

MIT Technology Review

Volume 125 Nov/Dec USD \$12.00
Number 6 2022 CAD \$14.00

Bots for the
brokenhearted

Reprogramming
aging bodies

What really happens
when you donate
your body to science

IS THERE A LIMIT



TO HUMAN LIFE?

MIT
Technology
Review
Insights

The Green Future Index 2022

The Green Future Index 2022 is the second edition of the comparative ranking of 76 nations and territories on their ability to develop a sustainable, low-carbon future. It measures the degree to which economies are pivoting toward clean energy, industry, agriculture, and society through investment in renewables, innovation, and green policy.

The index ranks the “green” performance of countries and territories across five pillars:

- Carbon emissions
- Energy transition
- Green society
- Clean innovation
- Climate policy



KEY

- Green leaders
- The greening middle
- Climate laggards
- Climate abstainers
- ↑ Countries that have gone up in the ranking since last year
- ↔ Countries that have retained the same ranking as last year
- ↓ Countries that have gone down in the ranking since last year

Overall top 10

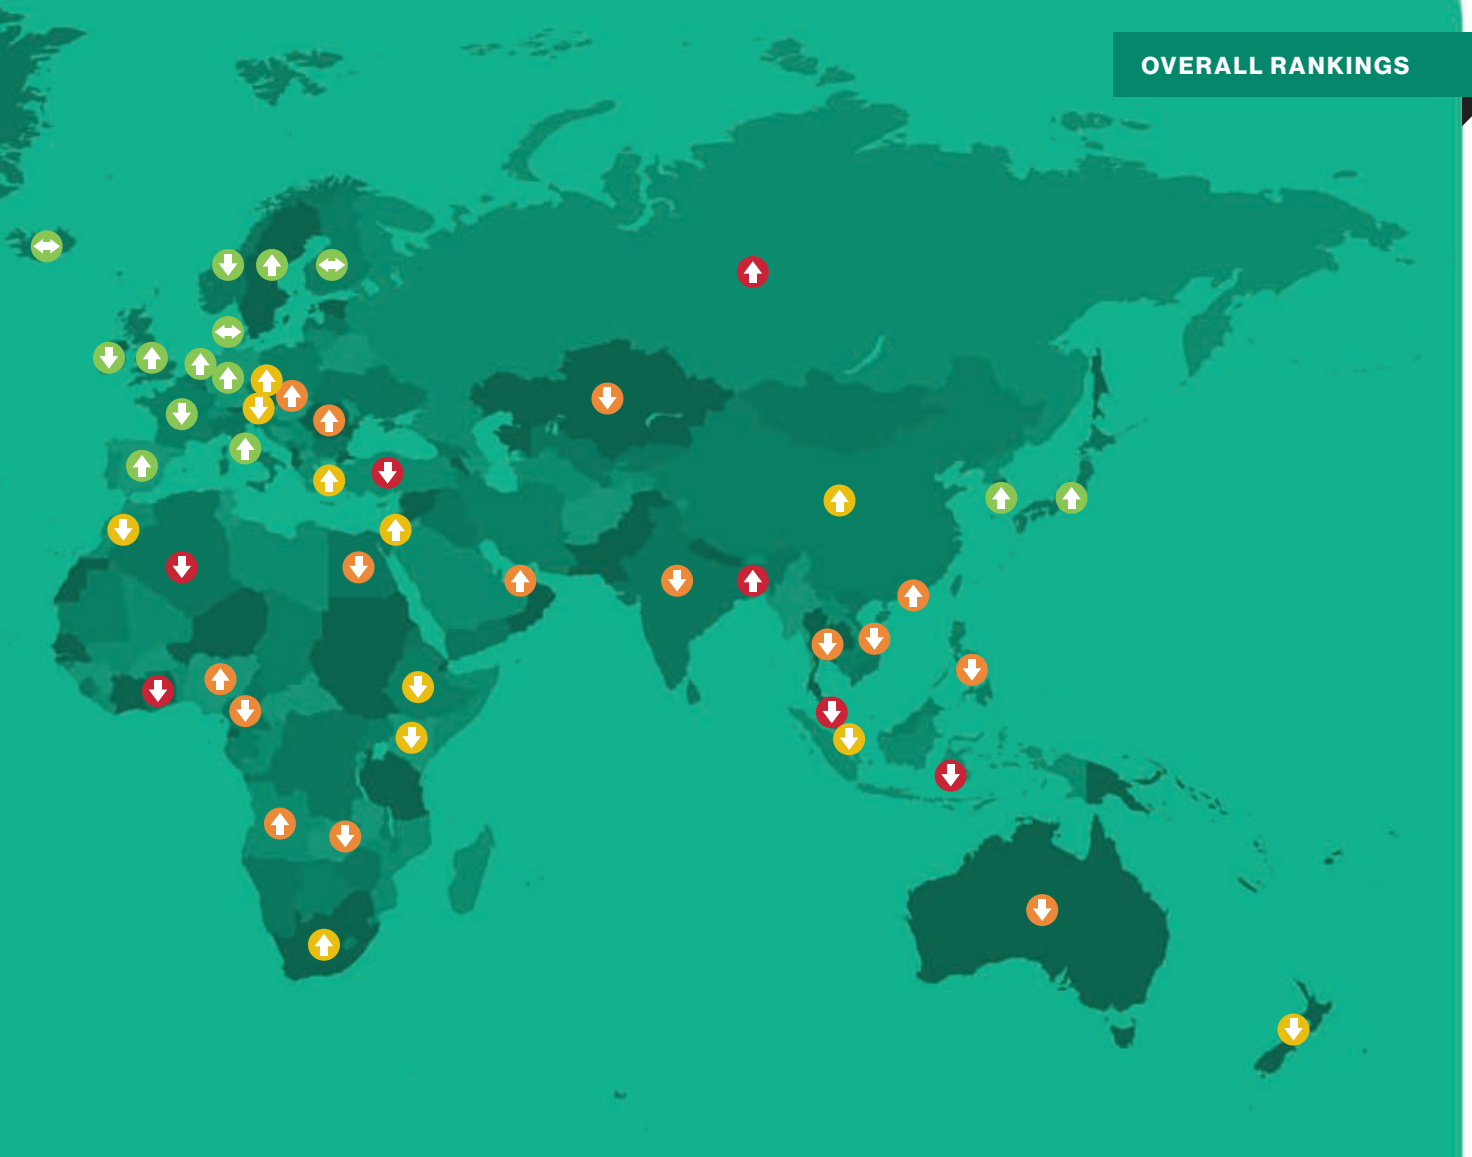
Rank 2022	Rank 2021	Territory	Score	Rank 2022	Rank 2021	Territory	Score
1	↔ 1	Iceland	6.9	6	↔ 6	Finland	6.2
2	↔ 2	Denmark	6.6	7	↓ 4	France	6.1
3	↑ 10	Netherlands	6.4	8	↑ 11	Germany	6.1
4	↑ 17	United Kingdom	6.3	9	↑ 12	Sweden	6.1
5	↓ 3	Norway	6.2	10	↑ 31	South Korea	6.0

- Nearly 36% of UK's power came from clean sources in the third quarter of 2021, with the aim of reaching 100% by 2035.
- In January 2022, Finland took on €217 million in pre-financing, which will partially go to efforts to decarbonize the energy sector.
- South Korea and Japan have seen significant rises in their innovation scores, thanks to their world-beating green IP contributions.

Experience the interactive index, view the data, and download the full report at:

technologyreview.com/gfi

OVERALL RANKINGS



Green society top 10

Rank 2022	Rank 2021	Territory	Score	Rank 2022	Rank 2021	Territory	Score
1	3	South Korea	7.0	6	17	Iceland	6.4
2	1	Singapore	6.8	7	4	Taiwan	6.4
3	2	Ireland	6.8	8	5	Philippines	6.3
4	8	Germany	6.5	9	6	Czech Republic	6.3
5	7	United States	6.5	10	13	Canada	6.2

- This pillar ranks each country on how well it is pivoting toward clean energy, industry, agriculture, and society.
- Leaders in the green society pillar are over-represented by nations that have incorporated strong civil planning and societal development goals into policy, regulation, and public infrastructure spending.
- Singapore and South Korea are the world's best-ranked recycling economies.

Interested in partnering with us? Contact: insights@technologyreview.com

Never say die

Hi. How are you? I have some news. You're going to die. We're all going to die.

I've been thinking about this not only because we've been putting this issue, the Mortality issue, together but also because I've hit one of those arbitrary milestones in life. Which is to say I turn 50 this month.

This is mostly pretty cool. Beats the alternative, as they say. But one decidedly uncool thing about being almost 50 is that my "bad" LDL cholesterol is too high. The obvious solution here is to improve my diet and get more exercise. But if that were so easy, no one would have high cholesterol, right? Another option might be to go on statins—use medication to bring it down. Lots of people do this, but it's a treatment, not a cure, and the side effects can be intense.

Or, just maybe, I could have my genes edited using CRISPR, as a patient in New Zealand recently did, to permanently reduce my LDL. Given that heart disease is a leading cause of death, that could meaningfully prolong my life. There's still much to learn about the effectiveness of such a treatment, but the mere prospect is pretty amazing.

The things we can do to alter our bodies today would have sounded like science fiction just a few years ago. Scratch that—some of them still sound like science fiction. Aging and death have long been our inevitable fates, constants of life that begin with birth. But what if aging isn't inevitable? What if we could slow or even reverse the process? This may be less far-fetched than you imagine.

A flotilla of billionaires and Silicon Valley elites, and even the Saudi royal family, have poured staggering amounts of money into this question, funding startups and research that seek to

extend our time on Earth. Antonio Regalado has the story of the quest for medical rejuvenation: researchers want to reprogram our cells to make them young again, letting us lead longer, healthier lives (page 54).

Or even more intriguing: What if we could cheat death and live forever? Because we can. Or at least our digital replicas can, as Charlotte Jee learned when she trained an AI to create simulacra of her parents (page 26). The technology she explored is one of several meant to let the living commune with the dead, in this case via Amazon Alexa.

But even if we can live forever in a smart speaker, at the moment there's still no way around death. Of course, there are a lot of people working on that. Perhaps you might want to have your corpse preserved at -196°C in the hopes of being brought back to life again one day. The field of cryonics is still trucking along, even if it is "a hopeless aspiration that reveals an appalling ignorance of biology," as one scientist puts it (page 34).

For those who have accepted their fate but want their death to help others live longer, there is always the possibility of donating one's body to science. While that's something you have surely heard of, it's nothing like what you have probably imagined. Abby Ohlheiser pulls back the curtain (page 62) to reveal how this intimate process truly works.

I hope this issue serves up a lot for you to think about. As always, I'd love to hear your thoughts. You can reach me at mat.honan@technologyreview.com.

See you on the other side,

Mat



Using data to elevate the business of
people flow.



www.technologyreview.com/thecloudhub

Infosys®

“Fundamentally, people don’t like grappling with the fact they are going to die.”



Front

- 2 Letter from the editor

THE DOWNLOAD

- 9 The artist dominating AI-generated art; job titles of the future; robotic bees; SpaceX engineers who never existed; pet life extension; the problem with EVs; and the impacts of insect extinction.

