Human-AI Interaction, Lecture 20

AI as Normal Technology

This is the sheet from **Lecture 1**.

Tool: _____ **Task:** _____

Gulf of execution

1. What did you want?

2. What did you type or do?

3. Where was it hard to get it to do that?

Gulf of evaluation

1. What did it give back?

2. How did you check it?

3. Where was it hard to tell what it did?

Circle the place the gulf was widest. **One change** that would make it smaller:

How do we design AI that people can **understand**, **steer**, and **delegate** to?

Today

1. Where is AI going? Two readings
2. Discussion of the readings
3. Three parts, and what you built
4. Redo your map
5. Your final paper
6. The two votes
7. What comes next

1. Where is AI going?

“[…] a separate species, a highly autonomous, potentially superintelligent entity”

Narayanan & Kapoor 2025, opening paragraph

“But this discourse about large models as intelligent agents is fundamentally **misconceived**.”

Farrell, Gopnik, Shalizi & Evans, Science 2025, p. 1153

ESSAYS AND SCHOLARSHIP

AI as Normal Technology

An alternative to the vision of AI as a potential superintelligence

BY ARVIND NARAYANAN & SAYASH KAPOOR
APRIL 15, 2025

Part I: “The Speed of Progress”

We articulate a vision of artificial intelligence (AI) as normal technology. To view AI as normal is not to understate its impact—even transformative, general-purpose technologies such as electricity and the internet are “normal” in our conception.

Narayanan & Kapoor, “AI as Normal Technology,” Knight First Amendment Institute, 2025

In their sense, a normal technology
can still change a lot, as electricity did.

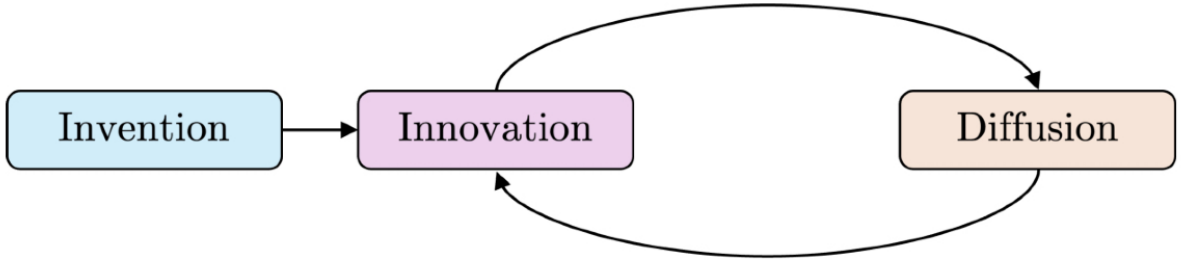
Description, prediction, prescription

> The statement “AI is normal technology” is three things: a description of current AI, a prediction about the foreseeable future of AI, and a prescription about how we should treat it. We view AI as a tool that we can and should remain in control of, and we argue that this goal does not require drastic policy interventions or technical breakthroughs.

Narayanan & Kapoor 2025

Their prescription, that people stay **in control** of AI, is also a design goal in this course.

Invention, innovation and diffusion



Stage	Methods / capabilities	Products / Applications	Early adoption	Adaptation
Example	LLMs	AI code editor	Individuals creating simple apps using AI	Retraining software engineers to emphasize AI skills
Appropriate metrics	Benchmarks	Uplift studies	Surveys	Economic indicators
Speed limits	Herding "Last mile" problem	Capability-reliability gap Dependence on diffusion	Learning curves Safety	Organizational changes Laws and norms

Narayanan & Kapoor, "AI as Normal Technology," 2025, Figure 1

Epic's sepsis prediction tool

“[…] despite having seemingly high accuracy when internally validated, performed far worse in hospitals, **missing two thirds of sepsis cases**”

Narayanan & Kapoor 2025, Part I, citing Wong et al., JAMA Internal Medicine 2021

Electricity in factories

As an example, Paul A. David's analysis of electrification shows that the productivity benefits took decades to fully materialize. Electric dynamos were "everywhere but in the productivity statistics" for nearly 40 years after Edison's first central generating station.

What eventually allowed gains to be realized was redesigning the entire layout of factories around the logic of production lines.

Narayanan & Kapoor 2025, Part I, citing Paul A. David 1990

The same can happen with an AI tool: it may not help much until the work is **redesigned** around it.

