


Conversation Pieces



MIT
Perspectives
on Ethics,
Computing,
and AI



Conversation Pieces

MIT Perspectives on Ethics, Computing, and AI

Excerpts from a series by the Office of the Dean
MIT School of Humanities, Arts, and Social Sciences

Prepared for Hello World, Hello MIT.
**A Celebration of the MIT Stephen A. Schwarzman
College of Computing**
February 26–28, 2019

The quotes in this booklet are just a taste of the full
commentaries, which you will find online at:
shass.mit.edu/aiethics.

TABLE OF CONTENTS

FOREWORD | Ethics, Computing, and AI

Melissa Nobles

Kenan Sahin Dean and Professor of Political Science
School of Humanities, Arts, and Social Sciences

INTRODUCTION | The Tools of Moral Philosophy

Caspar Hare and Kieran Setiya

Professors of Philosophy
School of Humanities, Arts, and Social Sciences

I. A Human Endeavor

Computing is embedded in cultural, economic, and political realities.

Computing is Deeply Human

Stefan Helmreich, Elting E. Morison Professor of Anthropology
Heather Paxson, William R. Kenan Jr. Professor of Anthropology
School of Humanities, Arts, and Social Sciences

When Programs Become Unpredictable

John Guttag, Dugald C. Jackson Professor of Computer Science
and Electrical Engineering
School of Engineering

Safeguarding Humanity in the Age of AI

Bernhardt Trout, Raymond F. Baddour Professor of Chemical Engineering
School of Engineering

II. Community Insights

Shaping ethical technology is a collective responsibility.

The Common Ground of Stories

Mary Fuller, Professor of Literature
School of Humanities, Arts, and Social Sciences

Who's Calling the Shots?

Leigh Hafrey, Senior Lecturer, Leadership and Ethics
MIT Sloan School of Management

In Praise of Wetware

Caroline A. Jones, Professor of Art History
School of Architecture and Planning

Blind Spots

David Kaiser, Germeshausen Professor of the History of Science,
and Professor of Physics
School of Humanities, Arts, and Social Sciences + School of Science

Assessing the Impact of AI on Society

Lisa Parks, Professor of Comparative Media Studies
School of Humanities, Arts, and Social Sciences

Clues and Caution from Biomedicine

Robin Wolfe Scheffler, Assistant Professor of the History of Science
School of Humanities, Arts, and Social Sciences

The Necessary Environment for Ethical Action

T.L. Taylor, Professor of Comparative Media Studies
School of Humanities, Arts, and Social Sciences

Biological Intelligence and AI

Matthew A. Wilson, Sherman Fairchild Professor of Neuroscience
School of Science + The Picower Institute

Machine Anxiety

Bernardo Zacka, Assistant Professor of Political Science
School of Humanities, Arts, and Social Sciences



III. A Structure for Collaboration

Thinking together is powerful.

Blending and Bilinguals

Hal Abelson, Class of 1922 Professor of Electrical Engineering
and Computer Science
School of Engineering

A Network of Practitioners

Nick Montfort, Professor of Digital Media
School of Humanities, Arts, and Social Sciences

A Dream of Computing

D. Fox Harrell, Professor of Digital Media and Artificial Intelligence
School of Humanities, Arts, and Social Sciences + Computer Science
and Artificial Intelligence Lab

Ethics, Computing, and AI

Melissa Nobles


**Kenan Sahin Dean
and Professor of Political Science
School of Humanities, Arts, and Social Sciences**

What kind of world do we want to make? Will the future be humane and livable? Will it be fair and just? What knowledge and values can sustain and guide us? The launch of the MIT Stephen A. Schwarzman College of Computing offers an extraordinary new opportunity for the MIT community to respond to today's most consequential questions in ways that serve the common good.

To support ongoing conversations about how to do this, my office invited faculty from all five MIT schools to offer perspectives on the role of ethics in computing and AI. This booklet of short quotes from their responses offers a taste of the full commentaries online (shass.mit.edu/aiethics.) These insights—drawn from expertise in fields as diverse as moral philosophy, neuroscience, history, engineering, and literature—also sound some common themes:

that the practice of computing is a human endeavor, one that—like all human endeavors—can be used for good or ill; that the humanistic disciplines are deep resources for research and education about societal and ethical issues; that we envision an MIT culture in which all of us are equipped and encouraged to think about the social and ethical implications of new research and new technologies.

These commentaries also implore us to be collaborative, foresighted, and courageous as we shape a new College, and to proceed with judicious humility. Rightly so. We are embarking on an endeavor that will influence nearly all aspects of the human future.




The Tools of Moral Philosophy

Caspar Hare and Kieran Setiya

**Professors of Philosophy
School of Humanities, Arts, and Social Sciences**

“We face ethical questions every day. Philosophy does not provide easy answers for these questions, nor even fail-safe techniques for resolving them. What it does provide is a disciplined way to think about ethical questions, to identify hidden moral assumptions, and to establish principles by which our actions may be guided and judged. Framing a discussion of the risks of advanced technology entirely in terms of ethics suggests that the problems raised are ones that can and should be solved by individual action. In fact, many of the challenges presented by computer science will prove difficult to address without systemic change. One responsibility that does fall on us as individuals is to work toward political conditions in which it is possible for us to live and work more ethically.”