EXPLAINED

- 18 **When we lose weight, where does it go?**
By Bonnie Tsui

PROFILE

- 20 **Diving into DAIR**
Alex Hanna aims to help right systemic issues in AI research. By Anmol Irfan

The mortality issue

- 23 **Is old age a disease in its own right?**

A formal diagnosis could help unlock further support for longevity research. Or it could just be a matter of semantics. BY SARAH SLOAT

- 26 **Bots for the brokenhearted**

Digital clones of the people we love could change how we grieve for them. BY CHARLOTTE JEE

- 34 **The hopeful dead**

The field of cryonics has survived for decades not on fact, but on faith in the possibility of human reanimation. Is there any reason to think things are different now? BY LAURIE CLARKE

- 36 **The bird is fine, the bird is fine, the bird is fine, it’s dead**

The pursuit of immortality is getting older. So are we. BY JONATHAN WEINER

- 42 **Effective altruism and its growing influence**

Can data really quantify the “best way” to help the world? BY REBECCA ACKERMANN

- 52 **The death of knowledge**

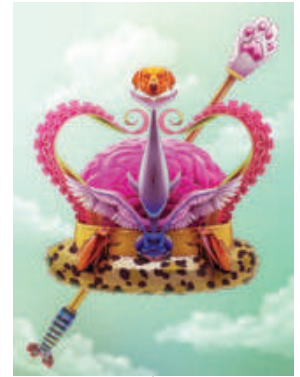
We can preserve a lot of information digitally, but the process is far from perfect. BY ERIK SHERMAN

- 54 **How to become young again**

Research labs are pursuing technology to “reprogram” aging bodies back to youth. BY ANTONIO REGALADO

- 62 **Taking care of the dead**

Donating one’s body to science can advance medicine and benefit the living—if handled properly. BY ABBY OHLHEISER



Back

- 72 **Life, death, and automation**

AI is being asked to help make life-or-death decisions. This mustn’t be an excuse to turn away from painful choices. By Will Douglas Heaven

- 78 **Tech’s downward spiral**

Innovation that truly serves us all is in scarce supply. That’s a problem. By Shannon Vallor

- 80 **Inside the minds of animals**

Three authors argue that we should spend more time understanding the creatures around us. By Matthew Ponsford

ARCHIVE

- 88 **Is there a limit to human life?**

We’ll likely never stop searching for the answer.

SAVE THE DATE

MIT Technology Review's signature AI conference returns in person March 2023.

MIT
Technology
Review

**EmTech
Digital**

**JOIN US IN CAMBRIDGE, MA
OR ONLINE**

March 28–29, 2023

**SUBSCRIBERS SAVE 10%
WITH CODE PRINTND22 AT**

EmTechDigital.com

Editorial

Editor in chief
Mat Honan

Executive editor, operations
Amy Nordrum

Executive editor, newsroom
Niall Firth

Editorial director, print
Allison Arieff

Editorial director, audio and live journalism
Jennifer Strong

Editor at large
David Rotman

News editor
Charlotte Jee

Features and investigations editor
Amanda Silverman

Managing editor
Timothy Maher

Commissioning editor
Rachel Courtland

Senior editor, MIT News
Alice Dragoon

Senior editor, biomedicine
Antonio Regalado

Senior editor, climate and energy
James Temple

Senior editor, digital culture
Abby Ohlheiser

Senior editor, AI
Will Douglas Heaven

Podcast producer
Anthony Green

Senior reporters
Tanya Basu (humans and technology)
Eileen Guo (features and investigations)
Jessica Hamzelou (biomedicine)
Melissa Heikkilä (AI)
Tate Ryan-Mosley (tech policy)

Reporters
Casey Crownhart (energy and climate)
Rhiannon Williams (news)
Zeyi Yang (China and East Asia)

Copy chief
Linda Lowenthal

Senior audience engagement editor
Abby Ivory-Ganja

Audience engagement editor
Juliet Beauchamp

Editorial fellows
Hana Kiros
Tammy Xu

Creative director, print
Eric Mongeon

Digital visuals editor
Stephanie Arnett

Corporate

Chief executive officer and publisher
Elizabeth Bramson-Boudreau

Assistant to the CEO
Alexandra Rivera-Palmieri

Human resources director
Alyssa Rousseau

Manager of information technology
Colby Wheeler

Office manager
Linda Cardinal

Technology

Chief technology officer
Drake Martinet

Vice president, product
Mariya Sitnova

Senior software engineer
Molly Frey

Associate product manager
Allison Chase

Digital brand designer
Vichhika Tep

Events

Senior vice president,
events and strategic partnerships
Amy Lammers

Director of event content and experiences
Brian Bryson

Head of international and custom events
Marcy Rizzo

Senior event content producer
Erin Underwood

Director of events
Nicole Silva

Event operations manager
Elana Wilner

Manager of strategic partnerships
Madeleine Frasca

Event coordinator
Bo Richardson

Finance

Vice president, finance
Enejda Xheblati

General ledger manager
Olivia Male

Accountant
Anduela Tabaku

Consumer marketing

Vice president, marketing and
consumer revenue
Alison Papalia

Director of retention marketing
Taylor Puskaric

Email marketing manager
Tuong-Chau Cai

Circulation and print production manager
Tim Borton

Advertising sales

Senior vice president, sales and
brand partnerships
Andrew Hendler
andrew.hendler@technologyreview.com
201-993-8794

Executive director, integrated marketing
Caitlin Bergmann
caitlin.bergmann@technologyreview.com

Executive director, brand partnerships
Marii Sebahar
marii@technologyreview.com
415-416-9140

Executive director, brand partnerships
Kristin Ingram
kristin.ingram@technologyreview.com
415-509-1910

Executive director, brand partnerships
Stephanie Clement
stephanie.clement@
technologyreview.com
214-339-6115

Senior director, brand partnerships
Debbie Hanley
debbie.hanley@technologyreview.com
214-282-2727

Senior director, brand partnerships
Ian Keller
ian.keller@technologyreview.com
203-858-3396

Senior director, brand partnerships
Miles Weiner
miles.weiner@technologyreview.com
617-475-8078

Digital sales strategy manager
Casey Sullivan
casey.sullivan@technologyreview.com
617-475-8066

Media kit
www.technologyreview.com/media

**MIT Technology Review Insights
and international**

Vice president, Insights and international
Nicola Crepaldi

Global director of custom content
Laurel Ruma

Senior manager of licensing
Ted Hu

Senior editor, custom content
Michelle Brosnahan

Senior editor, custom content
Kwee Chuan Yeo

Editor, custom content
Teresa Elsey

Senior project manager
Martha Leibs

Project manager
Natasha Conteh

Director of partnerships, Europe
Emily Kutchinsky

Director of partnerships, Asia
Marcus Ulvne

Board of directors

Cynthia Barnhart, Cochair
Alan Spoon, Cochair
Lara Boro
Peter J. Caruso II, Esq.
Whitney Espich
Sanjay E. Sarma
David Schmittlein
Glen Shor

**Customer service and
subscription inquiries**

National
877-479-6505

International
847-559-7313

Email
customer-service@technologyreview.com

Web
www.technologyreview.com/
customerservice

Reprints
techreview@wrightsmedia.com
877-652-5295

Licensing and permissions
licensing@technologyreview.com

MIT Technology Review

196 Broadway, 3rd Floor
Cambridge, MA 02139
617-475-8000

Our in-depth reporting reveals what's
going on now to prepare you for what's
coming next.

Technology Review, Inc., is an independent nonprofit 501(c)(3) corporation wholly owned by MIT; the views expressed in our publications and at our events are not always shared by the Institute.

GIVE THE PERFECT PRESENT

Look no further for the curious,
tech-obsessed people on your gift list!

Go to [TRsub.com/Gift](https://www.mitreview.com/gift)



The gift that gives **all year long**:

- In-depth, thoughtful stories on the latest technologies coming out of labs and innovation centers around the world
- Insights into the new technologies that will transform businesses and industries
- No hype, only the most important tech news curated and interpreted for you by MIT Technology Review
- All accessible in multiple formats (online, print, and a new app coming soon)

MIT
Technology
Review

Don't believe everything you hear.



The award-winning podcast **In Machines We Trust** thoughtfully examines the far-reaching impact of artificial intelligence on our daily lives.

Download it wherever you listen.



The Download

This artist is dominating AI-generated art. He's not happy about it.

Greg Rutkowski is a more popular prompt than Picasso.

By Melissa Heikkilä

Those cool AI-generated images you've seen across the internet? There's a good chance they are based on the works of Greg Rutkowski.

A Polish digital artist who uses classical painting styles to create dreamy fantasy landscapes, Rutkowski has made illustrations for games such as Sony's Horizon Forbidden West, Ubisoft's Anno, Dungeons & Dragons, and Magic: The Gathering. And he's become a sudden hit in the new world of text-to-image AI generation.

His distinctive style is now one of the most commonly used prompts in the new open-source AI art generator Stable Diffusion, which was launched in late August. The tool, along with other popular image-generation AI models, allows anyone to create impressive images based on text prompts. For example, type in "Wizard with sword and a glowing orb of magic fire fights a fierce dragon Greg Rutkowski," and the system will produce something that looks not a million miles away from works in Rutkowski's style.

But these open-source programs are built by scraping images from the internet, often without permission and proper attribution to artists. As a result, they are raising tricky questions about ethics and copyright. And artists like Rutkowski have had enough.

According to the website Lexica, which tracks over 10 million Stable Diffusion images and prompts, Rutkowski's name has been used as a prompt around 93,000 times. Some of the world's most famous artists, such as Michelangelo, Pablo Picasso, and Leonardo da Vinci, brought up around 2,000 prompts each or



less. Rutkowski's name also features as a prompt thousands of times in the Discord channel of another text-to-image generator, Midjourney.

Rutkowski was initially surprised but thought it might be a good way to reach new audiences. Then he tried searching for his name to see if a piece he had worked on had been published. The online search brought back work that had his name attached to it but wasn't his.

"It's been just a month. What about in a year? I probably won't be able to find my work out there because [the internet] will be flooded with AI art," Rutkowski said in September. "That's concerning."

Some artists may have been harmed in the process

Other artists besides Rutkowski have been surprised by the apparent popularity of their work in text-to-image generators—and some are now fighting back.

Karla Ortiz, an illustrator based in San Francisco who found her work in Stable Diffusion's data set, has been raising awareness about the issues around AI art and copyright. Artists say they risk losing income as people start using AI-generated images based on copyrighted material for commercial purposes. But it's also a lot more personal, Ortiz says, arguing that because art is so closely linked to a person, it could raise data protection and privacy problems.

"There is a coalition growing within artist industries to figure out how to tackle or mitigate this," says Ortiz. The group is in its early days of mobilization, which could involve pushing for new policies or regulation. One suggestion is that AI models could be trained on images in the public domain, and AI companies could forge partnerships with museums and artists, Ortiz says.

AI-generated art poses tricky legal questions. In the UK, where Stability.AI is based, scraping images from the internet without the artist's consent to train an AI tool could be a copyright infringement, says Gill Dennis, a lawyer at the firm Pinsent Masons. Copyrighted works can be used to train an AI under "fair use," but only for noncommercial purposes. The UK, which hopes to boost domestic AI development, wants to change laws to give AI developers greater access to copyrighted data. Under these changes, developers would be able to scrape works protected by copyright to train their AI systems for both commercial and noncommercial purposes.

While artists and other rights holders would not be able to opt out of this regime, they would be able to choose where they make their works available. The art community could end up moving into a pay-per-play or subscription model like the one used in the film and music industries. Rutkowski says he doesn't blame people who use his name as a prompt. For them, "it's a cool experiment but for me and many other artists, it's starting to look like a threat to our careers." *To read the full story, visit www.technologyreview.com.* ■



Robotic bees to the rescue

A team hopes to build robots to help honeybees and the ecosystems around them.

By Elizabeth Preston

Something was wrong, but Thomas Schmickl couldn't put his finger on it. It was 2007, and the Austrian biologist was spending part of the year at East Tennessee State University. As he made his daily walk across some fields to campus, "it felt unpleasant," he says. "And I didn't realize why until I heard a bumblebee."

Insects, he realized, had otherwise seemed conspicuously absent. "I was dragging the biology professors out of the building and saying, 'Look in the sky—there is nothing flying!'" he recalls.

Schmickl, who now leads the Artificial Life Lab at the University of Graz in Austria, wasn't wrong. Studies in various parts of the world have since found that insect populations are declining or changing. After working in the field of swarm robotics for several years—using nature to inspire robots—he decided to

flip his work around and design robots to help the natural world. He's focusing on bees. Honeybees and other pollinators face habitat loss, pesticide exposure, and other challenges, and Schmickl believes that coming to their aid could help strengthen entire ecosystems.

