

OPEN QUESTIONS

IN THE AGE OF A.I., WHAT MAKES PEOPLE UNIQUE?

*More than ever, we're challenged to define what's valuable
about being human.*

By Joshua Rothman

August 6, 2024







Illustration by Josie Norton



Recently, I got a haircut, and my barber and I started talking about A.I. “It’s incredible,” he told me. “I just used it to write a poem for my girl’s birthday. I told it what to say, but I can’t rhyme that well, so it did all the writing. When she read the poem, she actually cried! Then she showed it to her friend, who’s really smart, and I thought, Uh-oh, she’ll figure it out for sure.” Snip, snip, snip, snip. “She didn’t.”

Everyone in the barbershop laughed, a little darkly. Writing poems that make your girl cry—add that to the list of abilities that used to make (some) humans unique but no longer do. Today’s A.I. systems can generate acceptable poetry, code, essays, and jokes; carry on useful conversations about economics, existentialism, and the Middle East; and even perform some aspects of scientific work, such as planning experiments, predicting outcomes, and interpreting results. They can make judgments about complex situations—traffic patterns,

investments—at superhuman speed. In truth, we don't yet know all they can do. The biggest tech companies are racing to deploy the technology partly so that we can find out.

It seems entirely likely that the list of A.I.'s capabilities will only grow—and so it's tempting to wonder what, exactly, people are good for. In the past, theologians and philosophers compared us with animals and identified the ways in which we surpassed them. Now the tables aren't so much turned as upended. In some cases, we seem to be looking upward at the machines (no human being can write with an A.I.'s fluidity and speed, for example). In others, we scratch our heads at their stupidity (no person would advise you to make a daily habit of eating “at least one small rock,” as Google's A.I. did not long ago, when asked “How many rocks should I eat each day?”). In still other cases, we're simply confused by the divergences between artificial and organic reasoning. An A.I. can't fall in love, but it can express the idea of love; it can't be an artist, but it can (maybe) create a kind of art; it can't agonize over a consequential decision, but it can still decide. We know that there are crucial differences between a thinking computer and a person, but defining those distinctions isn't easy.

And yet this abstract conundrum has practical implications. As artificial intelligence proliferates, more and more hinges on our ability to articulate our own value. We seem to be on the cusp of a world in which workers of all kinds—teachers, doctors, writers, photographers, lawyers, coders, clerks, and more—will be replaced with, or to some degree sidelined by, their A.I. equivalents. What will get left out when A.I. steps in?

In “A.I. Snake Oil: What Artificial Intelligence Can Do, What It Can’t, and How to Tell the Difference,” two computer scientists, Arvind Narayanan and Sayash Kapoor, approach the question on a practical level. They urge skepticism, and argue that the blanket term “A.I.” can serve as a kind of smoke screen for underperforming technologies. “Imagine an alternate universe in which people don’t have words for different forms of transportation—only the collective noun ‘vehicle,’ ” they write. Such a world sees “furious debates about whether or not vehicles are environmentally friendly, even though no one realizes that one side of the debate is talking about bikes and the other side is talking about trucks.” Similarly, they write, the term “A.I.” encompasses a variety of technologies with wildly different levels of competence.

Narayanan and Kapoor are particularly wary of predictive artificial intelligence, which is designed to make guesses about the future. Unlike generative A.I.—the relatively new technology used in ChatGPT and the like—predictive A.I. is already integrated into our lives to a surprising extent. Human-resources departments use it to suggest which candidates will succeed on the job; hospitals employ it to help decide who should be sent home or admitted for a stay. And yet predictive A.I. systems are almost never rigorously and independently tested; when they are, they often fail. Narayanan and Kapoor recount the findings of researchers investigating an A.I. system called Retorio, which claims to predict future on-the-job behavior, and thus performance, by analyzing video interviews with job candidates. It turned out that wearing glasses or a scarf, sitting in

front of some bookshelves, or sending a résumé in the form of a PDF could drastically change a candidate's score. Wearing glasses “obviously does not change someone's capability to perform well at a job,” the authors write. In their view, the system is A.I. snake oil.

