Human–AI Interaction (Info 693)

College of Computing & Informatics, Drexel University

Course Syllabus, Fall 2022

Professor
Tim Gorichanaz, Ph.D.
tjg68@drexel.edu
Office: 3675 Market #1118
Student Hours: by appointment

Course Overview

Artificial Intelligence (AI) applications are widespread and already transforming nearly every aspect of society. As AI becomes even more broadly embedded in software, user interface considerations must be included in the AI system development process. This class introduces the unique interface design challenges presented by AI. It explores questions of usability and user experience specific to AI systems, and it reflects more broadly on the relationship between humanity and emerging technologies. Students will practice skills in design, research and writing regarding the human side of AI. Topics include system design, speculative design, confidence and error, explainability, HCI principles, human augmentation/ amplification, data ownership, AI ethics, and embodied AI. No programming experience required.

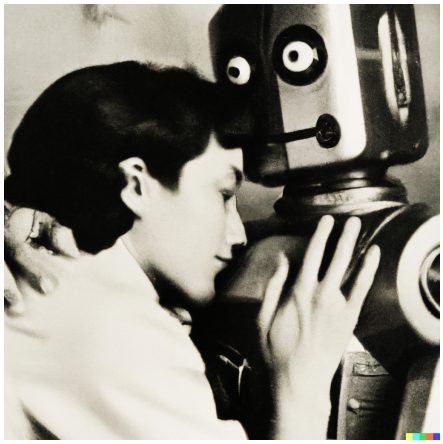

AI-generated image using Dall-E 2 with the prompt, "1960s film still of a human and a robot hugging"

Course purpose

The course will provide you with an opportunity to engage with the theory and practice of designing interfaces for AI-powered computer systems. It is an elective course in the **Human-Centered Computing** concentration designed for graduate students studying **Artificial Intelligence & Machine Learning** or **Human–Computer Interaction & User Experience**. That said, it may be taken as a Free Elective by any student with graduate standing. There are no formal prerequisites, but students with some background in human-centered design will get the most out of this class.

Learning targets

- As a result of your experience in this course, you will be able to:
- Describe the scope of human–AI interaction (HAII), including its foundations, current issues and future potential.
- Analyze debates in HAII using various theoretical frameworks and engaging with the scholarly literature.
- Explain the interface design challenges that AI systems present.
- Prototype and test a new human–AI interface using interface design methods such as speculative design and usability testing.

- Evaluate the arguments for and against the application of AI in various aspects of human society and participate in public discourse on these topics.
- Advocate for particular AI applications and design choices you care about using the knowledge and skills gained in this course, such as ensuring fairness and accessibility in AI systems.

In addition, as your instructor I will help you to:
- Engage in informed discussions with friends and family about artificial intelligence
- Self-reflect on your work and cultivate a growth mindset
- Receive and respond to feedback gracefully
- Develop your skills of good thinking and related habits of mind, such as love of learning, intellectual courage, intellectual humility, and practical wisdom

What I think about teaching and learning

I love being a teacher, I love the material I get to teach, and I love bringing new voices into the fold. So my goal is not to "cover course content," but to welcome you into a community of learning that will serve you well for life. This goal guides my teaching philosophy and my choices as an instructor.

For example, the vast majority of our time together will be in activities and conversations. You will be talking, making, doing, teaching… This is what we call *active learning,* and it will help you develop your skills in collaboration, solidify your learning for the long term, spark creative thinking, and get you to have fun along the way. Long story short, you will not be sitting in class listening to me read PowerPoint slides (though on occasion I may give short lectures to share basic information).

I realize active learning can be challenging, particularly for people who are introverted or have any kind of anxiety. I have been there! I myself am an introvert. I do my best to create a comfortable learning environment where everyone feels psychologically safe and able to participate and learn to their best ability. I hope, over time, you will feel at home in the learning community we will build this quarter.

How to be a part of this class

Learning happens best in community. For our purposes, this means being present and being caring. This quarter, let mutual respect guide our time together, so that everyone can participate and enjoy the class. To be a part of this class, strive to do the following:
- **Attend each class session**, arriving a few minutes early so we can start on time.
- **Read the assigned readings and give yourself time to mull them over**, so you can participate meaningfully in our activities.
- **Do your best to get to know your classmates as the course goes on**. These are your allies—and future colleagues in life. The world is small. If you're on the shy side, challenge yourself to step just outside of your comfort zone when you're able.