Already, some companies offer augmented beehives that monitor conditions inside, or even robotically tend the bees. Schmickl and his colleagues want to go a step further and use technology to manipulate the insects' behavior.

As part of a European Union-funded project called Hiveopolis, they're building prototype hives outfitted with sensors and cameras as well as devices that can create vibration inside the hive and adjust temperature or air flow. Such tools could ultimately direct the bees' traffic patterns: Schmickl's experiments have shown that vibration slows the bees down, while moving air encourages them to walk away.

Hiveopolis collaborator Tim Landgraf, a professor of artificial and collective intelligence at Freie Universität Berlin in Germany, is working on another kind of tool for these hives: a robotic dancing bee. When real honeybees return from foraging, they perform a distinctive "waggle dance" that communicates the location of the food. When enough bees are doing the same dance, they'll fly out to find the food.

In earlier research, Landgraf built a robot that could perform a waggle dance sufficiently convincing that other bees followed it—and, at least sometimes, flew in the direction the robot suggested. Now he's getting ready to test an improved version. The robot doesn't look very bee-like to a human eye. Its body is simply a small, flexible tube with a fluttering "wing." But it's connected to a motor outside the hive that can steer and shimmy it across the hive's dance floor.

In theory, such a robot could guide honeybees to a safe foraging site if humans determined that another site was contaminated with pesticides that were affecting the hive's health, Schmickl says. Or humans could direct honeybees away from a site that was being reserved for wild bees.

Elina L. Niño, a honeybee expert at the University of California, Davis, says it makes more sense to focus on creating healthier environments for bees, "so we don't have to be concerned with regulating their flight and their foraging." While the beekeepers she works with "would just kind of chuckle at this," Niño says, from a research perspective, the European research is "super exciting." Observing how honeybees interact with the waggle robot, for example, could give scientists new insights into the ways bees communicate. ■

Job titles of the future: Robotaxi safety operator

Liu Yang, 33

China requires a person to be present in autonomous vehicles. Liu's job at Baidu, the AI giant, is to sit in the car. The taxi moves tourists around Shougang Park, a 3.3-square-mile redevelopment that was part of the 2022 Winter Olympics.

Qualifications needed: A robotaxi minder needs a passion for being in cars. Liu has an associate's degree in human resources, but he worked as a driver. He spends seven hours a day in the taxi.

Weirdest moment: Liu found it hard to go back to driving his own car. "I instinctively went for the passenger seat," says the 33-year-old Beijing native, who joined the Chinese tech giant's autonomous-vehicle division in January 2021.

Career prospects: Robotaxi safety operator could actually be a short-lived occupation. Baidu moved Liu from behind the wheel to the passenger seat in July, after winning a license expansion. But companies hope to convince governments that they can do away with human intervention altogether.

"I want to stay in this industry. I have high hopes for it, and I like it a lot. I entered this industry because I was curious, and now I still believe it's advanced and powerful. I think it's very futuristic," says Liu. "I still want to stay in self-driving." ■



The 1,000 Chinese SpaceX engineers who never existed

LinkedIn users are being scammed out of millions of dollars by fake connections posing as graduates of prestigious universities and employees at top tech companies.

By Zeyi Yang

If you were just looking at his LinkedIn page, you'd certainly think Mai Linzheng was a top-notch engineer. With a bachelor's degree from Tsinghua, China's top university, and a master's degree in semiconductor manufacturing from UCLA, Mai began his career at Intel and KBR, a space tech company, before ending up at SpaceX in 2013. Having spent the past eight years and nine months working in the human race to space, he's now a senior technician.

Except all is not as it seems.

Upon closer inspection, there are plenty of red flags: Despite having been in the US for 18 years, Mai has written all his job titles, degrees, and company locations in Chinese. His bachelor's degree is in business management, even though his alma mater, Tsinghua, only offers that degree to student athletes, and Mai was not one. Besides, the man in his profile photo looks younger than Mai's stated age. The image, as it turns out, was stolen from Korean influencer Yang In-mo's Instagram. In fact, none of the information on this page is true.

The profile of "Mai Linzheng" is actually one of the millions of fraudulent pages set up on LinkedIn to lure users into scams, often involving cryptocurrency investments and targeting people of Chinese descent all over the world. Scammers like Mai claim affiliation with prestigious schools and companies to boost their credibility before connecting with other users, building a relationship, and laying a financial trap.

Since last year, such activities have been steadily on the rise on LinkedIn, following

years of proliferation on other social media platforms and dating apps. In the second half of 2021, LinkedIn removed 7% more profiles because of fraudulent identities than in the six months before that, according to Oscar Rodriguez, LinkedIn's senior director of trust, privacy, and equity. "Scammers are highly sophisticated and proactive in terms of how often they adapt tactics," he says. For instance, a week after the Biden administration announced its student loan forgiveness plan, LinkedIn started seeing scammers incorporating the news into their scripts.

By now, victims have lost millions of dollars through scams that originated on the platform. This summer, the FBI announced it would investigate these scams and work with victims to identify the bad actors and disable their accounts, even though the financial losses are almost impossible to recover.

Scammers "are always thinking about different ways to victimize people, victimize companies," Sean Ragan, the FBI's special agent in charge of the San Francisco and Sacramento field offices, told CNBC in June. "And they spend their time doing their homework, defining their goals and their strategies and their tools and tactics that they use." He called the work of these criminals a "significant threat."

A SpaceX "employee" invited you to connect

At one point in July, there were over 1,000 LinkedIn profiles for individuals who, like "Mai Linzheng," claimed to have graduated

from Tsinghua University and to work at SpaceX. The eye-popping number even triggered patriotic Chinese influencers to lament the brain drain and accuse Chinese university graduates of disloyalty to their country.

This caught the attention of Jeff Li, a Toronto-based tech influencer and columnist at Financial Times China. He confirmed on July 11 that he could find 1,004 Tsinghua graduates by searching for SpaceX employees on LinkedIn; this would have made the alumni group the largest at the company. But many accounts he saw claimed the exact same education and work experiences—suggesting that someone was mass-generating fake profiles.

"They all graduated from Tsinghua and went on to the University of Southern California or similar well-known universities," Li says. "Besides that, they all worked at a certain company in Shanghai. Obviously, I suspect these are fake, generated data."

(SpaceX did not reply to a request from MIT Technology Review asking to confirm the number of Tsinghua graduates working at the company.)

This wasn't the first time Li had noticed what he thought were fake LinkedIn accounts. Starting in late 2021, he says, he started seeing profiles with under a few dozen connections—rare for real users—and with profile photos that were always good-looking men and women, likely stolen from other websites. Most appeared to be of Chinese ethnicity and to live in the United States or Canada. In recent years, as China has cracked down on fraudulent online activities, these operations have pivoted to targeting people elsewhere who are of Chinese descent or speak Mandarin. The Global Anti-Scam Org (GASO) was established in July 2021 by one such victim, and the organization now has nearly 70 volunteers on several continents.

While these fake accounts are relatively new to LinkedIn, they have permeated other platforms for a long time. "Scammers started moving to LinkedIn maybe after dating sites tried to crack down on them,



LinkedIn scammers may target people of Chinese descent by claiming similar experiences or affiliations.

[like] Coffee Meets Bagel, Tinder,” says Grace Yuen, a GASO spokesperson.

In certain ways, LinkedIn is a great way for fraudsters to expand their reach. “You might be already married and you are not on the dating sites, but you probably have a LinkedIn account that you check occasionally,” says Yuen.

A scammer on LinkedIn may try to connect with someone through common work experience, a shared hometown, or the feeling of living in a foreign country. Over 60% of the victims who have reached out to GASO are Chinese immigrants or have Chinese ancestry, which these actors lean on to evoke nostalgia or a desire for companionship. The fake claims to have graduated from China’s top universities, which are notoriously difficult to get into, also help scammers earn respect.

On average, LinkedIn victims in particular tend to lose more money than

victims of fraud on other platforms—oftentimes over a million dollars, says Yuen.

“Unlike dating sites, which are where the first scam victims were coming from, LinkedIn actually has a lot of information that’s really useful for the scammers,” she says. “They know your earning potential based on the type of work you listed.”

The responsibility for preventing these scams, though, also falls on the sites where perpetrators hunt for their victims in the first place. After several media reports about the rampant scams on LinkedIn, the platform released a report in June that says it has been able to detect 96% of fake accounts before the people behind them make any contact with users.

LinkedIn does this through a mix of algorithms, industry expert suggestions, and human user reports, says Rodriguez. It looks for behavioral signals, like whether a

new account immediately starts to message other users, and whether any of these users block or flag the account. To LinkedIn’s credit, Li, who confirmed the presence of fake SpaceX engineers on the platform, says this year he has noticed that scam accounts are being taken down more quickly. “At the end of last year, the account might survive three or four days; now they’re being taken down in hours,” he says.

But anyone who searches LinkedIn today for SpaceX employees who graduated from Tsinghua University is still likely to find around 200 results—including “Mai Linzheng” and other fakes. Generally speaking, Rodriguez explains, the platform prioritizes identifying fake accounts that are actively engaging with real users; accounts that remain could have been dormant after registration. *To read the full story, visit www.technologyreview.com.* ■

These scientists are working to extend the life span of pet dogs—and their owners

Anti-aging drugs are being trialed in companion dogs—but the goal is to find ways to have people, as well as beloved pets, live longer, healthier lives.

by Jessica Hamzelou

Matt Kaeberlein is what you might call a dog person. He has grown up with dogs and describes his German shepherd, Dobby, as “really special.” But Dobby is 14 years old—around 98 in dog years. “I’m very much seeing the aging process in him,” says Kaeberlein, who studies aging at the University of Washington in Seattle.

Kaeberlein is co-director of the Dog Aging Project, an ambitious research effort to track the aging process of tens of thousands of companion dogs across the US. He is one of a handful of scientists on a mission to improve, delay, and possibly reverse that process to help them live longer, healthier lives.

The Dog Aging Project is just one of several groups seeking to understand and improve dog aging. Biotech company Loyal has plans to offer life extension for dogs. And a third group, running a project called Vaika, is looking for ways to lengthen life span through a study on retired sled dogs.

Matt Kaeberlein with Chloe and Dobby.



But dogs are just the beginning. Because they are a great model for humans, anti-aging or life-span-extending drugs that work for dogs could eventually benefit people, too. In the meantime, attempts to prolong the life of pet dogs can help people get on board with the idea of life extension in humans, say researchers behind the work. “It will go a long way to convincing people that this is possible [in humans],” says Kaeberlein. “Aging is modifiable.”

For the love of dogs

“I love dogs,” says Kate Creevy, who studies dog aging and infectious disease in animals at Texas A&M University. “You’re not motivated to do this work if you don’t love dogs.” Creevy, who is also chief veterinary officer of the Dog Aging Project, is one of around 40,000 people with a dog enrolled in the study.

All participants provide their pet dog’s medical history and complete detailed surveys on an annual basis. “It takes about three hours,” says Creevy. A subset of around 8,500 dogs will have their genomes sequenced, and some of these will have their hair, blood, and urine studied as well. Smaller groups of dogs are being more closely studied for specific disorders. The team will assess 200 dogs with a form of dementia known as canine cognitive dysfunction, or CCD, for example.

The idea is to find biological clues that might help identify which dogs might be



Kate Creevy with Poet and Sophie.

at risk of developing such diseases in the future—and eventually aid the discovery of drugs that might prevent or treat them. The team also hopes to find out which aspects of a dog’s lifestyle might help extend its “health span,” the number of years lived in good health.

“We expect to learn which types of diets, which types of exercise regimes, and which types of husbandry are associated with better long-term outcomes,” says Creevy, “so that we can do things that help them have a better quality of life into their later years.”