Benchmarks and real use

“For example, while GPT-4 reportedly achieved scores in the top 10% of bar exam test takers, this tells us **remarkably little** about AI's ability to practice law.”

Narayanan & Kapoor 2025, Part I, “Benchmarks do not measure real-world utility”

How to measure usefulness

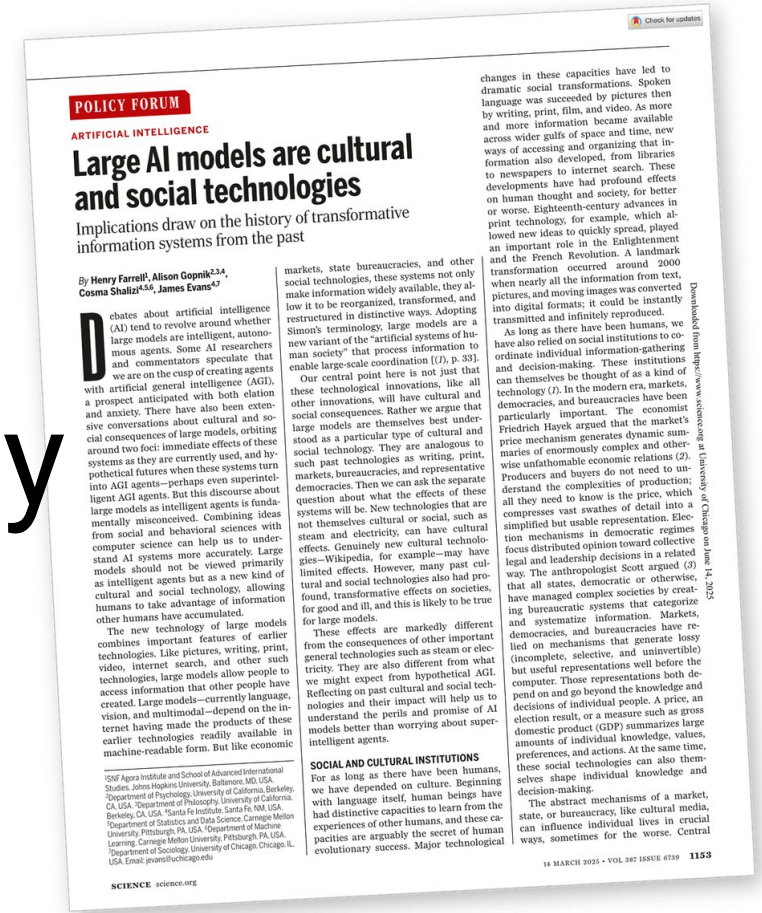
> The only sure way to measure real-world usefulness of a potential application is to actually build the application and to then test it with professionals in realistic scenarios (either substituting or augmenting their labor, depending on the intended use). Such 'uplift' studies generally do show that professionals in many occupations benefit from existing AI systems, but this benefit is typically modest and is more about augmentation than substitution [...]

Narayanan & Kapoor 2025, Part I

That is what milestones 2, 3 and 4 ask you to do.

A cultural and social technology

Reading 2: Farrell et al., *Science*, 2025



POLICY FORUM

ARTIFICIAL INTELLIGENCE

Large AI models are cultural and social technologies

Implications draw on the history of transformative information systems from the past

By Henry Farrell[1], Alison Gopnik[2,3,4], Cosma Shalizi[4,5,6], James Evans[4,7]

Debates about artificial intelligence (AI) tend to revolve around whether large models are intelligent, autonomous agents. Some AI researchers and commentators speculate that we are on the cusp of creating agents with artificial general intelligence (AGI), a prospect anticipated with both elation and anxiety. There have also been extensive conversations about cultural and social consequences of large models, orbiting around two foci: immediate effects of these systems as they are currently used, and hypothetical futures when these systems turn into AGI agents—perhaps even superintelligent AGI agents. But this discourse about large models as intelligent agents is fundamentally misconceived. Combining ideas from social and behavioral sciences with computer science can help us to understand AI systems more accurately. Large models should not be viewed primarily as intelligent agents but as a new kind of cultural and social technology, allowing humans to take advantage of information other humans have accumulated.