Course mechanics

Course structure

This course has two sections under one roof: one online and asynchronous (section 900); and another that meets in person on Thursdays (section 001). The course is organized into weeks, each running Monday to Sunday. You can find an at-a-glance view of each week's topic on the final page of this syllabus. Assignments are always due on Sunday by 9:00 p.m. Eastern. (After 9 o'clock you can do something fun or turn in for a good night's sleep.)

Each week will include readings, videos and activities. For those in the synchronous section (001), we'll do many of these things during class. For those in the online section, all of this can be done according to

your own schedule. There will be some elements of teamwork in the course; online students will be matched with teammates who have a similar schedule as much as possible.

Time commitment

This is a graduate course in a professional school. Drexel estimates that the total workload (readings, assignments, notes, exercises) will take *12 to 15 hours per week*. Be prepared for a serious commitment of attention and effort. It will pay off!

What book you need—none!

There is no required textbook in this course. We will read several articles each week, but these will be made available freely on Blackboard. If you would like to supplement your reading in this course with other books and podcasts on the topic of HAII, see the Additional Resources section below.

Materials and software

If you are interested in the topic of this course and want to deepen your learning with some additional reading, I recommend the following books and podcasts. These are all geared toward a general audience, and they're engaging and readable. If you know of other recommended books along these lines, please share with the class!

Books

- Bostrom, Nick. (2014). Superintelligence: Paths, Dangers, Strategies. Oxford University Press.
- Christian, Brian. (2020). The Alignment Problem: Machine Learning and Human Values. Norton.
- Coeckelbergh, Mark. (2020). AI Ethics. The MIT Press.
- Coleman, Flynn. (2019). A Human Algorithm: How Artificial Intelligence Is Redefining Who We Are. Counterpoint Press.
- Daugherty, Paul R., & H. James Wilson. (2018). Human + Machine: Reimagining Work in the Age of AI. Harvard Business Review Press.
- Harari, Yuval Noah. (2017). Homo Deus: A Brief History of Tomorrow. HarperCollins.
- Hawkins, Jeff. (2021). A Thousand Brains: A New Theory of Intelligence. Basic Books.
- Kissinger, Henry A., Eric Schmidt, & Daniel Huttenlocher. (2021). The Age of AI: And Our Human Future. Little, Brown.
- Marcus, Gary, & Ernest Davis. (2019). Rebooting AI: Building Artificial Intelligence We Can Trust. Vintage.
- Markoff, John. (2016). Machines of Loving Grace: The Quest for Common Ground Between Humans and Robots. HarperCollins.
- Rushkoff, Douglas. Team Human. Norton.
- Russell, Stuart. (2019). Human Compatible: Artificial Intelligence and the Problem of Control. Penguin.
- Tegmark, Max. (2017). Life 3.0: Being Human in the Age of Artificial Intelligence. Knopf.
- Webb, Amy. (2019). The Big Nine: How the Tech Titans and Their Thinking Machines Could Warp Humanity. PublicAffairs.
- Zuboff, Shoshana. (2019). The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power. PublicAffairs.

Podcasts

- AI Nation (from WHYY and NPR)
- AI4Society Dialogues
- Data Stories Podcast
- Radical AI

- Tech Society Podcast

Contacting me

Student–instructor interaction is an important part of any course. I am available to you, and I want to help you succeed in this course, in your program at Drexel—and in life. Please come to me with any questions, problems, discoveries or anything else you'd like to share.

If your question is general and may be of interest to others in the class (e.g., syllabus, readings, schedule, etc.), please post it on the *Questions About the Course* Discussion Board so that others can benefit, or ask during class. With **personal or urgent questions**, you should email me directly (tjg68@drexel.edu). If you have a **technical question**, you will be better off contacting the Instructional Technology Group. See http://www.drexel.edu/irt/help/learn or call (215) 895–1224.

I am available to meet with you by appointment, and feel free to email me at any time. I check email twice a day, and I don't do it at night or on weekends. In our always-on society, it is important to set boundaries—firstly because healthy lives require off-time, and also because our academic activities require uninterrupted periods of time for reading, writing and thinking. Moreover, taking time for rest and pursuing leisure activities have been shown to improve productivity, creativity and accomplishment, as Alex Pang discusses in his book Rest: Why You Get More Done When You Work Less. I hope you will join me in living with more balance.