But the research has another goal. Kaeberlein says his “lightbulb moment” occurred around 10 years ago, when he suddenly realized that not only would such research reveal how dogs age—it might identify ways to slow the process. “I was like, man, I would love if I could slow aging in my dog,” he recalls.

The Dog Aging Project will trial potential anti-aging drugs among groups of pets. The first is rapamycin, a drug that has been found to extend the lives of flies, worms, and mice in the lab. “I’m convinced that some of the interventions that we know extend life span and health span in mice will work in dogs,” says Kaeberlein. “It’s really just a matter of showing it through clinical trials.” *To read the full story, visit www.technologyreview.com.* ■

How climate change is screwing with insects

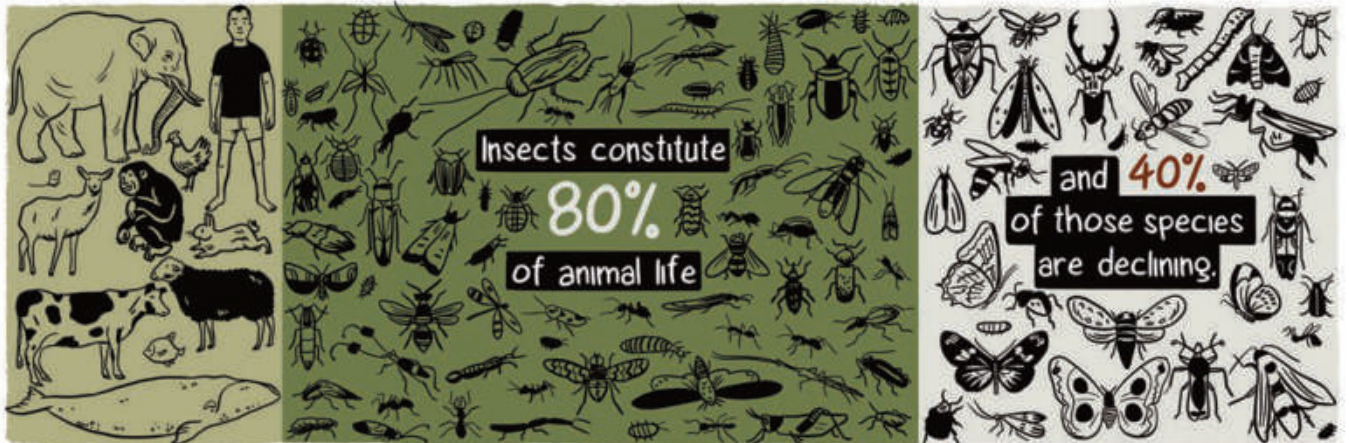
Hotter temperatures are wiping out species—but increasing invasive pest populations.

By Susie Cagle

We are in the midst of earth's sixth mass extinction -



its first caused by humans.



Insects constitute **80%** of animal life

and **40%** of those species are declining.

A new study published this year linked this rolling extinction to the climate crisis and industrial agriculture processes across the world.



DR. CHARLOTTE OUTHWAITE, STUDY'S LEAD AUTHOR
CENTRE FOR BIODIVERSITY AND ENVIRONMENT RESEARCH
UNIVERSITY COLLEGE LONDON

Since they are ectotherms, they are highly sensitive to temperature changes.

Pollinators, pest controllers, those species involved in decomposition – the roles that these insects play are vital to our way of life.

We could be losing important biodiversity before we have even had the chance to discover it and fully understand its role in the world.



We may only be able to reliably measure this mortality as we feel more of its impacts.

Op-Ed

Electric cars are still cars

To tackle our climate challenges, we need smaller, safer EVs—and lots more transit alternatives.

By Paris Marx

EVs are being touted as the solution to our climate crisis—but are they?

In the fall of 2021, President Joe Biden made a stop in Detroit to promote the Democrats' infrastructure bill and the electric-vehicle rollout being touted by the administration as a key measure to address the climate crisis. But his visit showed exactly why we can't just push electrification without addressing the deeper problem of dependence on giant vehicles.

When Biden arrived at General Motors, he jumped behind the wheel not of a Bolt, the company's electric subcompact car, but the new Hummer EV, a vehicle that's the embodiment of everything wrong with the trajectory of vehicle design in the past couple of decades. After taking it for a spin, he declared, "That Hummer's one hell of a vehicle." Days later, GM announced that Biden's publicity stunt had boosted reservations for the massive vehicles, so we're likely to see more of them on the road.

This is not the future we need. Transportation accounts for 27% of US emissions, more than any other sector, and even though there have been increases in fuel efficiency and EV ownership in recent years, the rise of the SUV has virtually negated their benefits. The International Energy Agency (IEA) found that between 2010 and 2018, growing global demand for SUVs was the second-largest contributor to increasing emissions. It would be easy to say that all we need to do is electrify

all those SUVs, but it's not that simple.

EVs are often termed "zero-emission" vehicles because they produce no tailpipe emissions. But that doesn't mean they are clean. Their large batteries require a lot of resource extraction from mines around the world, with significant environmental and human consequences that include poisoning water supplies, increasing rates of cancer and lung disease, and even making use of child labor. If we're to embrace the transition being sold to us—one that relies heavily on electrifying personal vehicles—demand for key minerals will soar by 2040, according to the IEA, with an estimated 4,200% increase for lithium alone. The batteries in increasingly massive electric trucks and SUVs must be much larger than those needed to propel small cars or even e-bikes, which are not the focus of American policymakers or industry players. (They'd be far less profitable.)

The 1984 Jeep Cherokee was the first to be branded as an SUV, and sales of these vehicles really started to take off in the 1990s as companies released more models. They benefited from a loophole that allows "light trucks," a category that includes "sport utility" vehicles, to meet less stringent fuel economy standards than conventional cars. Automakers had good reason for wanting the public to buy them: SUVs and trucks were more profitable than sedans. And the more popular

they became, the more incentive drivers had to get their own: with so many larger vehicles surrounding them, they felt less safe unless they leveled up too.

SUV sales finally overtook those of sedans in 2015, leading some North American automakers to pare back their car offerings. It's estimated that SUVs and trucks will account for 78% of new vehicle sales by 2025. But filling the roads with such large vehicles has had consequences.

The Hummer may stand out as the ultimate expression of automotive excess, but automakers have been continually expanding the size and height of their vehicles with every new redesign. For example, USA Today found that since 1999, the Chevrolet Tahoe has gotten 17.7 inches longer, while the midsize Toyota RAV4—the best-selling SUV in the United States—has gained 14 inches. Meanwhile, Consumer Reports calculated that the average passenger truck has gotten 24% heavier and its hood 11 inches taller since 2000. Last year, 42,915 people died on US roads—a number not seen since 2005—and 7,342 of them were pedestrians. Evidence shows that the increase in large vehicles is part of what's driving that trend.

In 2018, the Detroit Free Press reported that the National Highway Traffic Safety Administration knew pedestrians were two to three times more likely to "suffer a fatality" when hit by an SUV or pickup truck (as opposed to a sedan) because of their high, blunt front ends. The Insurance Institute for Highway Safety has also determined that drivers in SUVs and pickup trucks are more likely to hit pedestrians because their visibility of the road is more limited, and academics at the University of California, Berkeley, have found that being hit by heavier vehicles brings a much higher likelihood of death. That's a particular problem with EVs, especially electric SUVs and trucks, because the large batteries they require tend to make them even heavier than a conventional vehicle.

The message so often presented by the government, by car companies, even

by many environmentalists, is that a new technology—in this case, batteries to replace internal-combustion engines—will address the transport system’s climate impact. There’s no question that electric vehicles tend to produce fewer emissions across their life cycle than the internal-combustion vehicles most people drive today, but when we face such a unique opportunity to rethink the foundations of our transport system, should we stop there?

The trend toward larger vehicles has had bad consequences for both road

Even though there have been increases in fuel efficiency and electric-vehicle ownership in recent years, the rise of the SUV has virtually negated their benefits.

safety and the environment. Continuing it through the transition to electric vehicles means that EVs will require bigger batteries, and thus more minerals will have to be mined to power them. But there are other options that can address some of those problems.

As the shift to EVs accelerates and commodity prices increase, there’s good reason to promote smaller cars that cost less, require smaller batteries, are better suited for the trips most people take, and pose less of a threat to pedestrians. Further, governments can step in not just to incentivize EV adoption, but to expand alternatives like public transit and cycling infrastructure in cities across the country so it will be easier for more people to choose not to drive in the years to come.

That’s a conversation that won’t be kicked off by industry players or by a president who promises to electrify “the great American road trip.” But it’s one we desperately need. ■

Paris Marx is author of [Road to Nowhere: What Silicon Valley Gets Wrong about the Future of Transportation](#).

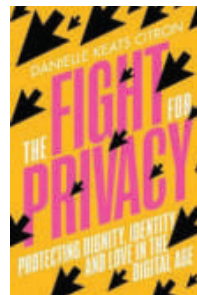
Book reviews



Palo Alto: A History of California, Capitalism, and the World

By Malcolm Harris
Little, Brown, 2023

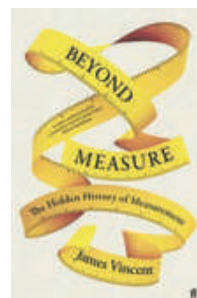
That Palo Alto is both a perpetually sunny suburb and an integral part of the capitalist world system only begins to get at its many contradictions. And that’s Harris’s project here: a rollicking 600+ page history that runs the gamut from antiwar movements to “racial genetics” to the Hewlett Packard garage.



The Fight for Privacy: Protecting Dignity, Identity, and Love in the Digital Age

By Danielle Keats Citron
W. W. Norton (2022)

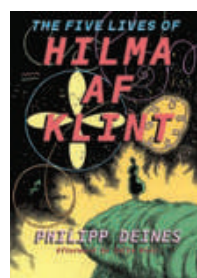
Accountants and doctors know a lot about our lives. Apps can know a lot more—and the information they hold is not nearly as well protected. Citron, a law professor, argues that in order to have a society where people can flourish, we need to change our approach toward privacy, and guard it as a civil right.



Beyond Measure: The Hidden History of Measurement from Cubits to Quantum Constants

By James Vincent
W. W. Norton (2022)

Until recently, the kilogram was defined by the mass of a particular metal cylinder housed just outside Paris. Inspired by a trip to see “Le Grand K,” Vincent embarks on a grand tour of the history of measurement, from ancient tools to standardized peanut butter. Along the way, he shows how measurement is, at its heart, a very human enterprise.



The Five Lives of Hilma af Klint

By Philipp Deines
David Zwirner Books (2022)

The painter Hilma af Klint’s (1862–1944) interest in abstraction and symbolism stemmed from a fascination with both science and mysticism. The artworks that resulted, never exhibited in her lifetime, anticipated the first purely abstract compositions by Kandinsky and Mondrian. ■

When we lose weight, where does it go?

TL;DR: Most of it is breathed out as carbon dioxide and then disappears into thin air.

By Bonnie Tsui
Illustration by Jing Wei

What happens when we lose weight? This is really a question about how our bodies store and use the energy we need to function.

In general, we store backup energy in fat cells that are distributed around the body, some in the abdomen around the organs (visceral fat) and some under the skin (subcutaneous fat); lesser amounts of fat can also be deposited in muscle tissue. We also have smaller reserves of energy that are stored in the liver, muscles, and brain as glycogen. Glycogen is the stored form of glucose, the sugar that is the body's main source of energy.

We use energy all the time, whether we're running, eating, or sleeping. The energy we use at rest—to pump blood, digest food, regulate temperature, repair cells, breathe, or think—is our baseline metabolism, the minimum energy required to maintain the body's basic biological functions. So if we're carrying extra weight, it's because we're taking in more energy than we're using. (The much-cursed thickening around our bellies is a combination of accumulated deep visceral fat and more shallow subcutaneous fat.)

