

CS592 Human-AI Interaction

Instructor: Tianyi Zhang, Assistant Professor of Computer Science

Email: tianyi@purdue.edu

Lecture: 10:30-11:20 MWF @ Lawson B134

Office Hours: Tue and Thurs 3pm-4pm

Office: Lawson 3154H

Instructional Modality: Face-to-Face

Course Credits: 3.0

Prerequisites: Basic programming skills are required.

Course Description

Have you ever wondered about these:

- What is the role of humans in the future of AI?
- Will programming jobs no longer exist because of large language models like GPT-3?¹
- How far are we from the “black art” of natural language programming as Dijkstra called it 40 years ago?²
- Why does IBM suddenly seek to sell Watson Health, their AI for healthcare division, after 10 years of huge investment?³
- Self-driving cars are coming, but are we ready?
- How can humans efficiently give feedback to AI and correctify its mistakes?
- How will humans and AI evolve together in the next decade?

This course will help you answer those questions. You will read and discuss research papers in human-AI interaction, including but not limited to research topics about (1) AI-based systems and applications working with---or clashing against---the strengths and weaknesses of human cognition, (2) how to design interactive, human-in-the-loop approaches that achieve human-AI symbiosis, and (3) how to support interpretability, transparency, trust, and fairness in AI-based systems. Specifically, we will look into recent research advances in several trending domains such as “AI for code”, healthcare, and autonomous driving.

¹ [OpenAI's GPT-3 Can Now Generate The Code For You](#). *Analytics India Magazine*, July 20, 2020.

² E. W. Dijkstra. [On the foolishness of “natural language programming”](#). In *Program Construction*, pages 51–53. Springer, 1979.

³ L. Cooper and C. Lombardo. [IBM Explores Sale of IBM Watson Health](#). *The Wall Street Journal*, Feb. 18, 2021.

This course is designed to introduce research topics in human-AI interaction. Students at all levels are welcome. You do not need to have a strong ML or HCI background to take this class. Senior undergrads, master and PhD students who are interested in writing a thesis or doing a research-based project are especially encouraged to take this class. If you are not sure about your qualification for this course, feel free to research out to the instructor via email or stop by 3154H.

Activities will include lectures, presentations of research papers, and discussion of relevant literature in each field. You should expect to present two research papers during the semester. You also need to write a short paper review (2 or 3 paragraphs) in the form of comments and questions and post it on Piazza before each paper discussion. There will be a course project, in which you will work in groups to design and carry out research projects related to human-AI interaction. During the semester, we will have two guest speakers from other universities and industry to present their current research in human-AI interaction. This is a good opportunity to connect with faculty from other universities and learn about their research.

Learning Resources, Technology, Texts

No textbook is required for this course. Useful online learning resources include:

- Christoph Molnar, [Interpretable Machine Learning: A Guide for Making Black Box Models Explainable](#)
- Jakob Nielson, [10 Usability Heuristics for User Interface Design](#)
- Google PAIR, [People + AI Guidebook](#)
- Max Welling, [A First Encounter with Machine Learning](#)
- Ian Goodfellow, Yoshua Bengio, Aaron Courville, [Deep Learning](#)
- [ML + Design](#) (a website with a collection of books, articles, tools, and tutorials on ML and design)
- Vishal Maini and Samer Sabri, [Machine Learning for Humans](#)

We will use Piazza to post paper review comments, ask & answer questions, make announcements, look for teammates, etc. Please join this course on Piazza. You can find this course by searching for "CS59200 Human-AI Interaction" on Piazza or use [this signup link](#).

All lecture slides, student presentation slides, and other materials will be posted on [Purdue BrightSpace](#).