The new technology of large models combines important features of earlier technologies. Like pictures, writing, print, video, internet search, and other such technologies, large models allow people to access information that other people have created. Large models—currently language, vision, and multimodal—depend on the internet having made the products of these earlier technologies readily available in machine-readable form. But like economic markets, state bureaucracies, and other social technologies, these systems not only make information widely available, they allow it to be reorganized, transformed, and restructured in distinctive ways. Adopting Simon's terminology, large models are a new variant of the "artificial systems of human society" that process information to enable large-scale coordination [(*1*), p. 33].

Our central point here is not just that these technological innovations, like all other innovations, will have cultural and social consequences. Rather we argue that large models are themselves best understood as a particular type of cultural and social technology. They are analogous to such past technologies as writing, print, markets, bureaucracies, and representative democracies. Then we can ask the separate question about what the effects of these systems will be. New technologies that are not themselves cultural or social, such as steam and electricity, can have cultural effects. Genuinely new cultural technologies—Wikipedia, for example—may have limited effects. However, many past cultural and social technologies also had profound, transformative effects on societies, for good and ill, and this is likely to be true for large models.

These effects are markedly different from the consequences of other important general technologies such as steam or electricity. They are also different from what we might expect from hypothetical AGI. Reflecting on past cultural and social technologies and their impact will help us to understand the perils and promise of AI models better than worrying about superintelligent agents.

SOCIAL AND CULTURAL INSTITUTIONS

For as long as there have been humans, we have depended on culture. Beginning with language itself, human beings have had distinctive capacities to learn from the experiences of other humans, and these capacities are arguably the secret of human evolutionary success. Major technological changes in these capacities have led to dramatic social transformations. Spoken language was succeeded by pictures then by writing, print, film, and video. As more and more information became available across wider gulfs of space and time, new ways of accessing and organizing that information also developed, from libraries to newspapers to internet search. These developments have had profound effects on human thought and society, for better or worse. Eighteenth-century advances in print technology, for example, which allowed new ideas to quickly spread, played an important role in the Enlightenment and the French Revolution. A landmark transformation occurred around 2000 when nearly all the information from text, pictures, and moving images was converted into digital formats; it could be instantly transmitted and infinitely reproduced.

As long as there have been humans, we have also relied on social institutions to coordinate individual information-gathering and decision-making. These institutions can themselves be thought of as a kind of technology (*1*). In the modern era, markets, democracies, and bureaucracies have been particularly important. The economist Friedrich Hayek argued that the market's price mechanism generates dynamic summaries of enormously complex and otherwise unfathomable economic relations (*2*). Producers and buyers do not need to understand the complexities of production; all they need to know is the price, which compresses vast swathes of detail into a simplified but usable representation. Election mechanisms in democratic regimes focus distributed opinion toward collective legal and leadership decisions in a related way. The anthropologist Scott argued (*3*) that all states, democratic or otherwise, have managed complex societies by creating bureaucratic systems that categorize and systematize information. Markets, democracies, and bureaucracies have relied on mechanisms that generate lossy (incomplete, selective, and uninvertible) but useful representations well before the computer. Those representations both depend on and go beyond the knowledge and decisions of individual people. A price, an election result, or a measure such as gross domestic product (GDP) summarizes large amounts of individual knowledge, values, preferences, and actions. At the same time, these social technologies can also themselves shape individual knowledge and decision-making.

The abstract mechanisms of a market, state, or bureaucracy, like cultural media, can influence individual lives in crucial ways, sometimes for the worse. Central

[1]SNF Agora Institute and School of Advanced International Studies, Johns Hopkins University, Baltimore, MD, USA. [2]Department of Psychology, University of California, Berkeley, CA, USA. [3]Department of Philosophy, University of California, Berkeley, CA, USA. [4]Santa Fe Institute, Santa Fe, NM, USA. [5]Department of Statistics and Data Science, Carnegie Mellon University, Pittsburgh, PA, USA. [6]Department of Machine Learning, Carnegie Mellon University, Pittsburgh, PA, USA. [7]Department of Sociology, University of Chicago, Chicago, IL, USA. Email: jevans@uchicago.edu

Check for updates

Downloaded from https://www.science.org at University of Chicago on June 14, 2025

Farrell et al.'s claim

stand AI systems more accurately. Large models should not be viewed primarily as intelligent agents but as a new kind of cultural and social technology, allowing humans to take advantage of information other humans have accumulated.

Farrell et al., Science 2025, p. 1153

Earlier cultural and social technologies: writing, print, markets, bureaucracies, and representative democracies.