Assignments and Grading

There are several assignments in this course: **quizzes**, **short papers**, and a **design project**.

First, there will be regular, low-stakes **quizzes** taken during class to gauge your comprehension of the readings for that week and to help us frame our discussion.

As well, you will submit four written **short papers** throughout the term centered around the readings. These papers will be short (about 2 pages each) and will give you the chance to demonstrate your understanding of the readings—and, later in the term, practice applying concepts from the readings.

- The first two papers will be **Discussion Papers,** in which you write a synthesis of the week's readings in Weeks 2 and 3.
- Next, a **Position Paper** will be due in Week 5, where you use evidence from the readings to support your argument for a particular position in a debate.
- Last, you will write an **Op-Ed** for Week 8's class, expressing an informed position on an HAII-related topic framed for a general audience.

The major assignment in this class is a term **design project**, in which you will design an interactive system that involves some form of AI. Examples include GUI software with an intelligent user interface, chatbots and voice assistants, data science applications, robotics, route navigators, etc. The project will involve prototyping the user interface and building a proof-of-concept of the AI algorithm (no coding necessary). You may work individually or in groups of up to 3 people. The project will be developed and submitted in phases throughout the term, as follows:

1. **Proposal**: Identify a problem space, conduct exploratory research, and list system requirements. (Due Week 4)
2. **Conceptual Design**: Create the conceptual model for your system and discuss ethical issues in the design. (Due Week 6)
3. **Concrete Design**: Prototype the user interface using an appropriate method and subject it to evaluation, and create a proof-of-concept of the AI portion of the design. (Due Week 9)
4. **Presentation**: Share your work with the class. (During class in Week 11)

Discussion and participation

Learning can be exciting and, at times, confusing. For students in the face-to-face section of the course, we'll be able to do this when we meet each week. For all students, the Discussion Boards will offer you an opportunity for you to share your journey with your classmates—to ask and answer questions, discuss the readings, share resources, etc. There are several forums that will be open throughout the term, and some forums will correspond to topics we'll discuss in particular weeks (e.g., for sharing our prototypes and giving feedback).

In the past, students have also benefited from creating informal learning groups via the Discussion Boards. If you learn better with colleagues, this is highly recommended! You can make use of Zoom or other software to meet or you could connect in person, geography permitting.

Grading

Research has demonstrated that grades diminish students' learning, decrease students' interest in the subject matter, and prevent students from taking creative risks. Moreover, many students experience anxiety about grades. In this class, we are here to learn, to become interested in design, and to take some creative risks by trying new things (not just following a rubric)—and we certainly don't need more anxiety in our lives. So in this class, we will take a different approach to assessment.

My intention with this class is to help you to work in an organic way, as you will after graduation. So while you will get a final grade in the class, I will not put quantitative grades on individual assignments. Rather, when I review your work, I will ask questions and make comments meant to engage your work rather than simply evaluate it. You, too, will reflect deeply on your work and that of your peers throughout this quarter, and we will discuss your learning and effort as the course progresses. We will do this throughout the course, but there are three things I will ask you to do as anchors for this process:

- **Goal-Setting**: During Week 1, we will set goals for our learning this quarter to give us each a concrete place to aim for in addition to the general course learning outcomes listed above.
- **Midterm Reflection**: Midway through the course, I will provide you with a link to an online form that will guide you through a reflection on your work thus far—particularly with respect to the goals you set for yourself. At the end, you'll be asked what letter grade you would give yourself for your work to date. This is your chance to assess yourself realistically and challenge yourself to improve in the second half of the term. I will respond to your reflection, and we'll have a conversation if our respective assessments do not match.
- **Final Reflection**: Ahead of Exam Week, I will provide you with another link to an online form where you'll complete your final self-reflection for the term. Again, I'll ask you what grade you would give yourself. I prefer to give everyone the grade they would give themselves, I do reserve the right to make adjustments.

For reference, here is my interpretation of the letter grades:

- *A – Excellent:* You have demonstrated significant progress toward the learning targets for this course while exceeding expectations in effort, participation and results. Exceeds expectations in class participation.
- *B – Good:* You have demonstrated progress toward the learning targets for the course while meeting expectations. Good participation.
- *C – Acceptable:* You are making progress toward most of learning targets for the course. Hit-or-miss participation.
- *D – Poor:* You do not meet the expectations of one or more of the learning targets for the course. Poor participation.
- *F – Failing:* You demonstrate no progress toward the learning targets for the course. Poor participation.

