# Crowdsourcing and Social Computing

CS59000, Fall 2018, Monday & Wednesday 4:30-5:45pm, LWSN B134

## Instructor

Ming Yin

Office: TBD
Email: mingyin@purdue.edu
Office Hours: By Appointment

## Course Description

Crowdsourcing and social computing sits at the intersection of computer science, economics and other social sciences. It concerns developing empirical understandings and designing computational systems and techniques to enable effective interactions between people and machines for solving complex problems. This course surveys the state of the art in this area. Topics of interests include incentive design, workflow design, quality control and intelligent management in crowdsourcing, computer supported collaborative work, and applications in various domains like artificial intelligence and citizen science.

## Course Schedule

| Date | Topic | Reading |
|---|---|---|
| Aug 20 | Introduction & Course overview | |
| Aug 22 | Crowdsourcing: Background and applications | Required<br>• Howe. The Rise of Crowdsourcing. WIRED. June 2006.<br>Optional<br>• von Ahn et al. reCaptcha: Human-based Character Recognition via Web Security Measures. Science, September 2008<br>• von Ahn and Dabbish. Labeling Images with a Computer Game. CHI'04<br>• Cooper et al. Predicting Protein Structures with a Multiplayer Online Game. Nature, August 2010. |
| Aug 27 | Crowd workers | Required<br>• Difallah et al. Demographics and Dynamics of Mechanical Turk Workers. WSDM'18 |

| | | |
|---|---|---|
| | | • Yin et al. The Communication Network Within the Crowd. WWW'16<br>Optional<br>• Martin et al. Being a Turker. CSCW'14<br>• Gray et al. The Crowd is a Collaborative Network. CSCW'16 |
| Aug 29 | Crowdsourcing platform: Amazon Mechanical Turk | Creating a HIT on MTurk<br>No required readings. Bring laptops to class. |
| Sep 3 | No class (Labor Day) | |
| Sep 5 | Crowdsourcing platform: Tasks and dynamics | Required<br>• Difallah et al. The Dynamics of Micro-Task Crowdsourcing: The Case of Amazon MTurk. WWW'15<br>Optional<br>• Gadiraju et al. A Taxonomy of Microtasks on the Web. HT'14<br>• Vakharia and Lease. Beyond Mechanical Turk: An Analysis of Paid Crowd Work Platforms. iConference'15<br>• Jain et al. Understanding Workers, Developing Effective Tasks, and Enhancing Marketplace Dynamics: A Study of a Large Crowdsourcing Marketplace. VLDB Endowment, March 2017 |
| Sep 10 | Crowdsourcing: Opportunities and challenges | Required<br>• Kittur et al. The Future of Crowd Work. CSCW'13<br>Optional<br>• Quinn and Bederson. Human Computation: A Survey and Taxonomy of a Growing Field. CHI'11<br>• Gadiraju et al. Human Beyond the Machine: Challenges and Opportunities of Microtask Crowdsourcing. IEEE Intelligent Systems, July 2015 |
| Sep 12 | Incentive design: Financial incentives | Required<br>• Ho et al. Incentivizing High Quality Crowdwork. WWW'15<br>Optional<br>• Mason and Watts. Financial Incentives and the “Performance of Crowds”. HCOMP'09<br>• Yin et al. The Effects of Performance-Contingent Financial Incentives in Online Labor Markets. AAAI'13 |

| | | |
|---|---|---|
| | | • Harris. The Effects of Pay-to-Quit Incentives on Crowdworker Task Quality. CSCW'15 |
| Sep 17 | Incentive design: Other incentives | Required<br>• Rogstadius et al. An Assessment of Intrinsic and Extrinsic Motivation on Task Performance in Crowdsourcing Markets. ICWSM'11<br>• Shaw et al. Designing Incentives for Inexpert Human Raters. CSCW'11<br>Optional<br>• Nov et al. Scientists@Home: What Drives the Quantity and Quality of Online Citizen Science Participation? PLOS ONE, April 2014<br>• Feyisetan and Simperl. Social Incentives in Paid Collaborative Crowdsourcing. TIST, September 2017 |
| Sep 19 | Incentive design: Intelligent management | Required<br>• Yin and Chen. Bonus or Not? Learn to Reward in Crowdsourcing. IJCAI'15<br>Optional<br>• Gao and Parameswaran. Finish Them!: Pricing Algorithms for Human Computation. VLDB Endowment, October 2014<br>• Feyisetan et al. Improving Paid Microtasks through Gamification and Adaptive Furtherance Incentives. WWW'15 |
| Sep 24 | Task assignment and recommendation | Required<br>• Ho and Vaughan. Online Task Assignment in Crowdsourcing Markets. AAAI'12<br>• Difallah et al. Pick-a-Crowd: Tell Me What You Like, and I'll Tell You What To Do. WWW'13<br>Optional<br>• Lin et al. Signals in the Silence: Models of Implicit Feedback in a Recommendation System for Crowdsourcing. AAAI'14<br>• Mavridis et al. Using Hierarchical Skills for Optimized Task Assignment in Knowledge-Intensive Crowdsourcing. WWW'16 |
| Sep 26 | Quality assurance: Empirical methods | Required<br>• Dow et al. Shepherding the Crowd Yields Better Work. CSCW'12<br>Optional |

