

Human-Computer Interaction

Intro to CS 770

Professor Bilge Mutlu

Today's Agenda

- Who we are
- The three strands, decoupled by day
- *Fundamentals are the medium* – the AI framing
- Syllabus, grading, crit model, AI-use policy
- Seeding the project + topic cards

Your Instructor

Bilge Mutlu

*Professor of Computer Sciences,
Psychology, & Industrial Engineering*

PhD, 2009, Carnegie Mellon University

bilge@cs.wisc.edu · bilgemutlu.com

Office hours: Fridays 11 am-12 pm ·
Morgridge Hall 7631



Your TA

Hailey Johnson

*PhD candidate, Computer Sciences –
People and Robots Lab*

hljohnson22@wisc.edu ·
haileyjohnson.github.io

Office hours: Tuesdays 1–2 pm ·
Morgridge Hall 4563



How about you?

*Give us your name, program,
year.*

Three strands, decoupled by day

Monday

Seminar – one enduring HCI topic per week

Friday

Methods – M1-M12, the research toolkit

Across both

Project – a term-long team study

Why decouple the strands?

The gap is the point.

- Seminar gives you the **questions** worth asking
- Methods give you the **instruments** to answer them
- The project is where they **collide** – you practice both on one real study

Watch this

A researcher delegates a literature search, a schedule, and a collaboration – to an agent.

As you watch, track one thing: *what does the **person** actually do?*

► [play \(YouTube\)](#) · [backup: Internet Archive](#)



1987.

Why doesn't it exist yet?

The durable questions

- **Delegation** – what do you hand over?
- **Common ground** – what does it know that you know?
- **Trust & reliance** – when

The barriers, 39 years on

- **Privacy** – always listening; no wake word, ever
- **Social & situational** – voice fails in an open office
- **Trust** – would you let it

^{kn} Newendorp, A.K., Sanaei, M., Perron, A.J., Sabouni, H., Javadpour, N., Sells, M.J., Nelson, K.T., Dorneich, M.C., & Gilbert, S.B. (2024). Apple's Knowledge Navigator: Why Doesn't that Conversational Agent Exist Yet? *CHI '24*, Article 776, 14 pp. doi:10.1145/3613904.3642739

**Fundamentals are the *medium*,
not the message**

The AI framing

This is a **research** course that uses enduring HCI fundamentals as the instrument for reading the current AI moment.

Every seminar week runs one frame:

Enduring fundamental → how AI is changing it → depth (a claim to argue with)

Two rules that hold all term

- Each seminar closes with a **guardrails coda** – oversight, accountability, appropriate reliance.
 - Design & usability content enters only as **"lens in, craft out"** – a theoretical lens or a research method, **never** as practice craft.
- 👉 We are not here to *build* good interfaces. We are here to *study* interaction.

What is HCI research?

"...a discipline concerned with the design, evaluation and implementation of interactive computing systems for human use and with the study of major phenomena surrounding them."

– ACM

770 lives in the second half: *the study of major phenomena.*

570 vs. 571 vs. 770

770

- Research methods
- Grads across campus
- Project-based
- No technical background

570

- Design methods
- Undergrads
- Project-based
- No technical background

571

- Design/building
- CS undergrads
- Assignment-based
- Needs CS-400 & JS

Learning Goals – *asking*

- Pose research questions, situate them in the literature, and identify gaps
- Meet seminal HCI research across topics
- Match a research approach to a question – qualitative, quantitative, hybrid

Learning Goals – *doing*

- Choose measures; design surveys & scales; assess reliability and validity
- Analyze qualitative and quantitative data
- Carry out an original project and **write it up as a paper**

Setting Expectations

- Read a lot – **1-3 papers + a book chapter** most weeks
- Budget **10-15 hrs/week** (a real 3-credit graduate load)
- A substantial **semester-long team project**
- Come ready to **discuss** – this is a high-interactivity room

