Computational Interaction with Bayesian Methods

Per Ola Kristensson
University of Cambridge
Cambridge, UK
pok21@cam.ac.uk

Nikola Banovic
University of Michigan
Ann Arbor, Michigan, USA
nbanovic@umich.edu

Antti Oulasvirta
Aalto University
Helsinki, Finland
antti.oulasvirta@aalto.fi

John Williamson
University of Glasgow
Glasgow, UK
JohnH.Williamson@glasgow.ac.uk

ABSTRACT

This course introduces computational methods in human–computer interaction. Computational interaction methods use computational thinking—abstraction, automation, and analysis—to explain and enhance interaction. This course introduces the theory of practice of computational interaction by teaching Bayesian methods for interaction across four wide areas of interest when designing computationally-driven user interfaces: decoding, adaptation, learning and optimization. The lectures center on hands-on Python programming interleaved with theory and practical examples grounded in problems of wide interest in human-computer interaction.

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI)**.

CHI'19 Extended Abstracts, May 4–9, 2019, Glasgow, Scotland UK

ACM ISBN 978-1-4503-5971-9/19/05.
https://doi.org/10.1145/3290607.3298820

KEYWORDS

computational interaction; optimization; inference; machine learning

ACM Reference Format:

Per Ola Kristensson, Nikola Banovic, Antti Oulasvirta, and John Williamson. 2019. Computational Interaction with Bayesian Methods. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI'19 Extended Abstracts), May 4–9, 2019, Glasgow, Scotland UK*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3290607.3298820

OVERVIEW

The course will cover:

Decoding: Principled and robust formulations for probabilistically decoding noisy sensor observations into user's intended actions.

Adaptation: Probabilistic adaptation demonstrated by sequential Monto Carlo approaches for modeling HCI tasks.

Learning: Using inverse reinforcement learning for learning computational models of human behavior.

Optimization: Using Bayesian optimization to perform human-in-the-loop optimization of a user interface problem.

THIS COURSE WILL:

- demonstrate how computational approaches can drive adaptation and design of interactive systems;
- extend researchers' capabilities to build robust interactions across a wide range of contexts and devices;
- show how computational approaches can focus interaction design on the interesting work of specifying the questions, and letting computational methods resolve the answers.

STRUCTURE

The course surveys computational methods for four wide areas of interest: decoding, adaptation, learning and optimization. Each area is covered in an 80-minute session and each area is contextualized in concrete HCI problems via practical Python examples. The course will be organized as interactive lecture sessions with interleaved and very concentrated practical work. The material will be provided as interactive Jupyter (IPython) Notebooks, to allow participants to live execute examples during the lectures and immediately extend and apply the material present to practical problems. These notebooks will provide enough of a skeleton to allow participants to make rapid progress in applying the ideas presented. Lectures are in-part based on chapters in the book *Computational Interaction* (Oxford University Press, 2018).

INTENDED AUDIENCE

We expect this to course to be primarily targeted at researchers and PhD students in HCI. Industrial practitioners, interested for example in data-driven design or new interface technologies, may also find the course valuable.

CONTENT

The course will focus on optimization and inference and on applying these techniques to *concrete* HCI problems. The course will specifically look at Bayesian methods for solving decoding, adaptation, learning and optimization problems in HCI.

1. DECODING

This section will cover the concept of decoding noisy observations from a variety of sensors, such as touchscreens and depth sensors, into users' intended commands or text. We will derive an initial formulation of the decoding problem from Bayes' rule and thereafter explain how to model data probabilistically and perform inference. The primary example will be text entry and attendees will be shown how to design a probabilistic auto-correct algorithm from scratch with iterative refinements. As part of this, attendees will be introduced to language modeling and provided with concrete examples of how probabilistic decoding can be used to solve a variety of HCI problems. **After completing this section**, attendees will be able to: understand how to design a probabilistic decoder; be aware of key design decisions that affect decoding performance.

Approximate timings. Introduction to probabilistic decoding (10); Key concepts in statistical language processing (10); Designing a probabilistic text decoder (25); Design examples and implications for HCI (10); Interactive exercise (20); Wrap-up (5).

2. ADAPTATION

We will discuss how probabilistic Bayesian models can be used to infer user intention robustly, framing interaction as a problem in inverse probability. This allows engineering user interfaces that dynamically adapt to user characteristics and input device variations, and cope with multiple competing hypotheses about intention. This is a rigorous and widely applicable approach which can be applied to interaction tasks from low-level motion estimation to high-level task models. We will specifically discuss the use of probabilistic filtering to update beliefs in real-time at a hierarchy of time scales, providing for adaptive interfaces. We will discuss and implement sequential Monte Carlo approaches as a concrete algorithm to implement these models for human-computer interaction tasks. **After completing this section**, attendees will be able to: represent interaction problems as inference problems; apply probabilistic methods to robustly estimate the evolution of user intentions over time, even with very noisy input devices.