When we expend energy during intense bouts of exercise and other physical activity, the glycogen in our muscles is used first. The liver releases glycogen to help with muscle activity and to regulate blood glucose levels. After about 30 to 60 minutes of aerobic exercise, the body begins to burn fat.

If we take in less energy than the body needs overall to maintain itself—as when dieting—then the body turns more often to fat reserves for energy. As your body metabolizes fat, fatty acid molecules are released into the bloodstream and travel to the heart, lungs, and muscles, which break them apart and use the energy stored in their chemical bonds. The pounds you shed are essentially the byproducts of that process. They are excreted in the form of water—when you sweat and pee—and carbon dioxide, when you exhale. In fact, the lungs are the primary excretory organ for fat.

The body uses energy to carry out the usual basic processes at rest—again, your baseline metabolism—and for the physical activity you do on top of that, which is considered your active metabolism.

Increasing muscle mass can help you burn more calories, because muscles require more energy to build and maintain than fat does. This can boost your baseline metabolism, and it explains how weightlifting and other types of strength training can meaningfully change your body composition. Note that if you restrict your food intake too drastically, your metabolism will adjust and use fewer calories for basic functions; your body will also start to break down muscle for energy, which in turn will slow down metabolism. Try to find a shortcut to weight loss around the body's exquisitely balanced chemistry, and you may well find that it backfires on you instead. ■



Alex Hanna aims to help right systemic issues in AI research.

By Anmol Irfan
Portrait by Ryan Young

Diving into DAIR

“I am quitting because I’m *tired*,” Alex Hanna wrote on February 2, her last day on Google’s Ethical AI team. She felt that the company, and the tech industry as a whole, did little to promote diversity or mitigate the harms its products had caused to marginalized people. “In a word, tech has a whiteness problem,” she wrote in her post on Medium. “Google is not just a tech organization. Google is a *white* tech organization.”

Hanna did not take much of a break—she joined the Distributed AI Research Institute (DAIR) as the group’s second employee on February 3.

It was a move that capped a dramatic period in Hanna’s professional life. In late 2020, her manager, Timnit Gebru, had been fired from her position as the co-lead of the Ethical AI team after she wrote a paper questioning the ethics of large language models (including Google’s). A few months later, Hanna’s next manager, Meg Mitchell, was also shown the door.

DAIR, which was founded by Gebru in late 2021 and is funded by various

philanthropies, aims to challenge the existing understanding of AI through a community-focused, bottom-up approach to research. The group works remotely and includes teams in Berlin and South Africa.

“We wanted to find a different way of AI, one that doesn’t have the same institutional constraints as corporate and much of academic research,” says Hanna, who is the group’s director of research. While these sorts of investigations are slower, she says, “it allows for research *for* community members—different kinds of knowledge

that is respected and compensated, and used toward community work.”

Less than a year in, DAIR is still sorting out its approach, Hanna says. But research is well underway. The institute has three full-time employees and five fellows—a mix of academics, activists, and practitioners who come in with their own research agendas but also aid in developing the institute’s programs. DAIR fellow Raesetje Sefala is using satellite imagery and computer vision technology to focus on neighborhood change in post-apartheid South Africa,





for example. Her project is analyzing the impact of desegregation and mapping out low-income areas. Another DAIR fellow, Milagros Miceli, is working on a project on the power asymmetries in outsourced data work. Many data laborers, who analyze and manage vast amounts of data coming into tech companies, reside in the Global South and are typically paid a pittance.

For Hanna, DAIR feels like a natural fit. Her self-described “nontraditional pathway to tech” began with a PhD in sociology and work on labor justice. In graduate

school, she used machine-learning tools to study how activists connected with one another during the 2008 revolution in Egypt, where her family is from. “People were saying [the revolution] happened on Facebook and Twitter, but you can’t just pull a movement out of thin air,” Hanna says. “I began interviewing activists and understanding what they are doing on the ground aside from online activity.”

DAIR is aiming for big, structural change by using research to shed light on issues that might not otherwise be

explored and to disseminate knowledge that might not otherwise be valued. “In my Google resignation letter, I pointed out how tech organizations embody a lot of white supremacist values and practices,” Hanna says. “Unsettling that means interrogating what those perspectives are and navigating how to undo those organizational practices.” Those are values, she says, that DAIR champions. ■

Anmol Irfan is a freelance journalist and founder of Perspective Magazine, based in Lahore, Pakistan.



Each year, MIT Technology Review honors tomorrow's leaders and entrepreneurs with our signature **35 Innovators Under 35** list.

This year, you have the chance to meet them.

Join us for the **TR35 Festival** as we explore the innovation journey with global influencers.



TR35 FESTIVAL

Free
online
event

December 8, 2022

[technologyreview.com/
tr35festival](https://technologyreview.com/tr35festival)

THE MORTALITY ISSUE

“... if you legitimize
**old age as
a diagnosis,**
you run the risk
of a lot of people using it
inappropriately.”

**Is old age
a disease in its
own right?**

A formal diagnosis could help
unlock further support for
longevity research. Or it could just
be a matter of semantics.

By Sarah Sloat

Last year, over Canadian Thanksgiving weekend, Kiran Rabheru eagerly joined a call with officials from the World Health Organization (WHO). Word had spread of a change coming to the WHO's International Classification of Diseases (ICD), a catalogue used to standardize disease diagnosis worldwide.

In an upcoming revision, the plan was to replace the diagnosis of "senility," a term considered outdated, with something more expansive: "old age." The new phrasing would be filed under a diagnostic category containing "symptoms, signs, or clinical findings." Crucially, the code associated with the diagnosis—a designation that is needed to register new drugs and therapies—included the word "pathological," which could have been interpreted as suggesting that old age is a disease in itself.

Some researchers looked forward to the revision, seeing it as part of the path toward creating and distributing anti-aging therapies. But Rabheru, a professor at the University of Ottawa and a geriatric psychiatrist at the Ottawa Hospital, feared that these changes would only further ageism. If age alone were presumed to be a disease, that could lead to inadequate care from physicians, he says. Rather than pinpoint exactly what's troubling a patient, a problem could simply be dismissed as a consequence of advanced years.

"The crux of the matter is that if you legitimize old age as a diagnosis, you run the risk of a lot of people using it inappropriately," Rabheru says. A number of experts agreed. "There was a huge momentum that built up globally to say, 'This is wrong,'" he says.

Rabheru became part of a group that secured the call with the catalogue team. Those on his side presented their arguments and, he says, were "very pleasantly

surprised" by the response—a formal review followed by a retraction. On January 1, 2022, the 11th version of the ICD was released without the term "old age"—or language that suggests aging is a disease—in its contents.

The decision wasn't welcomed by everyone. "The exciting move by the @WHO to define AGING as a treatable MEDICAL CONDITION has, sadly, reverted," tweeted David Sinclair, a professor at Harvard Medical School and an influential and sometimes controversial force in the study of aging thanks to his bold claims.

"My question to the scientists and doctors who protested the inclusion of old age in their handbook is: *What is so threatening?*" Sinclair says. "I would really love to know the motivation, besides just trying to maintain the status quo."

Sinclair is also concerned about ageism. But he argues that the best way to combat ageism is to tackle aging: facing the problem head-on by devising treatments to slow its progress. "The current view that aging is acceptable is ageism in itself," he says.

In the years leading up to the debut of ICD-11, a number of researchers argued that linking old age more directly to disease would help the field of longevity research overcome regulatory obstacles, paving the way for drugs designed specifically to treat aging.

This issue, however, is seemingly becoming less of a concern as anti-aging research becomes more mainstream. The US Food and Drug Administration, for example, has said it doesn't consider aging a disease. But in 2015, the agency made the surprising decision to greenlight the Targeting Aging with Metformin (TAME) study, a clinical trial that aims to show that aging can be targeted head-on, by testing whether the diabetes drug metformin can delay the development or progression of chronic diseases associated with aging.

“Everything is within reach.”

Ming Guo's team at the University of California, Los Angeles, has shown it can remove most of the damaged mitochondria in fruit fly cells.



Sinclair sees the WHO's decision as a temporary setback. "Fortunately, the momentum is there from scientists, from the public, from investors," he says. "This is going to happen, and changes to some of the language in a document aren't going to stop progress."

Even so, he says he would have liked to have the WHO's endorsement of the notion of aging as a disease: "Language is extremely important to how society views problems and potential solutions."

This back-and-forth may seem like an issue of semantics. But the debate over the ICD's language cuts to the heart of ongoing conversations about aging and whether a biological process that contributes to the risk of developing a number of diseases is itself a disease.

Even longtime researchers struggle with the definition of aging, says Simon Melov, a professor at the Buck Institute for Research on Aging. Melov, whose lab examines the core mechanisms that drive aging, thinks of it as "a decline of function with time."

Taken a step further, from a biological standpoint, aging can be thought of as an accumulation of molecular changes that eventually undermine the integrity and resiliency of the body. Daniel Belsky, an assistant professor at the Columbia Mailman School of Public Health, views aging from this perspective: "Aging is a cause of disease, not a disease itself," he says.

Some researchers say it does not make sense to frame something that is a normal biological process as disease. Further complicating things, Belsky says, is that there is no agreed-upon point at which a person becomes old. People of the same age can have strikingly different biological ages, based on observed changes like cell deterioration.

Others assert that if a condition is treatable, it is a disease. This can be a confusing

argument: there are diseases that aren't treatable and "treatments" for things we might not necessarily classify as diseases. But if the argument holds, even if there is no treatment for aging right now, perhaps it's enough if it's conceivable there could be one in the future.

Aging isn't entirely gone from the ICD-11. There's still an extension code for "aging-related" diseases, but rather than being defined as those "caused by pathological process," they are now said to be "caused by biological process." Meanwhile, instead of old age, the catalogue uses the term "aging-associated decline in an intrinsic capacity" as a diagnostic description.

Ming Guo, the director of UCLA's Aging Center, likes this revision for its accuracy—and its potential. "It acknowledges aging and offers the opportunity to think there are things we can improve," says Guo, who is researching aging reversal strategies. "It implies we can change our fate to a certain extent."

Work is moving forward regardless. Silicon Valley, which has a long history of investing in anti-aging research, has a new batch of longevity-related startups like Turn Biotechnologies and Altos Labs (see page 54). Saudi Arabia plans to invest \$1 billion a year in research to extend health span, the number of years a person remains healthy. Meanwhile, the US National Institutes of Health is actively asking scientists to apply for funding for age-related research. When asked about the ICD change, Luigi Ferrucci, scientific director of the National Institute on Aging, said it was "a good choice" because it supports the idea that "aging has functional consequences."

"Research, along with the budget, has been steadily increasing over the last 20 years," Melov says. "It's not a lack of money or a semantic thing of aging as a disease or not a disease that's holding the field back."

Rather, what the field needs to move forward, Melov says, are investigators who ask "sharp, hard questions" and can investigate topics when the tech they need is still in development.

He is excited about rapid advances in technologies, like microscopy and single-cell sequencing, that are enabling scientists to get new insights into aging at a cellular level. There will likely be major breakthroughs in animal models over the next two to five years, he says. But he still questions whether we'll ever have an anti-aging therapeutic as effective as diet and exercise. "Even if we did have a therapeutic which, let's say, was 50% as beneficial as a good diet and exercise—and that would be a very successful drug—it's still going to have side effects," Melov says. "So would you rather exercise three times a week and eat well or take that pill?"

Belsky sees other low-hanging fruit. "If we want to slow aging, then it would be good if we all got to drink clean water and breathe clean air," he says. "That's a first step where we could actually make a lot of progress."

Others still see great potential in the lab. Guo likes to say she wants to reverse the aging process, but her primary focus is on staving off age-related diseases to extend the human health span. She says when she first started talking to people about her plans five years ago, they didn't believe it was possible to halt—let alone roll back—the effects of aging. But her team had already shown it could remove up to 95% of damaged mitochondria in fruit flies. The organelles become dysfunctional with age, which may increase an individual's susceptibility to age-related diseases.