Learning Outcomes

At the end of this course, students should be able to:

- identify and understand the problem statement, research questions, methods, findings, and contributions in a research paper
- critically assess the contributions of a paper
- design and implement interactive systems with AI components
- evaluate an interactive AI-based system, especially through user studies in the lab or on a crowdsourcing platform like Amazon Mechanical Turks
- know the style of academic writing, especially in HCI and Software Engineering
- make and deliver academic presentations to the public

Grading

Reading assignments [20%]

Paper presentation [20%]

Final project [50%]

Quizzes [5%]

Class participation and discussion [5%]

Course Schedule

Week 1. Introduction to Human-AI Interaction	
Mon 8/23	Lecture: Introduction, logistics, and the design argument framework
Wed 8/25	Lecture: The role of humans in AI Paper reading: <ul style="list-style-type: none"> • Amershi et al., Power to the People: The Role of Humans in Interactive Machine Learning (AI Magazine 2014)
Fri 8/27	Lecture: A historical perspective of human-AI interaction Paper reading: <ul style="list-style-type: none"> • Shneiderman and Maes, Direct Manipulation vs. Interface Agents (Interactions 1997)
Sample course project ideas released by the instructor on 8/29	
Week 2: Human Needs, Perceptions, and Experiences of Using AI	
Mon 8/30	Lecture: Everyday AI applications Paper reading: <ul style="list-style-type: none"> • When People and Algorithms Meet: User-reported Problems in Intelligent Everyday Applications (IUI 2019)
Wed 9/1	Lecture: Conversational agents Paper reading:

	<ul style="list-style-type: none"> Luger and Sellen, “Like Having a Really Bad PA”: The Gulf between User Expectation and Experience of Conversational Agents (CHI 2016)
Fri 9/3	<p>Lecture: AI-based programming tools</p> <p>Paper reading:</p> <ul style="list-style-type: none"> Weisz et al., Perfection Not Required? Human-AI Partnerships in Code Translation (IUI 2021)
Paper presentation sign-up due on 9/3 midnight	
Week 3: Human Needs, Perceptions, and Experiences of Using AI	
Mon 9/6	No Class (Labor Day)
Wed 9/8	(Student Presentation) Hohman et al., Gamut: A Design Probe to Understand How Data Scientists Understand Machine Learning Models (CHI 2019)
Fri 9/11	(Student Presentation) Cai et al., "Hello AI": Uncovering the Onboarding Needs of Medical Practitioners for Human-AI Collaborative Decision-Making (CSCW 2019)
Week 4. Design Principles and Guidelines for Human-AI interaction (9/13-9/17)	
Mon 9/13	<p>Lecture: Guidelines for human-AI interaction</p> <p>Paper reading:</p> <ul style="list-style-type: none"> Amershi et al., Guidelines for Human-AI Interaction (CHI 2019)
Wed 9/15	<p>Lecture: Challenges of human-AI interaction</p> <p>Paper reading:</p> <ul style="list-style-type: none"> Yang et al., Re-examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design (CHI 2020)
Fri 9/17	<p>Lecture: A historical perspective of human-AI interaction design</p> <p>Paper reading:</p> <ul style="list-style-type: none"> Horvitz, Principles of Mixed-Initiative User Interfaces (CHI 1999)
Week 5. Communicating Model Confidence and Uncertainty	
Mon 9/20	<p>(Student Presentation) The effect of displaying system confidence information on the usage of autonomous systems for non-specialist applications: A lab study. (CHI 2016)</p> <p>Quiz 1 at the end of the class</p> <p>Quiz 1 will cover contents from Week 1 to Week 4</p>
Wed 9/22	(Student Presentation) Kay et al., When (ish) is My Bus? User-centered Visualizations of Uncertainty in Everyday, Mobile Predictive Systems (CHI 2016)
Fri 9/24	(Student Presentation) Bhatt et al., Uncertainty as a Form of Transparency: Measuring, Communicating, and Using Uncertainty (AEIS 2021)
Project proposal due on 9/27 midnight	