| | | |
|---|---|---|
| | | • Huang and Fu. Enhancing Reliability Using Peer Consistency Evaluation in Human Computation. CSCW'13<br>• Sampath et al. Cognitively Inspired Task Design to Improve User Performance on Crowdsourcing Platforms. CHI'14<br>• Doroudi et al. Toward a Learning Science for Complex Crowdsourcing Tasks. CHI'16 |
| Oct 1 | Quality assurance: Algorithmic approaches | Required<br>• Whitehill et al. Whose Vote Should Count More: Optimal Integration of Labels from Labelers of Unknown Expertise. NIPS'09<br>Optional<br>• Oyama et al. Accurate Integration of Crowdsourced Labels Using Workers' Self-Reported Confidence Scores. IJCAI'13<br>• Sunahase et al. Pairwise HITS: Quality Estimation from Pairwise Comparisons in Creator-Evaluator Crowdsourcing Process. AAAI'17<br>• Wang et al. Obtaining High-Quality Label by Distinguishing between Easy and Hard Items in Crowdsourcing. IJCAI'17 |
| Oct 3 | Quality assurance: intelligent management | Required<br>• Kamar et al. Combining Human and Machine Intelligence in Large-Scale Crowdsourcing. AAMAS'12<br>Optional<br>• Bragg et al. Optimal Testing for Crowd Workers. AAMAS'16<br>• Gurari and Grauman. CrowdVerge: Predicting If People Will Agree on the Answer to a Visual Question. CHI'17 |
| Oct 8 | No class (October Break) | |
| Oct 10 | Final project: Pitch | |
| Oct 15 | Engagement control: Empirical methods | Required<br>• Yu et al. A Comparison of Social, Learning, and Financial Strategies on Crowd Engagement and Output Quality. CSCW'14<br>• Law et al. Curiosity Killed the Cat, but Makes Crowdwork Better. CHI'16<br>Optional |

| | | |
|---|---|---|
| | | • Preist et al. Competing or Aiming to be Average?: Normification as a Means of Engaging Digital Volunteers. CSCW'14<br>• Dai et al. And Now for Something Completely Different: Improving Crowdsourcing Workflows with Micro-Diversions. CSCW'15 |
| Oct 17 | Engagement control: Intelligent management | Required<br>• Segal et al. Intervention Strategies for Increasing Engagement in Crowdsourcing: Platform, Predictions, and Experiments. IJCAI'16<br>Optional<br>• Segal et al. Optimizing Interventions via Offline Policy Evaluation: Studies in Citizen Science. AAAI'18 |
| Oct 22 | Workflow design: Use-case specific workflows | Required<br>• Bernstein et al. Soylent: A Word Processor with a Crowd Inside. UIST'10<br>• Chilton et al. Cascade: Crowdsourcing Taxonomy Creation. CHI'13<br>Optional<br>• Noronha et al. Platemate: Crowdsourcing Nutritional Analysis from Food Photographs. UIST'11<br>• Kim et al. Crowdsourcing Step-by-Step Information Extraction to Enhance Existing How-to Videos. CHI'14 |
| Oct 24 | Workflow design: General workflows | Required<br>• Little et al. TurKit: Human Computation Algorithms on Mechanical Turk. UIST'10<br>Optional<br>• Kittur et al. CrowdForge: Crowdsourcing Complex Work. UIST'11<br>• Kulkarni et al. Collaboratively Crowdsourcing Workflows with Turkomatic. CSCW'12 |
| Oct 29 | Workflow design: Intelligent management | Required<br>• Dai et al. Decision-theoretic Control of Crowd-sourced Workflows. AAAI'10<br>Optional<br>• Lin et al. Dynamically Switching between Synergistic Workflows for Crowdsourcing. AAAI'12 |