Course Policies

Grading

Assessment	Points
Seminar participation	20
Methods hands-on activities	30
Project milestones	18
Teamwork (peer evaluation)	6
Final paper	8
Final presentation	8
Attendance & participation	10
<i>Total</i>	<i>100</i>

Attendance & participation (10 pts)

- **Attendance (5)** – seat check against a fixed pod chart; **2 free absences**, no questions asked
- **Pod role (3)** – three roles rotate at your table each Monday:
Reporter · Skeptic · Connector
- **Self-assessment (2)** – two half-page reflections (Oct 26, Dec 7)
making your case **with evidence**

The crit model

You'll critique each other's work at the **method-design** and **analysis** milestones – self-run, in your pod.

- The **first crit is modeled live** so you see how it should go
- Etiquette: **specific / actionable / kind** – critique the work, not the team
- Your pod runs its own crit, self-run on a timer I call from the front:
4 min present · **2 min** silent written feedback · **5 min** round-robin
· **2 min** respond & commit

AI-Use Policy

Use AI to do your best work – not to cut corners.

- You are **expected** to use it. No disclosure notes on coursework.
- What matters isn't which tool you touched – it's whether the **intellectual contribution** is yours.

Mind the jagged frontier

AI is excellent at some tasks and **confidently wrong** at neighboring ones – and the boundary is invisible, and it moves.

- **+30%** quality on tasks *inside* the frontier
- **–19 points** correctness on a task *outside* it


Learning where that edge falls for **your** work is part of what you're here to develop.

^{ai1} Dell'Acqua, F., McFowland, E., Mollick, E., Lifshitz, H., Kellogg, K.C., Rajendran, S., Kraye, L., Candelon, F., & Lakhani, K.R. (2026). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality. *Organization Science*, 37(2), 403–423. doi:10.1287/orsc.2025.21838

THE JAGGED FRONTIER

The jagged frontier works as follows: on one end of the spectrum, we witness AI's remarkable prowess — tasks that once seemed insurmountable for machines are now executed with precision and efficiency (from generating creative content to predicting complex patterns), with AI showcasing capabilities that were once the exclusive domain of human cognition.

Yet, on the flip side, there are tasks where AI falters, struggling to match human intuition and adaptability. These are areas marked by nuance, context, and intricate decision-making — realms where the binary logic of machines (currently) meets its match.



The two hard lines

You own what you submit

Every claim, number, and citation is yours to stand behind.

Fabricated data, results, or citations are academic misconduct – regardless of what produced them.

Venue rules override mine

Journals and conferences set their own requirements, and they change. **Our final paper follows ACM's** – generative-AI use disclosed in a brief Acknowledgements note.


The Project

The project, end to end

A term-long, **3-person** research project – the place the seminar and methods strands meet.

- **Topic → RQ / lit review → method design → data → analysis → paper + lightning talk**
- Incrementally written up as a full **ACM CHI-format paper**
- Mixes qualitative *and* quantitative methods

Six milestones, two crits



Seeding the project (*today only*)

We're **priming**, not forming.

- Take home the **speculative-design topic cards** as a thinking aid
- Let ideas marinate over the week – no topics, no teams yet
- Everything real happens **next Friday (Sep 11)**: topics go on the board – you recruit into them over the weekend

Assigned readings (*for Fri Sep 11 / M1*)

- **Lazar, Feng & Hochheiser (2017), Ch. 1** – Introduction to HCI research
- **Grudin (2012)** – *A Moving Target: The Evolution of HCI* (survey depth)
- **Wobbrock & Kientz (2016)** – *Research Contributions in HCI* – the seven types

What's Next?

- **Fri Sep 11 – M1: What is HCI research? + Project kickoff**
 - Read **Lazar et al. Ch. 1**
 - Bring your **topic cards**
 - Come ready to **put a topic on the board** – teams form over the weekend, final Mon Sep 14