Lossy summaries

“A price, an election result, or a measure such as gross domestic product (GDP) summarizes large amounts of individual knowledge, values, preferences, and actions.”

Farrell et al., Science 2025, p. 1153

tions” (*1*) of this information at scale and in new ways, like markets, states, and bureaucracies. Just as market prices are lossy representations of the underlying allocations and uses of resources, and government statistics and bureaucratic categories imperfectly represent the characteristics of underlying populations, so too are large models “lossy JPEGs” (*6*) of the data corpora on which they have been trained.

Farrell et al., Science 2025, p. 1154

Other people's work in an output

active users. Someone asking a bot for help writing a cover letter for a job application is really engaging in a technically mediated relationship with thousands of earlier job applicants and millions of other letter writers and RLHF workers.

Farrell et al., Science 2025, p. 1155

RLHF is the learning from human preferences in Lecture 4.

Hallucinations and uneven accuracy

“‘hallucinations’ are an endemic problem in these systems because they have **no conception of truth and falsity**”

Farrell et al., Science 2025, p. 1154; see Lecture 8

“an intrinsic tendency to be **most accurate in situations most commonly found in their training data**”

Farrell et al., Science 2025, p. 1155; see Lectures 11 and 14

“Can large models be developed to **enhance** human creativity rather than to **dull** it?”

Farrell et al., Science 2025, p. 1156

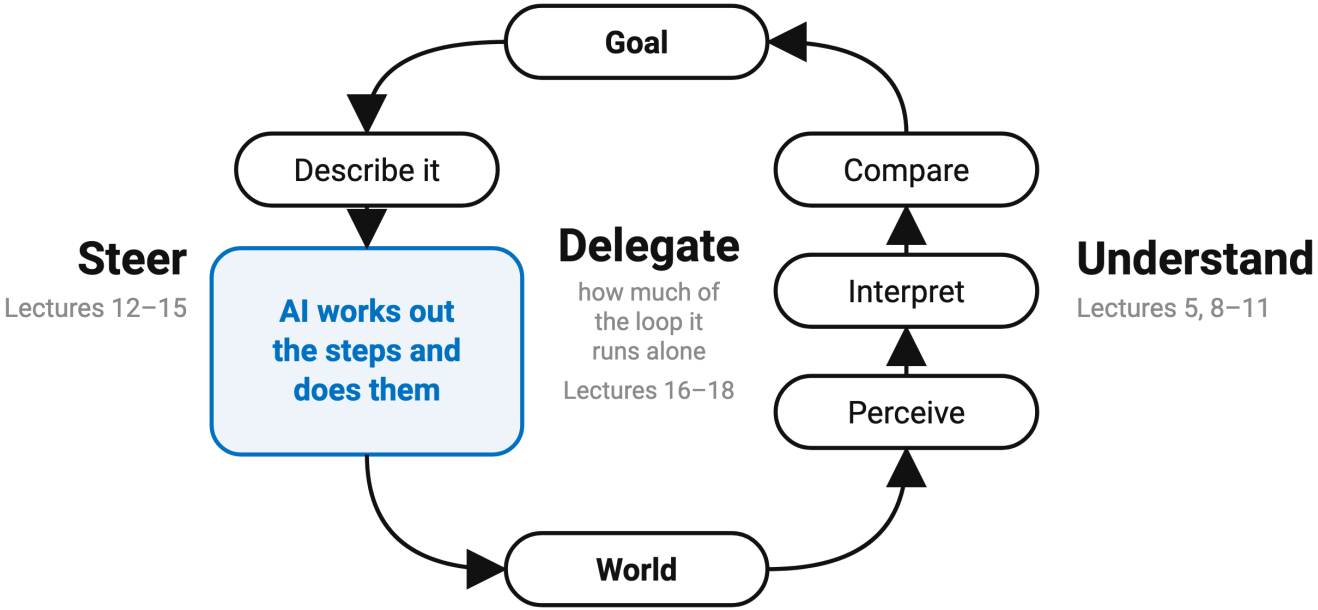
Two views of AI

	AI as normal technology	**AI as a cultural and social technology**
Like	electricity, the internet	writing, print, markets
For a designer	**Who has to change how they work for this to help?**	**Whose work is in the output?**

Discussion: the two readings

1. Which view fits what you saw in your own tests?
2. Pick one view. What would it make you **change** in your project?
3. What does each reading leave out that this course covered?

