Ch 6-7 in “AI Ethics” (Week 11)

Welcome to today’s discussion. Ch 6 looks at machine learning and data science, and Ch 7 looks at the real ethical challenges that arise when these technologies are deployed in society. Use this guide to prepare for small-group and whole-class conversation.

Part I — Chapter 6: Don’t Forget the Data (Science)

Quick Recap

Machine learning (ML) is the collection of algorithms that try to identify patterns in data without being told explicit rules. Data science is the broader practice of collecting, cleaning, and interpreting data. Both are built on statistics. Crucially, humans are involved at every stage—from deciding what data to collect to interpreting what the results mean.

Check Your Understanding

Before discussing, make sure you can answer these in your own words:

1. What is the difference between supervised and unsupervised machine learning?
2. Why does the author say data is now ‘active’ rather than ‘passive’?
3. What is a ‘spurious correlation’ and why does it matter for AI?
4. The author compares a statistical model to a map. What does this analogy mean?

Discussion Questions — Chapter 6

Q1. The Human Element

The author argues that ‘without programmers and data scientists, the technology simply doesn’t work’ and that humans are involved at every stage of the ML pipeline.

Does knowing this change how you think about who is responsible when an AI system makes a harmful decision? Why or why not? Can responsibility be shared between a human team and an algorithm?

Q2. Correlation vs. Causation

The book cites a correlation between the divorce rate in Maine and per-capita consumption of margarine—clearly not a causal relationship.

Can you think of a real-world AI application where acting on a spurious correlation could cause serious harm? What safeguards, if any, would prevent this? Who should bear responsibility for checking whether AI-found patterns are meaningful?

Q3. Choices in the Data

Designing a training data set always involves choices about ‘how to abstract from reality.’ Those choices are never neutral.

Think about an AI used in college admissions, hiring, or medical diagnosis. What choices about data collection might embed unfairness into the system before a single prediction is made? Who gets to make those choices, and who should?

Q4. Everyday ML

The chapter lists dozens of ML applications: Netflix recommendations, credit scoring, face recognition in stores, cancer diagnosis, Hello Barbie.

Which application on this list do you find most surprising or troubling, and why? Is there a meaningful ethical difference between ML used in entertainment versus in criminal justice or healthcare?

Part II — Chapter 7: Privacy and the Other Usual Suspects

Quick Recap

Chapter 7 surveys the major ethical problems raised by AI: privacy violations and data repurposing; manipulation and exploitation of users (including vulnerable ones); the spread of disinformation and erosion of trust; safety risks from physical AI systems; and security risks when AI is hacked or weaponized. A unifying theme is that these are not new problems, but AI makes them more intense and harder to see.

Check Your Understanding

Before discussing, make sure you can answer these in your own words:

5. What does ‘data repurposing’ mean and why is it ethically significant?
6. What did Cambridge Analytica do, and why did the author use it as an example?
7. Who are ‘vulnerable users’ in this context and why do they matter?
8. How does the author’s closing argument connect technological vulnerability to human existence?

Discussion Questions — Chapter 7

Q5. Consent and Control

The author notes that social media users face a coercive choice: consent to data collection or lose access to the platform. The text also notes that users often have no idea what data are collected or how they are used.

Is this kind of consent meaningful? What would ‘genuine’ informed consent look like in the context of a platform like Instagram or TikTok? Is it even achievable, given how complex these systems are?

Q6. Exploitation, Visible and Hidden

The chapter distinguishes two kinds of exploitation: users who produce free ‘digital labor’ (data), and workers who mine minerals, handle e-waste, and label data sets that make AI function.

Which of these forms of exploitation concerns you more, and why? Are they connected? What, if anything, do we as technology users owe to the invisible workforce behind our devices?

Q7. Vulnerable Users & the Internet of Toys

Hello Barbie records and analyzes everything a child says, sending responses based on what it has ‘learned’ about the child. The child likely has no idea.

Who bears responsibility here—the manufacturer, the parents, regulators, or society at large? Should children’s data be treated differently from adults’ data? Where would you draw the line?

Q8. Truth, Trust & Disinformation

The author argues that AI-powered disinformation risks a world in which ‘it is no longer clear what is true and what is false’—not through outright authoritarianism but through confusion and manipulation.

Have you personally experienced difficulty distinguishing AI-generated content from human-created content? What, if any, responsibility do platforms have to label AI-generated material? Does this responsibility differ between political and entertainment content?

Q9. Safety, Security & Dependency

The author’s closing claim: ‘To the extent that we become dependent on AI, AI is more than a tool we use; it becomes part of how we are, and how we are at risk, in the world.’

Do you agree that dependence on AI changes our existence as humans, or is this overstated? Think about a specific AI system you rely on. What would happen—to you, or to society—if it were suddenly unavailable or compromised?

Synthesis Question (In-class writing assignment)

Chapter 6 insists that humans are always involved in ML. Chapter 7 argues that AI causes serious harms. If humans are always in the loop, why do these harms keep occurring? What does your answer suggest about where reform should be focused—technical design, corporate policy, government regulation, or individual behavior? (Try to incorporate ideas from today's discussions)