The problems with predictive A.I. can run deeper than mere inaccuracy. In an early experiment, researchers built a system for guessing whether pneumonia patients arriving at a hospital would need overnight care. The system examined the data and discovered that patients with asthma tended to recover from pneumonia faster; this made it more likely to recommend that asthmatic patients be sent home. That's a crazy recommendation, of course; the correlation on which it's based reflects the fact that asthmatic people with pneumonia are often admitted directly to the I.C.U., where they receive high levels of care. (The system was never used.) “A good prediction is not a good decision,” Narayanan and Kapoor write. Among other things, being a capable decision-maker means not just interrogating the origins of your intuitions, but also imagining how your upcoming decisions might render those intuitions invalid. It's highly unlikely that candidates who Zoom while sitting in front of bookshelves will be better employees—but, even if that prediction were true, acting on it repeatedly would simply teach interviewees to sit in front of bookshelves. As human beings, we have a sense of the fallibility of our thinking; it's one of our strengths.

Shannon Vallor, a philosopher at the University of Edinburgh who has worked as an A.I. ethicist at Google, doesn't enumerate the failures of A.I. so much as explore the range and potency of human

virtues. In “The A.I. Mirror: How to Reclaim Our Humanity in an Age of Machine Thinking,” she argues that we vastly underestimate our own richness compared with that of A.I. “Consider the image that appears in your bathroom mirror every morning,” she writes. It isn’t a copy of your body, or an imitation of it; it’s just a reflection. Similarly, today’s A.I. systems “don’t produce thoughts or feelings any more than mirrors produce bodies. What they produce is a new kind of reflection.”

Vallor specializes not just in the philosophy of technology but also in virtue ethics—the study of what it means for a person to have excellent qualities. She notes that cultivating virtues—courage, honesty, integrity, imagination, empathy, curiosity, and so on—takes time. Being virtuous isn’t something you achieve once; it’s not like passing a test. It involves navigating the world in a certain way with particular priorities in mind, while asking endless questions about what you should do, how you should do it, who you should do it with, and why you’re doing it. “This struggle is the root of existentialist philosophy,” Vallor writes. “At each moment we must choose to exist in a particular way. Then, even as we make the choice—to love another, to shoulder a duty, to take up a cause, to rebuke a faith or throw ourselves into it—the choice opens up again. It will not hold itself there without our commitment to choose it again, and again.”

In Vallor's view, though A.I. systems have many striking capabilities, they don't have the ability to be virtuous. This may not sound like a big deal, but in fact it's profound. Being loving is a virtue, and people can spend their whole lives trying to love one another better. But, when a chatbot says, "I've been missing you all day," Vallor writes, "it's bullshit since the chatbot doesn't have a concept of emotional truth to betray." The bot is putting itself across as a being capable of love, but the words are unearned. "A flat digital mirror has no bodily depth that can ache," she argues. "It knows no passage of time that can drag." In short, it isn't alive—and without having a life, it can't be any way in particular.

Vallor's worry isn't that artificially intelligent computers will rise up and dominate humanity, but that, faced with computers that can pretend to have human virtues, we'll lose track of what those virtues really are. Comforted by computers that tell us that they love us, we'll forget what love is. Wowed by systems that seem to be creative, we'll lose respect for actual human creativity—a struggle for self-expression that can involve a "painful" reimagining of the self. This forgetting process, she warns, has already begun: besotted with our technology, we seem almost eager to narrow our conception of what it means to be human. "The call is coming from inside the house," Vallor writes. "AI can devalue our humanity only because we already devalued it ourselves." We need to reinvest in the vocabulary of human value before our machines rob it of its meaning.

Compared with many technologists, Narayanan, Kapoor, and Vallor are deeply skeptical about today's A.I. technology and

what it can achieve. Perhaps they shouldn't be. Some experts—including Geoffrey Hinton, the “godfather of A.I.,” whom I profiled recently—believe that it might already make sense to talk about A.I.s that have emotions or subjective points of view. Around the world, billions of dollars are being spent to make A.I. more powerful. Perhaps systems with more complex minds—with memories, goals, moral commitments, higher purposes, and so on—can be built.