3. Three parts, and what you built



Understand: the gulf of evaluation



Eslami et al. 2015



Kim et al. 2024



Liao et al. 2020



Lee & See 2004



Buçinca et al. 2021



Buolamwini & Gebru 2018

Milestone 2: Make it understandable

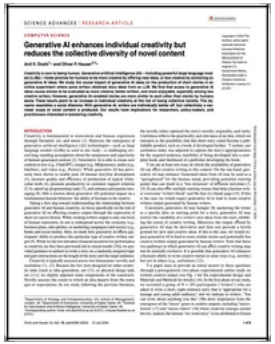
Steer: the gulf of execution


Zamfirescu-Pereira 2023


Wu et al. 2022


Shneiderman 1983


Chung & Adar 2023


Doshi & Hauser 2024


Horvitz 1999

Milestone 3: Make it steerable

Delegate: agents and automation



Maes 1994



Feng et al. 2025



Bainbridge 1983



Dell'Acqua et al. 2026



Weidinger et al. 2022



Samuelson 2023

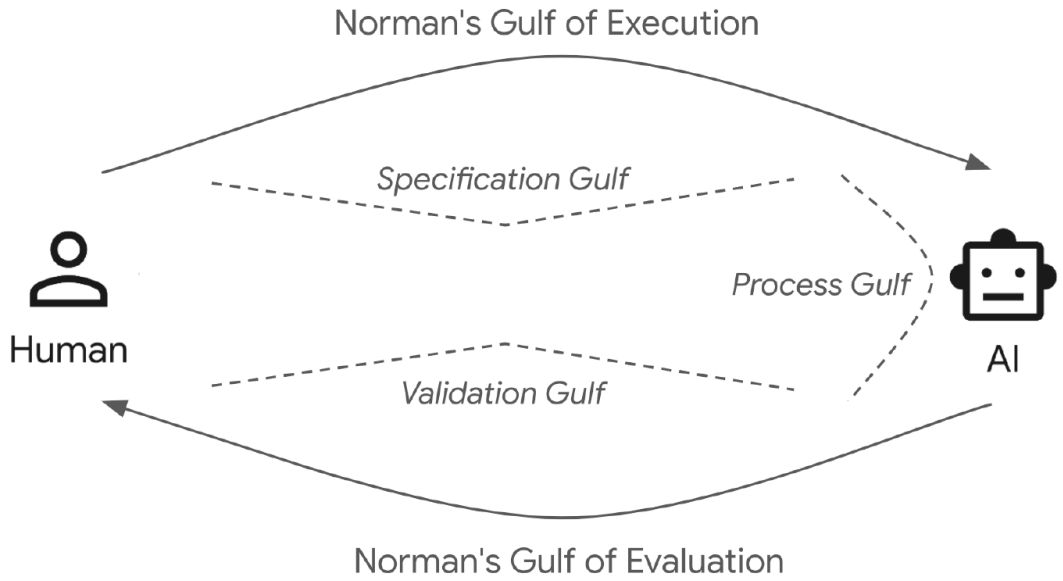
Milestone 4: Decide what it does on its own

Which part was **hardest** to design for in your project?

Redo your map: 15 minutes

1. Take out your sheet from Lecture 1
2. Same tool, same task: fill in a **new** sheet
3. Then put them side by side and circle what changed
4. Name the lecture or reading behind each change

Terry et al.: three gulfs for AI



Specification: telling the AI what to do

Process: how it does it

Validation: checking what it made

Terry et al., arXiv 2023 (v2, 2024), Fig. 2; optional reading

Final paper

Due two days after the fair

Paper outline: 8 minutes

1. The problem, and who it is for
2. Your tool, and the figure that shows it best
3. How you tested it, and what you found
4. Four readings, and what your tool adds
5. Where your design stops working

The vote in Lectures 2 and 16

Lecture 2, week 1

Lecture 16, week 8

A. Direct manipulation

B. Interface agents

C. It depends: **on what?**

How did the counts change?

Further courses and reading

Courses	HCI research methods, information visualization, machine learning, a design studio, technology ethics and policy
Papers	CHI, UIST, CSCW, DIS, Creativity & Cognition, FAccT
Research	write to a lab whose papers you liked this term

Next: finals week

- **Demo fair**: finals week
- **Final paper**: two days after the fair

Thank you for a good term. See you at the demo fair.