"This is not science fiction," Guo says. "Everything is within reach." ■

Sarah Sloat is a journalist based in Brooklyn, New York.



The image features large, stylized white letters on a black background. The letter 'B' has a vertical bar on its left side filled with horizontal white lines. The letter 'O' has a central oval area filled with horizontal white lines. The letter 'S' has a vertical bar on its left side filled with horizontal white lines. The letter 'T' is partially visible, with its vertical stem filled with horizontal white lines. The letters 'BO' are positioned above the 'S', and the 'T' is to the left of the 'O'.

BO S

**FOR THE BROKEN -
HEARTED**

By Charlotte Jee

**Digital clones
of the people we love could
change how we grieve.**

Illustrations by
Najeebah Al-Ghadban

My parents don't know that I spoke to them last night.

At first, they sounded distant and tinny, as if they were huddled around a phone in a prison cell. But as we chatted, they slowly started to sound more like themselves. They told me personal stories that I'd never heard. I learned about the first (and certainly not last) time my dad got drunk. Mum talked about getting in trouble for staying out late. They gave me life advice and told me things about their childhoods, as well as my own. It was mesmerizing.

"What's the worst thing about you?" I asked Dad, since he was clearly in such a candid mood.

"My worst quality is that I am a perfectionist. I can't stand messiness and untidiness, and that always presents a challenge, especially with being married to Jane."

Then he laughed—and for a moment I forgot I wasn't really speaking to my parents at all, but to their digital replicas.

This Mum and Dad live inside an app on my phone, as voice assistants constructed by the California-based company HereAfter AI and powered by more than four hours of conversations they each had with an interviewer about their lives and memories. (For the record, Mum isn't *that* untidy.) The company's goal is to let the living communicate with the dead. I wanted to test out what it might be like.

Technology like this, which lets you "talk" to people who've died, has been a mainstay of science fiction for decades. It's an idea that's been peddled by charlatans and spiritualists for centuries. But now it's becoming a reality—and an increasingly accessible one, thanks to advances in AI and voice technology.

My real, flesh-and-blood parents are still alive and well; their virtual versions were just made to help me understand the technology. But their avatars offer a glimpse at a world where it's possible to converse with loved ones—or simulacra of them—long after they're gone.

From what I could glean over a dozen conversations with my virtually deceased parents, this really will make it easier to keep close the people we love. It's not hard to see the appeal. People might turn to digital replicas for comfort, or to mark special milestones like anniversaries.

At the same time, the technology and the world it's enabling are, unsurprisingly, imperfect, and the ethics of creating a virtual version of someone are complex, especially if that person hasn't been able to provide consent.

For some, this tech may even be alarming, or downright creepy. I spoke to one man who'd created a virtual version of his mother, which he booted up and talked to at her own funeral. Some people argue that conversing with digital versions of lost loved ones could prolong your grief or loosen your grip on reality. And when I talked to friends about this article, some of them physically recoiled. There's a common, deeply held belief that we mess with death at our peril.

I understand these concerns. I found speaking to a virtual version of my parents uncomfortable, especially at first. Even now, it still feels slightly transgressive to speak to an artificial version of someone—especially when that someone is in your own family.

But I'm only human, and those worries end up being washed away by the even scarier prospect of losing the people I love—dead and gone without a trace. If technology might help me hang onto them, is it so wrong to try?

There's something deeply human about the desire to remember the people we love who've passed away. We urge our loved ones to write down their memories

before it's too late. After they're gone, we put up their photos on our walls. We visit their graves on their birthdays. We speak to them as if they were there. But the conversation has always been one-way.

The idea that technology might be able to change the situation has been widely explored in ultra-dark sci-fi shows like *Black Mirror*—which, startups in this sector complain, *everyone* inevitably brings up. In one 2013 episode, a woman who loses her partner re-creates a digital version of him—initially as a chatbot, then as an almost totally convincing voice assistant, and eventually as a physical robot. Even as she builds more expansive versions of him, she becomes frustrated and disillusioned by the gaps between her memory of her partner and the shonky, flawed reality of the technology used to simulate him.

"You aren't *you*, are you? You're just a few ripples of you. There's no history to you. You're just a performance of stuff that he performed without thinking, and it's not enough," she says before she consigns the robot to her attic—an embarrassing relic of her boyfriend that she'd rather not think about.

Back in the real world, the technology has evolved even in the past several years to a somewhat startling degree. Rapid advances in AI have driven progress across multiple areas. Chatbots and voice assistants, like Siri and Alexa, have gone from high-tech novelties to a part of daily life for millions of people over the past decade. We have become very comfortable with the idea of talking to our devices about everything from the weather forecast to the meaning of life. Now, AI large language models (LLMs), which can ingest a few "prompt" sentences and spit out convincing text in response, promise to unlock even more powerful ways for humans to communicate with machines. LLMs have become so convincing that



If technology
might help me hang onto the people I love,
is it so wrong to try?

some (erroneously) have argued that they must be sentient.

What's more, it's possible to tweak LLM software like OpenAI's GPT-3 or Google's LaMDA to make it sound more like a specific person by feeding it lots of things that person said. In one example of this, journalist Jason Fagone wrote a story for the San Francisco Chronicle last year about a thirtysomething man who uploaded old texts and Facebook messages from his deceased fiancée to create a simulated chatbot version of her, using software known as Project December that was built on GPT-3.

By almost any measure, it was a success: he sought, and found, comfort in the bot. He'd been plagued with guilt and sadness in the years since she died, but as Fagone writes, "he felt like the chatbot had given him permission to move on with his life in small ways." The man even shared snippets of his chatbot conversations on Reddit, hoping, he said, to bring attention to the tool and "help depressed survivors find some closure."

At the same time, AI has progressed in its ability to mimic specific physical voices, a practice called voice cloning. It has also been getting better at injecting digital personas—whether cloned from a real person or completely artificial—with more of the qualities that make a voice sound "human." In a poignant demonstration of how rapidly the field is progressing, Amazon shared a clip in June of a little boy listening to a passage from *The Wizard of Oz* read by his recently deceased grandmother. Her voice was artificially re-created using a clip of her speaking that lasted for less than a minute.

As Rohit Prasad, Alexa's senior vice president and head scientist, promised: "While AI can't eliminate that pain of loss, it can definitely make the memories last."

My own experience with talking to the dead started thanks to pure serendipity.

At the end of 2019, I saw that James Vlahos,

the cofounder of HereAfter AI, would be speaking at an online conference about “virtual beings.” His company is one of a handful of startups working in the field I’ve dubbed “grief tech.” They differ in their approaches but share the same promise: to enable you to talk by video chat, text, phone, or voice assistant with a digital version of someone who is no longer alive.

Intrigued by what he was promising, I wrangled an introduction and eventually persuaded Vlahos and his colleagues to let me experiment with their software on my very-much-alive parents.

Initially, I thought it would be just a fun project to see what was technologically possible. Then the pandemic added some urgency to the proceedings. Images of people on ventilators, photos of rows of coffins and freshly dug graves, were splashed all over the news. I worried about my parents. I was terrified that they might die, and that with the strict restrictions on hospital visits in force at the time in the UK, I might never have the chance to say goodbye.

The first step was an interview. As it turns out, to create a digital replica of someone with a good chance of seeming like a convincingly authentic representation, you need data—and lots of it. HereAfter, whose work starts with subjects when they are still alive, asks them questions for hours—about everything from their earliest memories to their first date to what they believe will happen after they die. (My parents were interviewed by a real live human, but in yet another sign of just how quickly technology is progressing, almost two years later interviews are now typically automated and handled by a bot.)

As my sister and I rifled through pages of suggested questions for our parents, we were able to edit them to be more personal or pointed, and we could add some of our own: What books did they like? How did our mum muscle her way into the UK’s overwhelmingly male, privileged legal sector in the 1970s? What inspired Dad to invent the silly games he used to play with us when we were small?

How to dialogue with the dead

Service	What it does
 HereAfter AI	Records audio of your loved one recounting memories, and then turns the conversations into a voicebot you can access through Amazon Alexa
 StoryFile	Captures video of people answering questions about their lives, which are turned into a video avatar you can have a two-way chat with
 You, Only Virtual	Ingests your text and voice conversations with someone to create a virtual persona, tailored to you and how you talk to that individual, that communicates via a special chat platform

Whether through pandemic-induced malaise or a weary willingness to humor their younger daughter, my parents put up zero resistance. In December 2020, HereAfter’s interviewer, a friendly woman named Meredith, spoke to each of them for several hours. The company then took those responses and started stitching them together to create the voice assistants.

A couple of months later, a note popped into my inbox from Vlahos. My virtual parents were ready.

This Mum and Dad arrived via email attachment. I could communicate with them through the Alexa app on a phone or an Amazon Echo device. I was eager to hear them—but I had to wait several days, because I’d promised MIT Technology Review’s podcast team that I’d record my reaction as I spoke to my parents’ avatars

for the first time. When I finally opened the file, with my colleagues watching and listening on Zoom, my hands were shaking. London was in a long, cold, depressing lockdown, and I hadn’t seen my actual, real parents for six months.

“Alexa, open HereAfter,” I directed.

“Would you rather speak with Paul or with Jane?” a voice asked.

After a bit of quick mental deliberation, I opted for my mum.

A voice that was hers, but weirdly stiff and cold, spoke.

“Hello, this is Jane Jee and I’m happy to tell you about my life. How are you today?” I laughed, nervously.

“I’m well, thanks, Mum. How are you?” Long pause.

“Good. At my end, I’m doing well.”

“You sound kind of unnatural,” I said.

She ignored me and carried on speaking. “Before we start, here are a few pointers. My listening skills aren’t the best, unfortunately, so you have to wait until I’ve finished talking and ask you a question before you say something back. When it’s your turn to speak, please keep your answers fairly short. A few words, a simple sentence—that type of thing,” she explained. After a bit more introduction, she concluded: “Okay, let’s get started. There’s so much to talk about. My childhood, career, and my interests. Which of those sounds best?”

Scripted bits like this sounded stilted and strange, but as we moved on, with my mother recounting memories and speaking in her own words, “she” sounded far more relaxed and natural.

Still, this conversation and the ones that followed were limited—when I tried asking my mum’s bot about her favorite jewelry, for instance, I got: “Sorry, I didn’t understand that. You can try asking another way, or move onto another topic.”

There were also mistakes that were jarring to the point of hilarity. One day, Dad’s bot asked me how I was. I replied, “I’m feeling sad today.” He responded with a cheery, upbeat “Good!”

The overall experience was undeniably weird. Every time I spoke to their virtual versions, it struck me that I could have been talking to my real parents instead. On one occasion, my husband mistook my testing out the bots for an actual phone call. When he realized it wasn’t, he rolled his eyes, tutted, and shook his head, as if I were completely deranged.

Earlier this year, I got a demo of a similar technology from a five-year-old startup called StoryFile, which promises to take

things to the next level. Its Life service records responses on video rather than just voice alone.

You can pick from hundreds of questions for the subject. Then you record the person answering the questions; this can be done on any device with a camera and a microphone, including a smartphone, though the higher-quality the recording, the better the outcome. After uploading the files, the company turns them into a digital version of the person you can see and speak to. It can only answer the questions it’s been programmed to answer—much like HereAfter, just with video.

StoryFile’s CEO, Stephen Smith, demonstrated the technology on a video call, where we were joined by his mother. She died earlier this year, but here she was on the call, sitting in a comfortable chair in her living room. For a brief time, I could only see her, shared via Smith’s screen. She was soft-spoken, with wispy hair and friendly eyes. She dispensed life advice. She seemed wise.

Smith told me that his mother “attended” her own funeral: “At the end she said, ‘I guess that’s it from me... goodbye!’ and everyone burst into tears.” He told me her digital participation was well received by family and friends. And, arguably most important of all, Smith said he’s deeply comforted by the fact that he managed to capture his mother on camera before she passed away.