I.

A Human Endeavor

Computing is embedded
in cultural, economic,
and political realities.



Computing is Deeply Human

Stefan Helmreich

**Elting E. Morison Professor of Anthropology
School of Humanities, Arts, and Social Sciences**

Heather Paxson

**William R. Kenan Jr. Professor of Anthropology
School of Humanities, Arts, and Social Sciences**


“Computing is a human practice that entails judgment and is embedded in politics. Computing is not an external force that has an impact on society; instead, society—institutional structures that organize systems of social norms—*is built right into* making, programming, and using computers. Therefore, as we found the MIT Schwarzman College of Computing, we need to think about promoting values that enable the creative flourishing of a greater diversity of programmers, who may in turn help make clear the many ways in which the computational is political.”

When Programs Become Unpredictable

John Guttag

**Dugald C. Jackson Professor of Computer Science and Electrical Engineering
School of Engineering**

“What happens when machine learning is used to build programs that make recommendations with important societal implications? Will they sometimes make bad recommendations? Almost certainly. So, while we should look forward to the many good things machine learning will bring to society, we should insist that technologists study the risks and explain them clearly. And society as a whole should take responsibility for understanding the risks and for making human-centric choices about using this ever-evolving technology.”





Safeguarding Humanity in the Age of AI








Bernhardt Trout


**Raymond F. Baddour Professor of
Chemical Engineering
School of Engineering**

“There seem to be two possibilities for how AI will turn out. In the first, AI will do what it is on track to do: slowly take over every human discipline. The second possibility is that we take the existential threat of AI with the utmost seriousness and completely change our approach. This means redirecting our thinking from a blind belief in efficiency to a considered understanding of what is most important about human life.”










II. **Community Insights** Shaping ethical outcomes is a collective responsibility.



The Common Ground of Stories

Mary Fuller

Professor of Literature

School of Humanities, Arts, and Social Sciences


“Stories are things in themselves, and they are also things to think with. Stories allow us to model interpretive, affective, ethical choices; they also become common ground. Reading about Milton’s angelic intelligences or William Gibson’s ‘bright lattices of logic’ won’t tell us what we should do with the future, but reading such stories at MIT may offer a conceptual meeting place to think together across the diversity of what and how we know.”

Who's Calling the Shots?

Leigh Hafrey

**Senior Lecturer, Leadership and Ethics
MIT Sloan School of Management**

“Efficiency’ is a perennial business value and a constant factor in corporate design, strategy, and execution. But in a world where the exercise of social control by larger entities is real, developments in artificial intelligence have yet to yield the ethics by which we might manage their effects. The integrity of our vision for the future depends on our learning from the past and celebrating the fact that people, not artifacts and institutions, set our rules of engagement.”





In Praise of Wetware

Caroline A. Jones

**Professor of Art History
School of Architecture and Planning**

“Humans have never been simple. Unlike the machines we build, humans are full of wet chemical signals, pulsing with fluids. Let’s face it, computational models of cognition have never been fully adequate to the wetware within, or the biological environment without. The problems we have brought to our planet require that every discipline contribute, harmonize, concertize, critique, and collaborate in bringing wetness into the equation.”





Blind Spots

David Kaiser

**Germeshausen Professor of the History of Science, and Professor of Physics
School of Humanities, Arts, and Social Sciences +
School of Science**


“No group of researchers, flushed with the excitement of building something new, can overcome the limitations of blind spots and momentum alone. How can we build an infrastructure that can address these limitations? One approach would be to institute forums in which people with many different backgrounds, fields, concerns, and experiences can brainstorm and debate together. When it comes to technologies as potent and far-reaching as AI, we will benefit greatly from engaging as many stakeholders as possible.”


Assessing the Impact of AI on Society

Lisa Parks

**Professor of Comparative Media Studies
School of Humanities, Arts, and Social Sciences**

“Given the power of AI tools to impact human behavior and shape planetary conditions, it is vital to conduct a political, economic, and materialist analysis of the technology’s relation to global trade, governance, natural environments, and culture. This will involve adopting an infrastructural disposition and specifying AI’s constitutive parts, processes, and effects as they take shape across diverse world contexts. Only then can the public understand the technology well enough to democratically deliberate its relation to ethics and policy.”





Clues and Caution from Biomedicine

Robin Wolfe Scheffler

**Assistant Professor of the History of Science
School of Humanities, Arts, and Social Sciences**

“The application of AI to the biomedical sciences follows upon an earlier series of efforts to make the fields of biology and medicine quantifiable and computable. The use of AI in these fields today deepens rather than resolves longstanding questions that are raised by the past intractability of biology and medicine to computation, and by the flawed assumptions that were adopted in attempting to make them so. The history of these efforts underlines two major points: ‘quantification is a *process of judgment and evaluation*, not simple measurement’ and ‘prediction is not destiny.’”