Week 6. Supporting Model Customization, Refinement, and Correction	
Mon 9/27	(Student Presentation) Kulesza et al., Principles of Explanatory Debugging to Personalize Interactive Machine Learning (IUI 2015)
Wed 9/29	(Student Presentation) Koh et al., Concept Bottleneck Models (ICML 2020)
Fri 10/1	(Student Presentation) Li et al., Multi-Modal Repairs of Conversational Breakdowns in Task-Oriented Dialogs (UIST 2020)
Week 7. Explaining Model Behavior	
Mon 10/4	(Student Presentation) Cai et al., The effects of example-based explanations in a machine learning interface (IUI 2019)
Wed 10/6	(Student Presentation) Wexler et al., The What-If Tool: Interactive Probing of Machine Learning Models (VAST 2019)
Fri 10/8	(Student Presentation) Ribeiro et al., "Why Should I Trust You?": Explaining the Predictions of Any Classifier (KDD 2016)
Week 8. Actively Eliciting and Incorporating User Feedback into ML	
Mon 10/11	No Class (October Break)
Wed 10/13	(Student Presentation) Siveraman et al., Active Inductive Logic Programming for Code Search (ICSE 2019)
Fri 10/15	(Student Presentation) Yao et al., Interactive Semantic Parsing for If-Then Recipes via Hierarchical Reinforcement Learning (AAAI 2019)
Week 9. Interactive Visual Analytics for Machine Learning	
Mon 10/18	(Student Presentation) Strobel et al., LSTMVis: A Tool for Visual Analysis of Hidden State Dynamics in Recurrent Neural Networks (TVCG 2017) Quiz 2 at the end of the class Quiz 2 will cover contents from Week 5 to Week 8
Wed 10/20	(Student Presentation) Xu et al., mTSeer: Interactive Visual Exploration of Models on Multivariate Time-series Forecast (CHI 2021)
Fri 10/22	(Student Presentation) Hohman et al., Visual Analytics in Deep Learning: An Interrogative Survey for the Next Frontiers (TVCG 2019)
Week 10. AI Ethics, Fairness, and Equity	
Mon 10/25	(Student Presentation) Buolamwini and Gebru, Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification . FACCT 2018

Wed 10/27	(Student Presentation) Cabrera et al., FAIRVIS: Visual Analytics for Discovering Intersectional Bias in Machine Learning (VIS 2019)
Fri 10/29	(Student Presentation) Galhotra et al., Fairness testing: testing software for discrimination (ESEC/FSE 2017)
Mid-point project summary due on 10/31 midnight	
Week 11. Human-AI Co-creation	
Mon 11/1	(Student Presentation) Clark et al., Creative writing with a machine in the loop: Case studies on slogans and stories (IUI 2018)
Wed 11/3	(Student Presentation) Swearngin et al., Scout: Rapid Exploration of Interface Layout Alternatives through High-Level Design Constraints (CHI 2020)
Fri 11/5	(Student Presentation) Louie et al., Novice-AI Music Co-Creation via AI-Steering Tools for Deep Generative Models (CHI 2020)
Week 12. Calibrating User Trust and Reliance on AI	
Mon 11/8	(Student Presentation) Bucinca et al., To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-assisted Decision-making (CSCW 2021) Quiz 3 at the end of the class Quiz 3 will cover contents from Week 9 to Week 11
Wed 11/10	(Student Presentation) Häuslschmid et al., Supporting Trust in Autonomous Driving (IUI 2017)
Fri 11/12	(Student Presentation) Kocielnik et al. Will You Accept an Imperfect AI? Exploring Designs for Adjusting End-User Expectations of AI Systems. (CHI 2019)
Week 13. More on XAI, IML, and Transparency	
Mon 11/15	(Student Presentation) Olah et al., The Building Blocks of Interpretability (Distill 2018)
Wed 11/17	(Student Presentation) Schneider et al., ExplAIn Yourself! Transparency for Positive UX in Autonomous Driving (CHI 2021)
Fri 11/19	(Student Presentation) Koh and Liang. Understanding Black-box Predictions via Influence Functions (ICML 2017)
Week 14. More on Human-AI Partnership	
Mon 11/22	(Student Presentation) Nguyen et al., Believe it or not: Designing a Human-AI Partnership for Mixed-Initiative Fact-Checking (UIST 2018)
Wed 11/24	No Class (Enjoy Thanksgiving!)