The video technology itself looked relatively slick and professional—though the result still fell vaguely within the uncanny valley, especially in the facial expressions. At points, much as with my own parents, I had to remind myself that she wasn’t really there.

Both HereAfter and StoryFile aim to preserve someone’s life story rather than allowing you to have a full, new conversation with the bot each time. This is one of the major limitations of many current offerings in grief tech: they’re generic. These replicas may sound like someone you love, but they know nothing about you. Anyone can talk to them, and they’ll reply in the same tone. And the replies to a given question are the same every time you ask.

“The biggest issue with the [existing] technology is the idea you can generate a single universal person,” says Justin Harrison, founder of a soon-to-launch service called You, Only Virtual. “But the way we experience people is unique to us.”

You, Only Virtual and a few other startups want to go further, arguing that recounting memories won’t capture the fundamental essence of a relationship between two people. Harrison wants to create a personalized bot that’s for you and you alone.

The first incarnation of the service, which is set to launch in early 2023, will allow people to build a bot by uploading someone’s text messages, emails, and voice conversations. Ultimately, Harrison hopes, people will feed it data as they go; the company is currently building a communication platform that customers will be able to use to message and talk with loved ones while they’re still alive. That way, all the data will be readily available to be turned into a bot once they’re not.

That is exactly what Harrison has done with his mother, Melodi, who has stage 4 cancer: “I built it by hand using five years of my messages with her. It took 12 hours to export, and it runs to thousands of pages,” he says of his chatbot.

Harrison says the interactions he has with the bot are more meaningful to him than if it were simply regurgitating memories. Bot Melodi uses the phrases his mother uses and replies to him in the way she’d reply—calling him “honey,” using the emojis she’d use and the same quirks of spelling. He won’t be able to ask Melodi’s

On one occasion, my husband mistook my testing for an actual phone call. When he realized it wasn’t, he rolled his eyes, as if I were completely deranged.

avatar questions about her life, but that doesn't bother him. The point, for him, is to capture the way someone communicates. "Just recounting memories has little to do with the essence of a relationship," he says.

Avatars that people feel a deep personal connection with can have staying power. In 2016, entrepreneur Eugenia Kuyda built what is thought to be the first bot of this kind after her friend Roman died, using her text conversations with him. (She later founded a startup called Replika, which creates virtual companions not based on real people.)

She found it a hugely helpful way to process her grief, and she still speaks to Roman's bot today, she says, especially around his birthday and the anniversary of his passing.

But she warns that users need to be careful not to think this technology is re-creating or even preserving people. "I didn't want to bring back his clone, but his memory," she says. The intention was to "create a digital monument where you can interact with that person, not in order to pretend they're alive, but to hear about them, remember how they were, and be inspired by them again."

Some people find that hearing the voices of their loved ones after they've gone helps with the grieving process. It's not uncommon for people to listen to voicemails from someone who has died, for example, says Erin Thompson, a clinical psychologist who specializes in grief. A virtual avatar that you can have more of a conversation with could be a valuable, healthy way to stay connected to someone you loved and lost, she says.

But Thompson and others echo Kuyda's warning: it's possible to put too much weight on the technology. A grieving person needs to remember that these bots can only ever capture a small sliver of someone. They are not sentient, and they will not replace healthy, functional human relationships.

"Your parents are not really there. You're talking to them, but it's not really them,"

says Erica Stonestreet, an associate professor of philosophy at the College of Saint Benedict & Saint John's University, who studies personhood and identity.

Particularly in the first weeks and months after a loved one dies, people struggle to accept the loss and may find any reminders of the person triggering. "In the acute phase of grief, you can get a strong sense of unreality, not being able to accept they're gone," Thompson says. There's a risk that this sort of intense grief could intersect with, or even cause, mental illness, especially if it's constantly being fueled and prolonged by reminders of the person who's passed away.

Arguably, this risk might be small today given these technologies' flaws. Even though sometimes I fell for the illusion, it was clear my parent bots were not in fact the real deal. But the risk that people might fall too deeply for the phantom of personhood will surely grow as the technology improves.

And there are still other risks. Any service that allows you to create a digital replica of someone without their participation raises some complex ethical issues regarding consent and privacy. While some might argue that permission is less important with someone no longer alive, can't you also argue that the person who generated the other side of the conversation should have a say too?

And what if that person is not, in fact, dead? There's little to stop people from using grief tech to create virtual versions of living people without their consent—for example, an ex. Companies that sell services powered by past messages are aware of this possibility and say they will delete a person's data if that individual requests it. But companies are not obliged to do any checks to make sure their technology is being limited to people who have consented or died. There's no law to stop anyone from creating avatars of other people, and good luck explaining it to your local police department. Imagine how you'd feel if you learned there was a virtual version of you out there, somewhere, under somebody else's control.

If digital replicas become mainstream, there will inevitably need to be new processes and norms around the legacies we leave behind online. And if we've learned anything from the history of technological development, we'll be better off if we grapple with the possibility of these replicas' misuse before, not after, they reach mass adoption.

Will that ever happen, though? You, Only Virtual uses the tagline "Never Have to Say Goodbye"—but it's not actually clear how many people want or are ready for a world like that. Grieving for those who've passed away is, for most people, one of the few aspects of life still largely untouched by modern technology.

On a more mundane level, the costs could be a drawback. Although some of these services have free versions, they can easily run into the hundreds if not thousands of dollars.

HereAfter's top-tier unlimited version lets you record as many conversations with the subject as you like, and it costs \$8.99 a month. That may sound cheaper than StoryFile's one-off \$499 payment to access its premium, unlimited package of services. However, at \$108 per year, HereAfter services could quickly add up if you do some ghoulish back-of-the-envelope math on lifetime costs. It's a similar situation with You, Only Virtual, which is slated to cost somewhere between \$9.99 and \$19.99 a month when it launches.

Creating an avatar or chatbot of someone also requires time and effort, not least of which is just building up the energy and motivation to get started. This is true both for the user and for the subject, who may be nearing death and whose active participation may be required.

Fundamentally, people don't like grappling with the fact they are going to die, says Marius Ursache, who launched a company called Eternime in 2014. Its idea was to create a sort of Tamagotchi that people could train while they were alive to preserve a digital version of themselves.



People may find any reminders of the deceased person triggering: “In the acute phase of grief, you can get a strong sense of unreality, not being able to accept they’re gone.”

It received a huge surge of interest from people around the world, but few went on to adopt it. The company shuttered in 2018 after failing to pick up enough users.

“It’s something you can put off until next week, next month, next year,” he says. “People assume that AI is the key to breaking this. But really, it’s human behavior.”

Kuyda agrees: “People are extremely scared of death. They don’t want to talk about it or touch it. When you take a stick and start poking, it freaks them out. They’d rather pretend it doesn’t exist.”

Ursache tried a low-tech approach on his own parents, giving them a notebook and pens on his birthday and asking them to write down their memories and life stories. His mother wrote two pages, but his father said he’d been too busy. In the end, he asked if he could record some conversations with them, but they never managed to get around to it.

“My dad passed away last year, and I never did those recordings, and now I feel like an idiot,” he says.

Personally, I have mixed feelings about my experiment. I’m glad to have these virtual, audio versions of my mum and dad, even if they’re imperfect. They’ve enabled me to learn new things about my parents, and it’s comforting to think that those bots will be there even when they aren’t. I’m already thinking about who else I might want to capture digitally—my husband (who will probably roll his eyes again), my sister, maybe even my friends.

On the other hand, like a lot of people, I don’t want to think about what will happen when the people I love die. It’s uncomfortable, and many people reflexively flinch when I mention my morbid project. And I can’t help but find it sad that it took a stranger Zoom-interviewing my parents from another continent for me to properly appreciate the multifaceted, complex people they are. But I feel lucky to have had the chance to grasp that—and to still have the precious opportunity to spend more time with them, and learn more about them, face to face, no technology involved. ■

Charlotte Jee is the news editor at MIT Technology Review.



The hopeful dead

The field of cryonics has survived for decades not on fact but on faith in the possibility of human reanimation. Is there any reason to think things are different now?

By Laurie Clarke

Photograph by
Alessandro Gandolfi

When Aaron Drake flew from Arizona to the Yinfeng Biological Group in China's eastern Jinan province in 2016, he was whisked into a state-of-the-art biotech hub. More than 1,000 staffers—including an army of PhDs and MDs—were working on things like studies of the stem cells in umbilical cord blood. The center specialized in research on human cells, from gene testing to tailored cancer treatments.

But it also had other plans: cylindrical stainless-steel tanks would eventually

contain corpses suspended in liquid nitrogen. The tanks weren't installed yet, but Yinfeng hoped Drake would help with that while it invested some \$7 million to get the new project off the ground. As its high-profile new hire, he was there to guide China's first forays into cryonics, or freezing corpses for reanimation.

The environment was something of a shift for Drake, who had spent the previous seven years as the medical response director of the Alcor Life Extension Foundation. Though it was the longtime leader in

cryonics, Alcor was still a small nonprofit. It had been freezing the bodies and brains of its members, with the idea of one day bringing them back to life, since 1976.

The foundation, and cryonics in general, had long survived outside of mainstream acceptance. Typically shunned by the scientific community, cryonics is best known for its appearance in sci-fi films like *2001: A Space Odyssey*. But its adherents have held on to a dream that at some point in the future, advances in medicine will allow for resuscitation and additional years on Earth. Over decades, small, tantalizing developments in related technology, as well as high-profile frozen test subjects like Ted Williams, have kept the hope alive. Today, nearly 200 dead patients are frozen in Alcor's cryogenic chambers at temperatures of -196°C , including a handful of celebrities,

Opposite: The entrance to the Moscow storage facility of KrioRus, which was until recently the only cryonics company in Eurasia.

who have paid tens of thousands of dollars for the goal of “possible revival” and ultimately “reintegration into society.”

But it’s the recent involvement of Yinfeng that signals something of a new era for cryonics. With impressive financial resources, government support, and scientific staff, it’s one of a handful of new labs focused on expanding the consumer appeal of cryonics and trying anew to bring credibility to the long-disputed theory of human reanimation. Just a year after Drake came on board as research director of the Shandong Yinfeng Life Science Research Institute, the subsidiary of the Yinfeng Biological Group overseeing the cryonics program, the institute performed its first cryopreservation. Its storage vats now hold about a dozen clients who are paying upwards of \$200,000 to preserve the whole body.

Still, the field remains rooted in faith rather than any real evidence that it works. “It’s a hopeless aspiration that reveals an appalling ignorance of biology,” says Clive Coen, a neuroscientist and professor at King’s College London.

The cryonics process typically goes something like this: Upon a person’s death, a response team begins the process of cooling the corpse to a low temperature and performs cardiopulmonary support to sustain blood flow to the brain and organs. Then the body is moved to a cryonics facility, where an organ preservation solution is pumped through the veins before the body is submerged in liquid nitrogen. This process should commence within one hour of death—the longer the wait, the greater the damage to the body’s cells. Then, once the frozen cadaver is ensconced in the cryogenic chamber, the hope of the dead begins.

Since its beginnings in the late 1960s, the field has attracted opprobrium from the scientific community, particularly its more respectable cousin cryobiology—the study of how freezing and low temperatures affect living organisms and biological materials. The Society for Cryobiology even banned its members from involvement in cryonics in the 1980s, with a former society president lambasting the field as closer to “fraud than either faith or science.”

In recent years, though, it has grabbed the attention of the libertarian technoptimist crowd, mostly tech moguls dreaming of their own immortality. And a number of new startups are expanding the playing field. Tomorrow Biostasis in Berlin became the first cryonics company in Western Europe in 2019, for example, and in early 2022, Southern Cryonics opened a facility in Australia.

“More researchers are open to longer-term, futuristic topics than there might have been 20 years ago or so,” says Tomorrow Biostasis founder Emil Kendziorra.