Approximate timings. Introduction to Bayesian modeling for interaction problems (20); Probabilistic filtering (25); Simultaneous inference across multiple sensors. (10); Interactive exercise (20); Wrap-up (5).

3. LEARNING

This section will review current computational approaches to describe, simulate, and predict human behavior from empirical behavior traces data stored in large behavior logs. It will contrast and compare computational modeling with existing behavior modeling in HCI. Attendees will learn how train computational models of human behavior and leverage them to create user interfaces that can automatically reason about and act in response to people's behaviors to help increase their productivity, safety, and health. The section will focus on high-level behaviors, such as routines and habits, represented by sequences of situations people find themselves in and actions they perform in those situations. Models will be trained using Inverse Reinforcement Learning. **After completing this section**, attendees will be able to: formulate behavior modeling HCI problems as computational modeling problems; learn computational models of human behavior.

Approximate timings. Review of existing methods for modeling human behavior and contrast computational modeling approaches with algorithm-driven, black-box Machine Learning models (20); Computational modeling of human behavior using Inverse Reinforcement Learning (35); Interactive exercise (20); Wrap-up (5).

4. OPTIMIZATION

This section will review Bayesian optimization, which is a principled approach to optimization of design and model parameters under noise. In HCI, it can be used to fit improve design parameters with humans in the loop and to fit cognitive models and simulator models to empirical data. The benefit of Bayesian optimization is that it is a sample-efficient algorithm for global optimization of stochastic functions. Relatively few iterations are needed with the model to obtain good parameter estimates. The principles are demonstrated in human-in-the-loop optimization of a design. **After completing this section**, attendees will be able to: formulate design problems such that an algorithm can help in exploring it; understands the requirements of this approach; understands how it can be tied to user-centered design and data-driven design.

Approximate timings. Overview of optimization approaches in HCI (10); Introduction to Bayesian optimization (20); Human-in-the-loop optimization of a user interface design problem (25); Interactive exercise (20); Wrap-up (5).

IMPROVEMENTS SINCE CHI 2017 AND CHI 2018

This course builds on and revises two successful previous CHI Courses "Computational Interaction Methods for HCI" at CHI 2017 and "Computational Interaction: Theory and Practice" at CHI 2018. Based on feedback from attendants at CHI 2017 we made following improvements: (1) the density of the material was reduced, such that more time was left for students to solve problems; (2) more preparatory materials were made available in advance; (3) we had more applications and less theory and each element was clearly tied to concrete interaction problems; and (4) we provided more diverse views from subdisciplines and practitioners to highlight application values. Feedback from CHI 2018 was overwhelmingly positive. However, one request was additional time for practical coding exercises. To address this each section will contain a 20-minute interactive exercise.

INSTRUCTOR BIOGRAPHIES

Per Ola Kristensson is a University Reader in the Department of Engineering at the University of Cambridge. His PhD thesis (at Linköping University/IBM Almaden) was on gesture keyboard technology for touchscreens and in 2007 he co-founded ShapeWriter, Inc. to commercialize this technology. In 2013 he was recognized as an Innovator Under 35 (TR35) by MIT Technology Review and appointed a Member of the Royal Society of Edinburgh Young Academy of Scotland. He has been awarded the ACM User Interface Software and Technology (UIST) Lasting Impact Award and the Royal Society of Edinburgh Early Career Prize in Physical Sciences—the Sir Thomas Makdougall Brisbane Medal.

Nikola Banovic is an Assistant Professor of Electrical Engineering and Computer Science at the University of Michigan. His research focuses on creating computational models of human behavior to study, describe, and understand complex human behaviors and enable technology that automatically reasons about and acts in response to people's behavior to help them be productive, healthy, and safe. He has been named NSERC Post-graduate Fellow and Yahoo! Fellow and published award-winning research on methods to study and model human behavior in premier HCI conferences.

Antti Oulasvirta is an Associate Professor at Aalto University where he leads the User Interfaces research group. He was previously a Senior Researcher at the Max Planck Institute for Informatics. Antti received his doctorate in Cognitive Science from the University of Helsinki in 2006, after which he was a Fulbright Scholar at the School of Information in University of California-Berkeley in and a Senior Researcher at Helsinki Institute for Information Technology HIIT. He was awarded an ERC Starting Grant (2015-2020) for research on computational design of user interfaces.

John Williamson: is a Lecturer at the University of Glasgow. He received his PhD from the University of Glasgow in 2006, after which he was awarded a SICSA Research Fellowship and a Lord Kelvin Adam Smith Fellowship in Sensor Systems. His research is focused on continuous interaction systems and computational approaches to HCI, including gesture recognition and motion sensor based interfaces, brain-computer interaction, and new feedback mechanisms for such systems. He has received multiple Honorable Mention awards at CHI. He co-founded the international summer school series on Computational Interaction.