The Necessary Environment for Ethical Action

T.L. Taylor

**Professor of Comparative Media Studies
School of Humanities, Arts, and Social Sciences**

“Many technological developments—from facial recognition to algorithmic decision systems—are causing real harm, good intentions notwithstanding. These are not trivial matters; they impact everything from the justice system to medical care. Unfortunately, there are limits to individualistic models of ethics. We can cultivate our students as ethical thinkers but if they aren’t working in (or studying in) structures that support advocacy and intervention, they will be stymied. We must, as both a university and a high-profile stakeholder in the overall ecology of technological development, think sociologically about critical ethical issues.”



Biological Intelligence and AI

Matthew A. Wilson

**Sherman Fairchild Professor of Neuroscience
School of Science + The Picower Institute**


“An understanding of biological intelligence is relevant to the development of AI, and the effort to develop artificial general intelligence (AGI) magnifies its significance. AGIs will be expected to conform to standards of behavior in situations that may not have been initially considered. This is the standard by which we evaluate biological intelligence: adaptive and appropriate behavior in novel contexts. Should we hold AIs to the same standards as the average human? Or will we expect AIs to perform at the level of an ideal human?”


Machine Anxiety


Bernardo Zacka

**Assistant Professor of Political Science
School of Humanities, Arts, and Social Sciences**

“To someone who studies bureaucracy, the anxieties surrounding AI have an eerily familiar ring. So too does the excitement. For much of the 20th century, bureaucracies were thought to be intelligent machines, with all the positive and negative connotations the term carries. As we examine the ethical and political implications of AI, the comparison with bureaucracy—the devil we know—is instructive. There are at least two insights to draw from bureaucracy’s history: that it is worth studying our anxieties whether or not they are realistic; and that in doing so we should not write off human agency.”









III.

A Structure for Collaboration

Thinking together is
powerful.





Blending and Bilinguals

Hal Abelson

**Class of 1922 Professor of Electrical
Engineering and Computer Science
School of Engineering**


“When we study society today, we can no longer separate humanities—the study of what’s human—from computing. So, while there is discussion underway about building bridges between computing and the humanities, arts, and social sciences, what the MIT Schwarzman College of Computing needs is blending, not bridges. MIT’s guideline should be President L. Rafael Reif’s goal to ‘educate the bilinguals of the future,’ experts in many fields who are also skilled in modern computing. Can MIT achieve such an aspiration? To do so, we’ll need joint research and joint teaching.”


A Network of Practitioners

Nick Montfort

**Professor of Digital Media
School of Humanities, Arts, and Social Sciences**

“Computing is not a single discipline nor even a set of disciplines; it is a practice. To reflect that, we could reconsider the initial metaphor for the College structure: that is, ‘bridge’ hires which connect a ‘core’ of computer science to other disciplines. A more inclusive metaphor would resonate with the fact that those working with computing today in, for example, the visual and literary arts, music, theater, and gaming often connect directly. Let’s use this chance to build a robust network, one that supports many types of connections.”






A Dream of Computing

D. Fox Harrell

**Professor of Digital Media and Artificial Intelligence
School of Humanities, Arts, and Social Sciences +
Computer Science and Artificial Intelligence Lab**

“We must reimagine our shared dreams for computing technologies as ones where their potential social and cultural impacts are considered intrinsic to the engineering practices of inventing them. To support such work within the new College, we need a model for its endeavors. Metaphorically, one could think of the College as a rhizome, with roots representing work in foundational disciplines and shoots representing work in emerging disciplines and transdisciplinary practices. Alternatively, one might think of the new College as an aquifer providing transformative outcomes (water) if tapped into with the artesian wells of research, practice, and development.”



Conversation Pieces: MIT Perspectives on Ethics, Computing, and AI is a collection of short excerpts from a series developed by the Office of the Dean, MIT School of Humanities, Arts, and Social Sciences.

Full publication online: shass.mit.edu/aiethics

© 2019 MIT School of Humanities, Arts, and Social Sciences

ADDITIONAL MIT PUBLICATIONS

The MIT Stephen A Schwarzman College of Computing

Article by Susan Silbey, Chair of the MIT Faculty, in the MIT Faculty Newsletter.

“The task before us now is to develop and propose designs for the MIT Schwarzman College of Computing.”

shass.mit.edu/aiethics/silbey

Ethical AI by Design


Interview with Abby Everett Jaques, MIT Philosophy postdoctoral associate.

“In this hands-on course, we teach an ethics protocol, a step-by-step process that students can apply to projects of their own.”

shass.mit.edu/aiethics/jaques



SCHOOL OF
HUMANITIES
ARTS AND
SOCIAL SCIENCES





Office of the Dean
MIT School of Humanities,
Arts, and Social Sciences

MIT