogarty, Jacob O. Wobbrock
- 14 Epidemiology as a Model for Large-Scale Application Accessibility Assessment**
Anne Spencer Ross, Xiaoyi Zhang, James Fogarty, Jacob O. Wobbrock
- 15 Examining Self-Tracking by People with Migraine: Goals, Needs, and Opportunities in a Chronic Health Condition**
Jessica Schroeder, Chia-Fang Chung, Daniel A. Epstein, Ravi Karkar, Adele Parsons, Natalia Murinova, James Fogarty, Sean A. Munson
- 16 How Feedback and Gender Affect Learners' Decisions to Share Creative Computing Projects**
Emilia Gan, Benjamin Mako Hill, Sayamindu Dasgupta
- 56 LabintheWild: Large-scale online experimentation with diverse uncompensated samples**
Nigini Oliveira, Eunice Jun, and Katharina Reinecke
- 17 Rewire: Interface Design Assistance from Examples**
Amanda Swearngin, Mira Dontcheva, Wilmot Li, Morgan Dixon, Joel Brandt, Andy Ko
- 18 Studying the Effects of Learning Disabilities on Programming Education with Kokopelli's World**
Rob Thompson, Steve Tanimoto
- 19 Yarn: Designing for Telling Stories with Personal Data**
Daniel Epstein, Koko Nakajima, Mira Dontcheva, James Fogarty, Sean Munson
- 20 Women's Safety in Public Spaces: Examining the Efficacy of Panic Buttons in New Delhi**
Naveena Karusala and Neha Kumar

ML + data science (5th floor landing)

- 21 A unified approach to interpreting model predictions**
Scott Lundberg and Su-In Lee
- 22 Cosette: An Automated Prover for SQL**
Shumo Chu, Chenglong Wang, Konstantin Weitz, Daniel Li, Alvin Cheung, Dan Suciu
- 24 Delta-Based Neural Network Inference for Fast Video Analytics**
Jingjing Wang and Magdalena Balazinska
- 25 Fast Bayesian Network Structure Learning Given Domain-Derived Constraints**
Jacob Schreiber and William Noble
- 26 PerfEnforce: A Dynamic Scaling Engine for Analytics with Performance Guarantees**
Jennifer Ortiz, Brendan Lee, Magdalena Balazinska, Joseph L. Hellerstein and Johannes Gehrke
- 27 Probabilistic Database Summarization for Interactive Data Exploration**
Laurel Orr, Magdalena Balazinska, Dan Suciu
- 28 Sketching for Join Cardinality Estimation**
Walter Cai, Magdalena Balazinska, Dan Suciu

29 Supervising Music Transcription

John Thickstun, Zaid Harchaoui, Dean Foster, and Sham Kakade.

Computer Systems / Systems & Networking / Architecture (CSE591)

59 AdaRes: Adaptive Resource Management for Virtual Machines

Ignacio Cano, Karan Gupta, Arvind Krishnamurthy

30 Approximating Fair Queuing on Reconfigurable Switches

Naveen Kr. Sharma, Ming Liu, Kishore Atreya, and Arvind Krishnamurthy

31 Building a High-Performance Consistent Distributed Key-Value Store Using Programmable NICs

Ming Liu, Arvind Krishnamurthy, Karan Gupta

32 Clustering Billions of Reads for DNA Data Storage

Cyrus Rashtchian, Konstantin Makarychev, Miklos Z. Racz, Siena Dumas Ang, Djordje Jevdjic, Sergey Yekhanin, Luis Ceze, Karin Strauss

34 Eris: Coordination-Free Consistent Transactions Using In-Network Concurrency Control.

Jialin Li, Ellis Michael, and Dan R. K. Ports

35 Hyperkernel: Push-Button Verification of an OS Kernel

Luke Nelson, Helgi Sigurbjarnarson, Kaiyuan Zhang, Dylan Johnson, James Bornholt, Emina Torlak, Xi Wang