I know this process is quite different from how we usually think about grades. If any of this causes more anxiety than it alleviates, contact me at any time to discuss your progress in the course. As the course goes on, you'll be able to track your progress in the My Grades section on Drexel Learn.

How to get a good grade

It's easy for school to become a game. Check the boxes, don't rock the boat, say nice things, and get that A. In this class, we're going to shake that up a bit. Like all shake-ups, it may be uncomfortable but in the end it will be a chance for growth. To get a good grade in this class, pay special attention to these areas:

- **Attendance and Participation**: A lot of learning will happen in our time together, and I don't want you to miss out. Plan to attend regularly, coming prepared and with energy to participate.
- **Assignments**: Use your assignments as an opportunity to demonstrate to me what you are learning. In your group project submissions, mini-essays and self-reflections, strive to integrate insights that have come up in classroom conversations as well as the course readings.
- **Professionalism**: School is practice. Here's your chance to try out new systems and figure out what will work for you in life. Give it a serious try, and bring your best self to this class. That way you'll get the most out of this experience.

If you are worried about getting a good grade in this class, your best strategy should be to do the readings, ask questions often, complete the assignments diligently and on time, and engage earnestly with all your classmates.

Policies

Academic integrity

You are expected to conduct yourself in a respectful manner as befitting the university environment. This includes academic integrity. In this course, as with any Drexel course, cheating will not be tolerated. This includes plagiarism (using others' intellectual work without reference). All work you submit must be your own work, with sources properly cited. Any plagiarism or other academic dishonesty will result in a sanction that may extend to failing the course. I am obligated to report incidents of cheating (including plagiarism) to Drexel administration. A student who is found in violation twice (even if in two different courses) will be expelled from the university. For more information, please refer to the Provost academic integrity policy or to resources regarding Student Conduct and Community Standards.

Dropping the course

If you are considering whether to continue your enrollment in the course, please refer to the Course Add/Drop Policy and the Course Withdrawal Policy.

Changes to the syllabus

I reserve the right to make changes to this course or its syllabus during the quarter if circumstances warrant such a change. Topics, readings and dates are subject to change, but only if necessary. Additional topics may be discussed as issues and ideas arise in the news and in discussion. All changes will be provided to students in writing as far in advance as possible.

Student conduct

Drexel University adopted a student conduct policy requiring that all students have the responsibility to be aware of, and abide by, the University's policies, rules, regulations, and standards of conduct. The Student Conduct and Community Standards policy information is available in the Official Student Handbook.

Appropriate use of course materials

It is important to recognize that some or all of the course materials provided to you may be the intellectual property of Drexel University, the course instructor, or others. Use of this intellectual property is governed by Drexel University policies, including the Acceptable Use Policy. Briefly, this policy states that course materials, including recordings, provided by the course instructor may not be copied, reproduced, distributed or re-posted. Doing so may be considered a breach of this policy and will be investigated and addressed as possible academic dishonesty, among other potential violations. Improper use of such materials may also constitute a violation of the University's Code of Conduct.

Participating in course evaluations

Student evaluations are a required element of every course. Evaluation forms are completely anonymous. They are confidentially used to make improvements in our curriculum and teaching. They are also used by administration in evaluating faculty performance, and in decisions about promotion, tenure and retention. Please take part in course evaluations.

Support and Recommendations

If you are experiencing anxiety, depression or other issues

Drexel offers free and confidential support for anxiety-related problems, depression, family concerns, relationship issues, adjustment issues, eating disorders, alcohol- and drug-related problems, and questions about gender and sexual identity, all through the Drexel Counseling Center. The Counseling Center is located at Suite 201 in the Creese Student Center at 3210 Chestnut. The telephone number is (215) 895–1415. **For emergencies, or to reach an on-call counselor after regular business hours, please call (215) 416–3337**. Learn more on the Counseling Center website.

If you need technical support

Get 24/7 technical support for Blackboard Learn from the Instructional Technology group online or by calling (215) 895–1224. For any other technical support (email, logins, etc.), Drexel University IT is here for you. You can contact them through email at consult@drexel.edu, by phone at (215) 895–2020, or by submitting the online Problem Report Form.