Fri 11/26	No Class (Enjoy Thanksgiving!)
Week 15. Heuristics, Biases, and Mental Models of AI Agents	
Mon 11/29	Lecture: Heuristics and biases in human decision making Paper reading: <ul style="list-style-type: none"> Human biases Kahneman and Tversky, Judgment under Uncertainty: Heuristics and Biases (Science 1974)
Wed 12/1	Lecture: The role of mental models in HAI Paper reading: <ul style="list-style-type: none"> Beyond Accuracy: On the Role of Mental Models in Human-AI Teams (HCOMP 2019)
Fri 12/3	Lecture: Class summary and review all papers we have read Quiz 4 at the end of the class Quiz 4 will cover contents from Week 12 to Week 15
Week 16. Final Project Presentation	
Final project report due on 12/15 midnight	

Paper Reading Assignment, Presentation, and Discussion

For each lecture, you should expect to read one research paper on a specific topic in Human-AI Interaction. I will provide some optional readings related to the topic but you are not required to read them. The optional readings are mostly for students who are particularly interested in the topic or who are doing a course project in the topic.

For each required paper reading, you need to submit a short paper review (one or two paragraphs) in the form of questions and comments on Piazza before the class. The grading of your paper review will depend on the overall quantity and quality of your questions and comments. As you read a paper or write your review, focus on the following perspectives.

- **Motivation of the work.** If the paper presents a new tool, who are the target users? Do they really need such a tool? What pain points does this tool address for those users? If the paper presents an empirical study, what are the research questions this study aims to answer? How important are these studies? Who will care about the findings and why should they care?
- **Novelty and significance of the work.** What is new here? What are the main contributions of the paper? What did you find most interesting?

- **Limitations, flaws, and blind spots.** Are there any unrealistic or false assumptions about the target users or the approach? Are there flaws or mistakes in the tool design, technical approach, or the study design?
- **Future work.** How would you improve on this work? Does this paper inspire any new ideas in your own research?

You should expect to present two research papers during the course. The instructor will ask students to sign up papers to present by the end of the second week. The instructor will present the rest of the unselected papers during the course.

Each paper presentation should be no more than 40 minutes, so we can have enough time for discussion. The presentation should focus on elaborating the motivation, related work, tool/study design, research questions, findings, limitations, and future work of the assigned paper. To make your presentation more insightful, try to center your presentation based on the literature and tell the audience why this work is proposed in the first place, how it advances people's understanding about a topic, and how it is different from other related work in the past. You are also encouraged to connect the assigned paper to your own research. You should prepare for a set of questions (either came up by yourself or based on questions other students post on Piazza) and co-lead an in-class discussion with the instructor based on these questions after the presentation.

The in-class discussion will follow the think-pair-share format.

- 1) Think. The presenter or the instructor will provoke students' thinking with a question. The students should take one or two minutes just to THINK about the question.
- 2) Pair. Using designated partners (such as with Clock Buddies), nearby neighbors, or a deskmate, students PAIR up to talk about their answers with each other. They compare their mental or written notes and identify the answers they think are best, most convincing, or most unique.
- 3) Share. After students talk in pairs for a few minutes, the presenter or instructor will call for pairs to SHARE their thinking with the rest of the class.

Course Project Instructions

You are expected to work on a course project either alone or in groups (no more than 3 students in a group). You can pick any topics related to human-AI interaction, e.g., building an interactive intelligent application, building a new user interface for an existing ML technique, evaluating existing AI-base techniques with user studies, etc. At

the end of Week 1, I will release a list of sample project ideas with paper references, but feel free to work on your own ideas. Between Week 2 and Week 5, please stop by during office hours to discuss your project ideas with the instructor to get early feedback on the relevance, novelty, feasibility, and significance of your ideas.

A short project proposal is due on Sept 27 midnight (Week 5). This proposal should describe the project idea, the motivation of this idea, and (optional) a usage scenario if you propose to build a new tool. The proposal could be any length but no longer than 4 pages. It will be evaluated based on the quality of the idea and writing, not the length of your writing.

A mid-point project summary is due on Oct 31 midnight (Week 9). This summary should describe the envisioned approach/methodology/design as well as which parts have been done so far. The summary could be any length but no longer than 4 pages.