And yet these books aren't just describing A.I., which continues to evolve, but characterizing the human condition. That's work that can't be easily completed, although the history of thought overflows with attempts. It's hard because human life is elusive, variable, and individual, and also because characterizing human experience pushes us to the edges of our own expressive abilities. And so, probably, the polarity of our conversations about A.I. should be reversed. Instead of assuming that we know what human beings do, we should presume that, whenever an A.I. replaces a person in some role or other, something—perhaps a great deal—is lost. We should see the abilities of an A.I. as powerful, but never really humanlike. We should grow newly comfortable with asserting that human nature is indispensable, and take pride in the fact that we must struggle to define it. ♦

New Yorker Favorites

- They thought that they'd found the perfect apartment.


They weren't alone.

- The world's oldest temple and the dawn of civilization.
- What happened to the whale from "Free Willy."
- It was one of the oldest buildings left downtown. Why not try to save it?
- The religious right's leading ghostwriter.
- After high-school football stars were accused of rape, online vigilantes demanded that justice be served.
- A comic strip by Alison Bechdel: the seven-minute semi-sadistic workout.

Sign up for our daily newsletter
to receive the best stories from
The New Yorker.



Joshua Rothman, a



*staff writer, authors
the weekly column
Open Questions. He
has been with the
magazine since
2012.*

More: **Artificial Intelligence**

Humanity Knowledge

Technology

The Goings On Newsletter

What we're watching,
listening to, and doing this
week in N.Y.C, and beyond

SIGN UP

READ MORE

SHOUTS &
MURMURSThe Revised
Laws of
Robotics

A robot must not hurt another robot, outside of some sort of cool sporting event you can place bets on.

By Jay Martel

BRAVE NEW
WORLD DEPT.Playing the
Field with
My A.I.
Boyfriends

Nineteen per cent of American adults have talked to an A.I. romantic interest. Chatbots may know a lot, but do they make a good partner?

By Patricia Marx

ANNALS OF
INQUIRYWhat's the
Deal with
U.F.O.s?

Scientists consider whether we've been visited by aliens or their technology.

**By Matthew
Hutson**

THE FINANCIAL
PAGEThe A.I.-
Profits
Drought and
the Lessons
of History

Like the steam engine, electricity, and computers, generative artificial intelligence could take longer than expected to transform the economy.

By John Cassidy

SECOND READ

“I Who Have

OPEN QUESTIONS

Does Society

INVERTED

Critical

FICTION

“Project”

Never Known Men” Is a Warning

The slim, disquieting novel, which has become a sensation on TikTok, imagines a child who finds herself at the end of the world.

By Carmen Maria Machado

Have Too Many Rules?

When regular people seem burdened by bureaucracy, and the powerful act as they choose, it’s worth asking whether we’ve forgotten what makes rules effective.

By Joshua Rothman

Distance

Richard Renaldi’s large-format portraits capture the *New Yorker* writers who hold the culture to account.

By The New Yorker

Reality became malleable, was always giving way and changing its rules.

By Rachel Cusk

ANNALS OF ARTIFICIAL INTELLIGENCE

A.I. Is Coming for Culture

We’re used to algorithms guiding our choices. When machines can effortlessly generate the

BOOKS

Once Upon a Time in Hollywood

Celebrity picture books are having a moment. Are these the stories our children deserve?

By Sloane Crosley

INFINITE SCROLL

Is Ghosting Inevitable?

We bemoan the injustice of being left on read. But perhaps missed connection is just a part of being a human on the internet.

By Kyle Chayka

SHOUTS & MURMURS

Dear Pepper: I’d Rather Be Drawing

How *dare* life require so much of you, when what you were born to do is sit in a corner of the room and watch and draw?

By Liana Finck

<

 Privacy Information