Support for equality and diversity

Drexel University strives to promote an environment of equality of opportunity and compliance with university policies and federal, state and local laws prohibiting discrimination based upon race, color, religion, gender, marital status, pregnancy, national origin, age, disability and veteran status. If you have a question or complaint concerning discrimination, harassment, and/or retaliation, contact the Office of Equality and Diversity online or at (215) 895–1405.

Coaching, mentorship and tutoring

The Center for Learning and Academic Success Services (CLASS) serves as the organizing department for a variety of programs and services that promote coaching, peer mentoring and tutoring at Drexel. The Center is located on campus at the Creese Student Center (3210 Chestnut Street), Suite 050.

Campus activities and community

Find the Student Handbook, conduct and community standards, and the Counseling Center at on the Student Life website. Consult this site for information on campus activities and student programs.

English help

The English Language Center offers English language instruction and support services to students, especially those who speak English as a second language. They are located at 229 N. 33rd Street. The telephone number is (215) 895-2022.

If you have a disability or are facing other challenges

The Office of Disability Resources (ODR) team coordinates reasonable accommodations for all Drexel students with disabilities to ensure a level playing field on which they can succeed. ODR can also facilitate temporary adjustments for students with short-term impairments such as those due to accident, injury or illness. If you are a student with a disability, you are encouraged to register with ODR to request reasonable accommodations. This office is here to work with you, so reach out to them for assistance.

For any accommodations to be made, you will need to request a current Accommodations Verification Letter (AVL) in the ClockWork database (if you are new to the system, start by clicking "Online Intake"). These requests are received by ORD, who then issues the AVL to the appropriate contacts, such as

professors. For additional information, visit the DR website, reach them by phone at (215) 895–1401 or email at disability@drexel.edu, or visit them in person in Suite 228 in the Main Building.

Free health services

The Student Health Center is located at 3401 Market St, Ste 105. You can call them at (215) 220–4700.

Career counseling

CCI Career Services offers help with job placement, job postings and credentialing. Outside our college, the Steinbright Career Development Center (SDLC) offers individualized career counseling, career fairs, career programs and resume workshops. The office is located at 3201 Arch Street, Suite 250. The telephone number is (215) 895–2185.

Course Schedule

In this class, all assignments are due on Sunday at the end of the given week by 6 p.m. on Wednesday (the start of class each week). You can find an at-a-glance table of the course schedule on the last page of this syllabus. Here you can find a detailed list of the topics and readings to be explored in each week. I hope these readings, along with our other activities, will spark lively discussions—perhaps even debate! Across this list, I recommend reading the texts in the order they are listed.

Week 1: Introduction

Readings:
- Langley, P. (2006). Intelligent behavior in humans and machines [technical report]. Computational Learning Laboratory, Stanford University.
- Markoff, J. (2015). Between human and machine. In Machines of loving grace: The quest for common ground between humans and robots (ch. 1; pp. 1–18). HarperCollins.
- Shneiderman, B., & Maes, P. (1997). Direct manipulation vs. interface agents. Interactions, 4(6), 42–61.
- Koponen, J. M. (2019). UI + AI: Combine user experience design with machine learning to build smarter products. VentureBeat.
- Lin, J. (2020). Rethinking human–AI interaction [blog post]. Perrell.

Week 2: Interactive AI/ML Systems

Due:
- Goal Setting
- Discussion Paper

Readings:
- Jhaver, S., Birman, I., Gilbert, E., & Bruckman, A. (2019). Human-machine collaboration for content regulation: The case of Reddit automoderator. ACM Transactions on Computer-Human Interaction, 26(5), article 31.
- Simard, P. Y., Amershi, S., Chickering, D. M., Pelton, A. E., Ghorashi, S., Meek, C., … & Wernsing, J. (2017). Machine teaching: A new paradigm for building machine learning systems [technical report]. https://arxiv.org/abs/1707.06742
- Laird, J. E., Gluck, K., Anderson, J., Forbus, K. D., Jenkins, O. C., Lebiere, C., … & Kirk, J. R. (2017). Interactive task learning. IEEE Intelligent Systems, 32(4), 6–21.
- MacLellan, C. J., Harpstead, E., Marinier III, R. P., & Koedinger, K. R. (2018). A framework for natural cognitive system training interactions. Advances in Cognitive Systems, 6(1), 1–16.
- Wang, G. (2019). Humans in the loop: The design of interactive AI systems [blog post]. Institute for Human-Centered AI, Stanford University.