In Week 16 (Dec 6-Dec 10), each team will deliver a presentation of their project. The presentation will be about 20 minutes. You will get another 5 minutes for Q&A.

A final project report is due on Dec 15 midnight (max 10 pages plus references). Your final project report should be built upon your proposal and project summary. Feel free to reuse sections from those two reports in your final report. You may include an appendix beyond 10 pages, but your paper must be understandable without it. Submissions should be in the [ACM format](#).

Your final report should be structured like a conference paper. It should contain:

- Abstract
- A well-motivated introduction
- Related work with proper citations
- Description of your methodology
- Evaluation results
- Discussion of your approach, threats to validity, and additional experiments
- Conclusions and future work

If you are doing a project that involves implementation, please include a link to your Github repository in your final report. Please also add a README file in your repository to describe how to run and test your code.

Important dates:

- Project proposal due on 9/27 midnight
- Mid-point project summary due on 10/31 midnight

- Final project presentation on Week 16 (12/6-12/10)
- Final project report due on 12/15 midnight

Quizzes

We will have four quizzes during this course. Each quiz will assess your understanding about the research topics we have covered in the previous two or three weeks. The quizzes will include multiple-choice questions and open-ended questions. To prepare for the quizzes, make sure (1) you have read the papers, (2) review the slides from the instructor or the paper presenter, (3) understand the methodologies, findings, and contributions of each paper. The instructor cannot accommodate quizzes on a different date. However, we will count the best 3 quizzes out of 4 quizzes. Each quiz should take about 5 to 10 minutes.

Course Policies and Expectations

Late submissions

All paper review comments must be submitted by 8AM before the class. No late submission is allowed.

For the course project proposal, mid-point project summary, and final project report, late submissions will be accepted with 7.5% decaying credit per day.

Attendance

This course follows Purdue's academic regulations regarding attendance, which states that students are expected to be present for every meeting of the classes in which they are enrolled. While we will not check attendance in each class, we will use other ways including quizzes and class discussion to check your attendance and participation in the class. Your final grade will depend on your participation in the class. Please come to the class continuously, read the assigned papers, and participate in discussions.

If you feel sick, have any symptoms associated with COVID-19, or suspect you have been exposed to the virus, you should stay home and contact the [Protect Purdue Health Center](#). Please also notify the instructor so that the instructor can arrange remote participation for you. If you miss classes because of COVID-related reasons, your final grade will not be affected by your absence of classes. For more guidance on class attendance related to COVID-19 are outlined in the [Protect Purdue Pledge for Fall 2021](#) on the Protect Purdue website.

For other conflicts or absences, such as for many University-sponsored activities and religious observations, the student should inform the instructor of the situation as far in

advance as possible. When the student is unable to make direct contact with the instructor and is unable to leave word with the instructor's department because of circumstances beyond the student's control, and in cases falling under excused absence regulations, the student or the student's representative should contact or go to the [Office of the Dean of Students website](#) to complete appropriate forms for instructor notification. Under academic regulations, excused absences may be granted for cases of grief/bereavement, military service, jury duty, and parenting leave. For details, see the [Academic Regulations & Student Conduct section](#) of the University Catalog website.

Academic Guidance in Event of Quarantine or Isolation

If you must miss class at any point in time during the semester, please reach out to me via email so that we can communicate about how you can maintain your academic progress. If you find yourself too sick to progress in the course, notify your adviser and notify me via email or Brightspace. We will make arrangements based on your particular situation. Please note that, according to [Details for Students on Normal Operations for Fall 2021](#) announced on the Protect Purdue website, "individuals who test positive for COVID-19 are not guaranteed remote access to all course activities, materials, and assignments."

Time commitment

Students are expected to spend no more than 12 hours per week in class or on coursework per week on average. I suggest, early in the semester, setting aside 3 hours to read and complete any assigned work related to the assigned papers before each class, leaving approximately 3 hours per week for group formation, early scouting of research project topics, need-finding research and brainstorming. As we get deeper into the semester, I suggest spending 2 hours per class reading assigned papers and completing related assigned work, leaving approximately 5 hours for building prototypes, conducting user studies, and preparing presentations.