36 Lightweight Data Center TCP Packet Processing

Antoine Kaufmann, Tim Stamler, Simon Peter (UT Austin), Naveen Kr. Sharma, Thomas Anderson, Arvind Krishnamurthy

58 Molecular Processing-in-Memory

Kendall Stewart, Max Willsey, Yuan-Jyue Chen, Karin Strauss, Luis Ceze

38 Profiling a GPU Database Implementation

Emily Furst, Mark Oskin, Bill Howe

39 PHub: High performance parameter server for efficient distributed DNN training

Liang Luo, Jacob Nelson(MSR), Amar Phanishayee(MSR), Luis Ceze, Arvind Krishnamurthy

40 Puddle: High-Level Microfluidic Programming

Max Willsey, Sharon Newman, Pranav Vaid, Sarang Joshi, Chris Takahashi, Luis Ceze, Karin Strauss

41 Real-Time VR Video Processing with the Hardware-Friendly Bilateral Solver

Amrita Mazumdar, Armin Alaghi, Luis Ceze, Mark Oskin

33 TVM: End-to-End IR Stack for Deep Learning Systems

Tianqi Chen, Thierry Moreau, Haichen Shen, Ziheng Jiang, Luis Ceze, Carlos Guestrin, Arvind Krishnamurthy

UbiComp (Gates Commons)

42 A Smartphone-Based System for Assessing Intraocular Pressure + Non-invasive Approach

Alex Mariakais, Edward Wang, Shwetal Patel, Joanne Wen + Chunjong Park

43 BiliScreen: Smartphone-Based Scleral Jaundice Monitoring for Liver and Pancreatic Disorders

Alex Mariakakis, Megan A. Banks, Lauren Phillipi, Lei Yu, James Taylor, Shwetak Patel

44 Carpacio: Driver Detection

Edward Jay Wang¹, Jake Garrison¹, Eric Whitmire, Mayank Goel, Shwetak Patel

45 DigiTouch: Reconfigurable Thumb-to-Finger Input and Text Entry on Head-mounted Displays

Eric Whitmire, Mohit Jain, Divye Jain, Greg Nelson, Ravi Karkar, Shwetak Patel, Mayank Goel

46 Heterogenous Bitwidth Binary Networks

Josh Fromm, Matthai Phillipose, Shwetak Patel

47 IDCam: Precise Item Identification for AR-Enhanced Object Interactions

Hanchuan Li, Eric Whitmire, Alex Mariakakis, Victor Chan, Alanson Sample, Shwetak Patel

48 RoyalFlush: Non-invasive toilet water overflow detector

Farshid Salemi Parizi, Josh Fromm, Shwetak Patel

50 Techniques for Cough Sound Analysis

Elliot Saba, Clara Lu, Shwetak Patel

51 The UbiComp Lab Group Poster

UbiComp Lab

52 The UbiComp Lab Health Poster

UbiComp Lab

53 Using CM EMI signals for appliance disaggregation

Manoj Gulati, Shobha Sundar Ram, Amarjeet Singh, Shwetak Patel

54 VR feedback

Morelle Arian, Farshid Salemi Parizi, Shwetak Patel

Building Maps

2nd floor



Building Maps

4th floor



Building Maps

5th floor



Building Maps

6th floor



Network Access for Allen Center Visitors

How Do I Connect to the UW Wireless Network?

The UW's network is named "University of Washington" and is the one you should use when on a UW campus. You may see other networks listed, but they are not supported by UW Information Technology and may be unsafe to use.

Logging In

No login is required to use UW's wireless service unless you want to connect to services or sites outside the UW.

- If you use your browser to connect to a non-UW Web site, you automatically will be prompted to enter your UW NetID and password.
- If you are using a non-Web service (e.g., a secure terminal or file transfer tool) to communicate outside the UW, you must FIRST run a browser and view a Web page outside the UW to bring up the authentication page.

UW NetID: event0522

Password: 5C7x-5C8p-6C5p

Once you have successfully logged in (authenticated) you will be able to use services outside the UW for up to 12 hours without having to re-authenticate.