| | | |
|---|---|---|
| | | • Bragg et al. Crowdsourcing Multi-Label Classification for Taxonomy Creation. HCOMP'13<br>• Tran-Thanh et al. Crowdsourcing Complex Workflows under Budget Constraints. AAAI'15 |
| Oct 31 | Beyond independent: Cooperative work examples | Required<br>• Drapeau et al. MicroTalk: Using Argumentation to Improve Crowdsourcing Accuracy. HCOMP'16<br>Optinoal<br>• Suzuki et al. Atelier: Repurposing Expert Crowdsourcing Tasks as Micro-Internships. CHI'16<br>• Chang et al. Revolt: Collaborative Crowdsourcing for Labeling Machine Learning Datasets. CHI'17 |
| Nov 5 | Final project: Peer feedback (Session 1) | |
| Nov 7 | Final project: Peer feedback (Session 2) | |
| Nov 12 | Cooperative work: Complex task | Required<br>• Retelny et al. Expert Crowdsourcing with Flash Teams. UIST'14<br>Optional<br>• Hahn et al. The Knowledge Accelerator: Big Picture Thinking in Small Pieces. CHI'16<br>• Valentine et al. Flash Organizations: Crowdsourcing Complex Work by Structuring Crowds As Organizations. CHI'17 |
| Nov 14 | Cooperative work: Intelligent management | Required<br>• Salehi et al. Huddler: Convening Stable and Familiar Crowd Teams Despite Unpredictable Availability. CSCW'17<br>Optional<br>• Singla et al. Learning to Hire Teams. HCOMP'15<br>• Zhou et al. In Search of the Dream Team: Temporally Constrained Multi-Armed Bandits for Identifying Effective Team Structures. CHI'18 |
| Nov 19 | Crowd-Powered Systems | Required<br>• Lasecki et al. Real-Time Captioning by Groups of Non-Experts. UIST'12 |

| | | |
|---|---|---|
| | | • Kokkalis et al. MyriadHub: Efficiently Scaling Personalized Email Conversations with Valet Crowdsourcing. CHI'17<br>Optional<br>• Vashistha et al. Respeak: A Voice-based, Crowd-powered Speech Transcription System. CHI'17<br>• Nguyen et al. An Interpretable Joint Graphical Model for Fact-Checking from Crowds. AAAI'18 |
| Nov 21 | No class (Thanksgiving) | |
| Nov 26 | Crowdsourcing: Future Ideas | Required<br>• Whiting et al. Crowd Guilds: Worker-led Reputation and Feedback on Crowdsourcing Platforms. CSCW'17<br>Optional<br>• Morris et al. Subcontracting Microwork. CHI'17<br>• Vaish et al. Crowd Research: Open and Scalable University Laboratories. UIST'17 |
| Nov 28 | No class (Project day) | |
| Dec 3 | Final project presentation 1 | |
| Dec 5 | Final project presentation 2 | |

## Grading

- Assignment: 10%
- Reading responses: 15%
- Class discussion: 10%
- Paper presentation: 20%
- Final project: 45% (proposal + pitch: 10%, milestone presentation: 10%, final presentation: 10%, final report: 15%)

## Paper Reading, Presentation, and Discussion

Most classes in this course consist of paper reading, presentation and discussion. Specifically, in a typical class, we will cover 1-2 papers on one topic, and 1-2 students will be assigned to give a presentation on this topic. Responsibility of presenters of one class include:

- Read all required paper(s), and at least one optional paper, for that class.
- After discussing with the instructor (one week before the class), post 2-3 conversation-provoking questions related to the required paper(s) of that class.

- Give a presentation in class, which should review all required paper(s) of that class, and also briefly introduce the optional paper(s) that they have read.
- Lead the discussion in class.

Responsibility of non-presenters of one class include:
- Read all required paper(s) for that class.
- Before class, provide reading responses to all questions that the presenters of that class post.
- Participate in the discussion in class.

## Final Project

Final project serves as an opportunity for students to get hands-on experience in crowdsourcing and social computing research. Projects are open-ended; sample projects include:
- Design and develop crowdsourcing workflows, platforms or systems for innovative use cases.
- Design and conduct online experiments to understand the behavior of crowd workers.
- Propose new methods to improve the efficiency of crowdsourcing processes.
- Understand dynamics in current crowdsourcing or social computing systems through theoretical or empirical data analysis.

Students are also encouraged to connect the final project with their own research.

Students can complete the project either individually or in a group of two. Tasks related to the final project include:
- Submit a project proposal which identifies the problem that the project aims to solve.
- Give a pitch presentation on the project proposal in class.
- Give a milestone presentation on the project progress and get peer feedback.
- Give a final presentation on the project in class, reporting the results of the project.
- Submit a final project report summarizing the project.

More detailed instruction on the final project will be provided through project guidelines.

## Prerequisites

Basic programming skills required. Students should be comfortable with at least one programming language (e.g., C, Java, Python, etc.). Knowledge with artificial intelligence and machine learning is welcome.

## Required Texts

No textbook is required for this course.