Week 3: UX Design for AI

Due:
- Discussion Paper

Readings:

- Verganti, R., Vendraminelli, L., & Iansiti, M. (2020). Design in the age of AI [technical report]. Harvard Business School.
- Osipova, A. (2018). How to design an artificial intelligent system, part 1: Concept development. Toward Data Science.
- Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., ... & Horvitz, E. (2019, May). Guidelines for human–AI interaction. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. https://doi.org/10.1145/3290605.3300233
- Horvitz, E. (1999, May). Principles of mixed-initiative user interfaces. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99) (pp. 159–166). https://doi.org/10.1145/302979.303030
- Auger, J. (2014). Living with robots: A speculative design approach. Journal of Human–Robot Interaction, 3(1), 20–42.

Week 4: UX Prototyping for AI

Due:
- Project Proposal

Readings:
- Yang, Q., Steinfeld, A., Rosé, C., & Zimmerman, J. (2020). Re-examining whether, why, and how human–AI interaction is uniquely difficult to design. Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. https://doi.org/10.1145/3313831.3376301
- Druga, S., Williams, R., Breazeal, C., & Resnick, M. (2017). "Hey Google is it OK if I eat you?" Initial explorations in child–agent interaction. Proceedings of the 2017 Conference on Interaction Design and Children (pp. 595-600). https://doi.org/10.1145/3078072.3084330
- Tsech, N. (2020). The design of AI-based products: 13 things to consider. UX Planet.
- Bek, S. (2018). How to use speculative prototyping to explore the future. Spotless Says.
- Explore the HAX (Human–AI Interaction) Toolkit from Microsoft: https://www.microsoft.com/en-us/haxtoolkit/toolkit-overview/

Week 5: Human Values and AI

Due:
- Position Paper

Readings:
- Guszcza, J., Lee, M. A., Ammanath, B., & Kuder, D. (2020). Human values in the loop: Design principles for ethical AI. Deloitte Review, 26.
- Keyes, O., Hutson, J., & Durbin, M. (2019). A mulching proposal: Analysing and improving an algorithmic system for turning the elderly into high-nutrient slurry. Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems. https://doi.org/10.1145/3290607.3310433
- Zhu, H., Yu, B., Halfaker, A., & Terveen, L. (2018). Value-sensitive algorithm design: Method, case study, and lessons. Proceedings of the ACM on Human-Computer Interaction, 2(CSCW), article 194. https://doi.org/10.1145/3274463
- Kim, Y., Reza, M., McGrenere, J., & Yoon, D. (2021). Designers characterize naturalness in voice user interfaces: Their goals, practices, and challenges. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. https://doi.org/10.1145/3411764.3445579

Week 6: Explainable AI

Due:
- Project: Conceptual Design
- Midterm Self-Reflection

Readings:
- Burnett, M. (2020). Explaining AI: Fairly? Well? Proceedings of the 25th International Conference on Intelligent User Interfaces (IUI '20). https://doi.org/10.1145/3377325.3380623
- Krause, J., Perer, A., & Ng, K. (2016). Interacting with predictions: Visual inspection of black-box machine learning models. Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems. https://doi.org/10.1145/2858036.2858529
- Langley, P. (2019). Explainable, normative, and justified agency. Proceedings of the AAAI Conference on Artificial Intelligence, 33(1), 9775–9779.
- Johs, A. J., Agosto, D. E., & Weber, R. O. (2020). Qualitative investigation in explainable artificial intelligence: A bit more insight from social science [working paper]. https://arxiv.org/abs/2011.07130

Week 7: When Things Go Wrong

Readings:
- Kocielnik, R., Amershi, S., & Bennett, P. N. (2019). Will you accept an imperfect AI? Exploring designs for adjusting end-user expectations of ai systems. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. https://doi.org/10.1145/3290605.3300641
- Cai, C. J., Reif, E., Hegde, N., Hipp, J., Kim, B., Smilkov, D., ... & Terry, M. (2019). Human-centered tools for coping with imperfect algorithms during medical decision-making. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. https://doi.org/10.1145/3290605.3300234
- Shane, J. (2018). Machine learning failures—for art! [video lecture]. https://www.youtube.com/watch?v=yneJIxOdMX4