Feedback to the instructor

During this course, I will be asking you to give me feedback on your learning in both informal and formal ways. Occasionally, at the end of a lecture, I will hand out index cards to collect anonymous comments and questions about this class and your learning experience. In the middle of the semester, I will send out an anonymous midpoint survey about how my teaching strategies are helping or hindering your learning. It is very important for me to know your reaction to what we are doing in the class, so I encourage you to respond to these surveys, ensuring that we can create an environment effective for teaching and learning.

Classroom Guidance Regarding Protect Purdue

Any student who has substantial reason to believe that another person is threatening the safety of others by not complying with Protect Purdue protocols is encouraged to report the behavior to and discuss the next steps with their instructor. Students also have the option of reporting the behavior to the [Office of the Student Rights and Responsibilities](#). See also [Purdue University Bill of Student Rights](#) and the Violent Behavior Policy under University Resources in Brightspace.

Academic Integrity

Please read and follow the academic integrity policy from the CS department ([link](#)) and from Purdue University ([link](#)). For paper reading assignments and the course project, you are encouraged to discuss with your classmates and course instructor. However, you should ensure that any written work you submit for evaluation is the result of your own research and writing and that it reflects your own approach to the topic. You must also adhere to standard citation practices in this discipline and properly cite any books, articles, websites, lectures, etc. that have helped you with your work.

Nondiscrimination Statement

Purdue University is committed to maintaining a community which recognizes and values the inherent worth and dignity of every person; fosters tolerance, sensitivity, understanding, and mutual respect among its members; and encourages each individual to strive to reach his or her potential. In pursuit of its goal of academic excellence, the University seeks to develop and nurture diversity. The University believes that diversity among its many members strengthens the institution, stimulates creativity, promotes the exchange of ideas, and enriches campus life. A hyperlink to Purdue's full Nondiscrimination Policy Statement can be found [here](#).

Accessibility

Purdue University is committed to making learning experiences accessible. If you anticipate or experience physical or academic barriers based on disability, you are welcome to let me know so that we can discuss options. You are also encouraged to contact the Disability Resource Center at: drc@purdue.edu or by phone: 765-494-1247.

Mental Health/Wellness Statement

If you find yourself beginning to feel some stress, anxiety and/or feeling slightly overwhelmed, try [WellTrack](#). Sign in and find information and tools at your fingertips, available to you at any time.

If you need support and information about options and resources, please contact or see the [Office of the Dean of Students](#). Call 765-494-1747. Hours of operation are M-F, 8 am- 5 pm.

If you find yourself struggling to find a healthy balance between academics, social life, stress, etc. sign up for free one-on-one virtual or in-person sessions with a [Purdue Wellness Coach at RecWell](#). Student coaches can help you navigate through barriers and challenges toward your goals throughout the semester. Sign up is completely free and can be done on [BoilerConnect](#). If you have any questions, please contact Purdue Wellness at evans240@purdue.edu.

If you're struggling and need mental health services: Purdue University is committed to advancing the mental health and well-being of its students. If you or someone you know is feeling overwhelmed, depressed, and/or in need of mental health support, services are available. For help, such individuals should contact [Counseling and Psychological Services \(CAPS\)](#) at 765-494-6995 during and after hours, on weekends and holidays, or by going to the CAPS office on the second floor of the Purdue University Student Health Center (PUSH) during business hours.

Basic Needs Security

Any student who faces challenges securing their food or housing and believes this may affect their performance in the course is urged to contact the Dean of Students for support. There is no appointment needed and Student Support Services is available to serve students 8 a.m.-5 p.m. Monday through Friday. Considering the significant disruptions caused by the current global crisis as it related to COVID-19, students may submit requests for emergency assistance from the [Critical Needs Fund](#).

Emergency Preparation

In the event of a major campus emergency, course requirements, deadlines and grading percentages are subject to changes that may be necessitated by a revised semester calendar or other circumstances beyond the instructor's control. Relevant changes to this course will be posted onto the course website or can be obtained by contacting the instructors or TAs via email or phone. You are expected to read your @purdue.edu email on a frequent basis.