Week 8: Data Ownership and Privacy

Due:
- Op Ed

Readings:
- Fiesler, C., & Hallinan, B. (2018). "We are the product": Public reactions to online data sharing and privacy controversies in the media. Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. https://doi.org/10.1145/3173574.3173627
- Kerry, C. F. (2020). Protecting privacy in an AI-driven world [technical report]. Artificial Intelligence and Emerging Technology Initiative, The Brookings Institution. https://www.brookings.edu/research/protecting-privacy-in-an-ai-driven-world/
- Ramsøy, J. (2018). Why data ownership matters in the age of AI. Machine Design. https://www.machinedesign.com/automation-iiot/article/21837140/why-data-ownership-matters-in-the-age-of-ai

- Teich, D. A. (2020). Artificial intelligence and data privacy—turning a risk into a benefit. Forbes. https://www.forbes.com/sites/davidteich/2020/08/10/artificial-intelligence-and-data-privacy--turning-a-risk-into-a-benefit

Week 9: History and Perspectives

Due:
- Project: Concrete Design

Readings:
- Myers, B. A. (1996). A brief history of human computer interaction technology. Interactions, 5(2), 44–54.
- Turing, A. (1950). Computing machinery and intelligence. Mind, 59(236), 433.
- Licklider, J. C. (1960). Man-computer symbiosis. IRE Transactions on Human Factors in Electronics, 1(1), 4–11.
- Wiener, N. (1960). Some moral and technical consequences of automation. Science, 131(3410), 1355–1358.

Week 10: No Class, Thanksgiving Break

What are you grateful for?

Week 11: Future of Human–AI Interaction

Due:
- Project: Presentation

Readings:
- Ito, J. (2017). Resisting reduction: A manifesto. Journal of Design and Science. https://doi.org/10.21428/8f7503e4
- Bigham, J. (2019). The coming AI autumn [blog post]. https://jeffreybigham.com/blog/2019/the-coming-ai-autumnn.html
- Schwarz, J. (2020). No user interface and data-driven design: How AI is changing the UI/UX landscape [blog post]. DVG Interactive. https://www.dvginteractive.com/no-user-interface-and-data-driven-design-how-ai-is-changing-the-ui-ux-landscape/
- Rainie, L., Anderson, J., & Vogels, E. A. (2021). Experts doubt ethical AI design will be broadly adopted as the norm within the next decade. Pew Research Center. https://www.pewresearch.org/internet/2021/06/16/experts-doubt-ethical-ai-design-will-be-broadly-adopted-as-the-norm-within-the-next-decade/

Exam Week

Due:
- Final Self-Reflection

Acknowledgements

The design of this course is indebted to the courses on Human–AI Interaction offered at

Carnegie Mellon University, Virginia Tech, the University of Texas at Austin, and Williams College. These courses cover additional topics and readings that may be interesting to you as you continue your journey learning about HAII.

Course Schedule

In this course, the weeks run Monday to Sunday. Remember, assignments listed in a given week are always due by Thursday at 9:00 p.m. Eastern. (After 9 o'clock you can do something fun or turn in for a good night's sleep.)

Wk	Dates	Meeting	Topic	Assignments
1	**Sep 19–25**	**Sep 22**	Introduction	Goal Setting
2	**Sep 26–Oct 2**	**Sep 29**	Interactive AI/ML Systems	Discussion Paper
3	**Oct 3–9**	**Oct 6**	UX Design for AI	Discussion Paper
4	**Oct 10–16**	**Oct 13**	UX Prototyping for AI	Project: Proposal
5	**Oct 17–23**	**Oct 20**	Human Values and AI	Position Paper
6	**Oct 24–30**	**Oct 27**	Explainable AI	Project: Conceptual Design Midterm Self-Reflection
7	**Oct 31–Nov 6**	**Nov 3**	When Things Go Wrong	
8	**Nov 7–13**	**Nov 10**	Data Ownership and Privacy	Op-Ed
9	**Nov 14–20**	**Nov 17**	History and Perspectives	Project: Concrete Design
10	**Nov 21–27**	–	–	
11	**Nov 28–Dec 4**	**Dec 1**	The Future of HAII	Project: Presentation
Ex	–	–	–	Final Reflection *due Tue, Dec 6*