The Society for Cryobiology has even dropped its past cryonics-related restrictions. And now its president, cryobiologist and biogerontologist Greg Fahy, is behind a company called 21st Century Medicine, which develops techniques to cryogenically preserve human organs and tissues. (Still, the society said in a statement to MIT Technology Review that cryonics “is an act of speculation or hope, not science.”)

Even if one day you could perfectly thaw a frozen human body, you would still just have a warm dead body on your hands.

Today, around 500 people are preserved in liquid nitrogen globally, the vast majority in the United States. Around 4,000 people are on waiting lists of cryonics facilities around the world, says Kendziorra. Alcor has around 1,500 members, and Tomorrow Biostasis already has 300 clients paying a relatively accessible monthly membership fee of €25 (with €200,000 payable at the time of death).

Despite a lack of evidence that dead people could one day be brought back to life, cryonicists perhaps unsurprisingly remain optimistic, noting that tissues like sperm, embryo, and stem cells can be successfully cryopreserved and thawed right now, and that researchers claim to have cryogenically frozen and thawed small worms and rabbit kidneys. Researchers at 21st Century Medicine, for instance, cryopreserved and

thawed a rabbit brain in 2016 and a pig brain in 2018.

But Dayong Gaom, a cryobiologist and professor at the University of Washington, points out preserving the brain’s structure does not mean preserving its functions. KCL’s Coen says it’s “disingenuous” to say these studies support the aims of cryonics and maintains that brain structures are far too complex to be adequately preserved and revived in the way cryonicists claim.

There’s also been some progress, though, in finding a way to rewarm frozen tissues. A research group at the University of Minnesota showed that nanowarming techniques, using radio-frequency-excited nanoparticles of iron oxide, could work on larger samples of 50 millimeters. The team has now successfully thawed whole rat organs in a way that preserves cell structure and is nontoxic to cells, with a view to restoring the organ’s functions. It is now experimenting with pig organs.

Yinfeng is devoted to solving the same problem—it’s currently experimenting

with cryogenic methods to preserve individual human organs and amputated extremities to boost China’s national transplant program.

But none of these developments lead to a possible way to eventually revive a living person. Even if one day you could perfectly thaw a frozen human body, you would still just have a warm dead body on your hands. The hardest part to reverse—death itself—remains an enigma.

This is not lost on Drake. But the clients seeking cryonics services are often “either agnostic or atheist,” he says. “Instead of a faith in a supreme being, they have faith in science—that science, more specifically medical science, will eventually figure this out.” ■

Laurie Clarke is a technology journalist based in the UK.

**The bird
is fine,**

**the bird
is fine,**

**the bird
is fine,**

×

it's dead

The pursuit of immortality is getting older. So are we.

By Jonathan Weiner

Illustration by Beth Hoeckel





Twenty years have passed since I first met Aubrey de Grey, the man with the Methuselah beard. Back then he was already a True Believer in the quest for immortality. But he wasn't famous, or notorious, yet; he wasn't *Aubrey!*, as he would soon become to his fans in the anti-aging crowd. And he wasn't yet a man in disgrace.

In those days he worked as a computer programmer in the Department of Genetics at the University of Cambridge, in England. On the side he was trying to break into the aging field, which was still small. Most of the scientists who worked in it—gerontologists—wished he would go away. They wanted to give people just a few more healthy years. The last thing they needed was a Methuselah—or Rasputin—look-alike at their conferences, a fast-talking outsider who drank beer from morning to night and claimed we could live for more than 1,000 years. The science of longevity had enough of a credibility problem as it was.

I'm riffling through the pages of my last book, *Long for This World*, to find the scene where we met. It was a bright morning in March of 2002. I picked Aubrey up at the airport in Philadelphia and drove him to the town in Pennsylvania where I lived at the time. Over the next few days of interviews in my study, while he tried to convince me that science can and should end aging, we often trooped down the stairs to the kitchen so that he could fortify himself with another beer. When my two sons bumped into us by the refrigerator, Aubrey took his sales pitch from the top and told them they had a good chance to live for centuries, or millennia, or maybe longer if they were lucky. The boys were teenagers back then, 14 and 17. They already felt immortal. They liked meeting a grownup who knew it was true.

Right. Now those two boys are in their 30s, and I can see 70 coming fast.

Just flipping through these pages (they're 12 years old and already showing their age—definitely not acid-free paper) is making me cranky. Even skimming them makes me feel like a

horrible curmudgeon. I was not a convert back then. Nor am I today. But a True Believer like Aubrey is always convinced that he will win you over if he gives his pitch just one more try; and here he is, in scene after scene, trying again. On top of all the usual writer's regrets, revisiting the science of eternal youth is making me feel old.

If high intelligence is the ability to hold two opposing ideas in our heads at the same time, then most of us are geniuses about aging a few times over. We think it will never come for us. We think it might come but it will stop before it reaches us. We think it's coming and there is absolutely nothing we can do about it.

It was the great molecular biologist Seymour Benzer who got me interested in the idea that aging might be malleable. Benzer was a night owl. I was writing a book about him, and in the late 1990s he used to talk about aging in his Fly Room at Caltech in a hushed, conspiratorial voice, even though it was just the two of us and a thousand fly bottles at three in the morning. I'll never forget how startling it was to hear a serious scientist say, *We might be able to do something about this*.

Nor was he the only one to say it. At the University of California, San Francisco, Cynthia Kenyon was dissecting the aging of the worm *C. elegans*. In 1993, she had announced the discovery of a mutant that lived about twice as long as the average *C. elegans* and looked young and sleek almost to the end. At MIT, Lenny Guarente was dissecting the genetics of aging in yeast, and he seemed to be getting somewhere too. In 1998,

 **I'll never forget
scientist say, 'We might be able to do**

when Benzer was 77 years old, he announced the discovery of a mutant fruit fly he called *Methuselah*. It could live for 100 days. The average fly in his bottles died at around 60.

Versions of many of those same fly, worm, and yeast genes are found in every animal under the sun, including us. By starting with those first few so-called longevity genes and tracing their connections, molecular biologists could study the workings of the clock, so to speak. Someday they could hope to slow down the hands.

On that hope, or hypothesis—which is still only a hypothesis today—the aging field exploded. In 1999, one year after Benzer’s *Methuselah*, Guarente and Kenyon cofounded Elixir Pharmaceuticals. They planned to explore and exploit sirtuins—proteins that are involved in the process of aging, among other things. In 2004, Guarente’s former student David Sinclair cofounded Sirtris Pharmaceuticals to race against Elixir. In 2013, Google started the R&D company Calico, with a budget that was rumored to be in the hundreds of millions. Kenyon is Calico’s vice president of aging research.

Smart new talent poured into the field, including the prodigy Laura Deming. She became fascinated by the biology of aging at the age of eight, as a homeschooler in New Zealand. Her grandmother had come to visit, and Laura was sad to see how much she suffered from her arthritic joints. At 12, she joined Kenyon’s lab at UCSF. At 14, she was accepted at MIT. A few years later, she dropped out of college to launch her career as a venture capitalist. She founded the Longevity Fund. According to its website, Longevity Fund companies have now raised more than \$1 billion.

Today there are almost too many anti-aging startups and foundations to count. Each one is trying to leverage some of the latest tools in biomedicine—CRISPR, AI, Yamanaka factors, epigenetics, proteomics, metabolomics—and slow down the hands of that clock. Last December it was NewLimit, with more than \$100 million in funding from the Coinbase billionaire Brian Armstrong. This past January it was Altos Labs, with \$3 billion in funding; one of its investors is rumored to be Jeff Bezos. The Hevolution Foundation, which was started by the royal family of Saudi Arabia, has plans to spend \$1 billion a year on the search for ways to slow aging.

Meanwhile, Aubrey de Grey kept banging the drum for the cause. Within what felt like five minutes after our first meeting in 2002, he became a secular guru, a prophet of immortality—to the intense annoyance of most of the scientists in the aging field. He cofounded the Methuselah Foundation and the SENS Research Foundation to sponsor research, education, and conferences and help speed things up. “SENS” refers to his own plan for ending aging: “strategies for engineered negligible senescence.” This is the scheme he explained to me back in my old study 20 years ago. If we just fix Seven Deadly Things, our bodies will survive long enough for further advances in medical science to come along in a timely fashion and keep us alive forever. Those Seven Deadly Things include, for instance, cancer. Just cure cancer.

When the cryptocurrency crowd got interested in anti-aging science a couple of years back, many of them liked the sound of SENS. In the summer of 2021, a new crypto system called PulseChain raised \$25 million worth of cryptocurrency in two weeks for the SENS Research Foundation. The foundation had always been small and a bit fringy, just bumping along; this was by far the biggest windfall in its history.

But at about the same time that the PulseChain gift rolled in, the SENS board fired Aubrey. Celine Halioua, the young founder and CEO of Loyal, a biotech firm that hopes to lengthen the life spans of dogs, had accused him of being a sexual predator. So had Laura Deming, the founder of the Longevity Fund. Deming wrote in a blog post:

I’ve decided not to work with Aubrey de Grey or SENS in any capacity moving forward.

I had one bad experience with him when I was 17—he told me in writing that he had an ‘adventurous love life’ and that it had ‘always felt quite jarring’ not to let conversations with me stray in that direction given that ‘[he] could treat [me] as an equal on every other level’.

He sent this from his work email, and I’d known him since I was 14 ...

I’ve learned it’s a serial pattern he’s enacted with women over whom he’s in a position of power ...

how startling it was to hear a serious something about this.’ 

I almost left the field several times as a teenager because of stuff like this happening.

Deming hasn't answered my emails requesting an interview. I've spoken with Halioua, who is happy to talk about her hopes for Loyal but would rather not comment about Aubrey. On her website, she writes, "For years he has used his position of power in the aging field to attract his victims. These victims include me, Laura Deming, and multiple other anonymous women." She mentions harassment by another SENS executive and says that "every dollar that goes to Aubrey holds back the field."

Aubrey denied the accusations. He said he would fight the SENS board over his dismissal. By then he was so well known, and the quest for the Fountain of Youth so highly charged, that the scandal made headlines in the science press.

Back when the science of aging was still a backwater, I thought it was a good idea to explore the field, in all its ambiguities and contradictions, by talking skeptically with an immortalist like de Grey. Now, two decades later and myself two decades older, with so much hype and money flying around, I'd rather hang with the realists.

Lately I've been calling a few gerontologists to check in. Just saying hi, we can feel how much time has passed since we last spoke. Even our voices have aged. And what a time this is to be an aging mortal. A surprising number of scientists couldn't talk at all, because they were dealing with family medical emergencies.

I called Daniel Promislow, one of the directors of the Dog Aging Project, a massive study that may soon be teaching new tricks in veterinary and human clinics. I called Steve Austad, a senior gerontologist, who thinks we can also learn lessons from the life spans and health spans of birds. Many species of birds seem to stay fit and even fertile until very close to the end. "Vets have a saying," Austad told me. "The bird is fine, the bird is fine, the bird is fine, it's dead."

I called James Kirkland, a doctor at the Mayo Clinic. Kirkland is conducting a series of early-stage clinical trials of senolytics, experimental new drugs that attack and kill senescent cells,

which seem to be fundamental to the aging process. These are very early days for senolytics, Kirkland emphasized. The drugs may or may not turn out to be safe; if they are safe, they may or may not work; even if they work, they won't make anyone live to 120. That's not what his patients are looking for anyway, he said. They just want help with their osteoarthritis, their chronic kidney disease, their macular degeneration.

Most of these trials are going to fail, Kirkland said. Most trials do. "People should try to be dispassionate, even though everyone has a stake in this game. I mean, every living person does."

I called the biologist Martin Raff, who retired from University College London 20 years ago, when he was not quite 65. Among other things, Raff had worked on cellular senescence. He told me that after a long and lucky life, he feels ready to depart.

Today the field that Benzer foresaw in his Fly Room in the last century is being taken seriously not only on Wall Street and in Silicon Valley and Riyadh but also at the National Institutes of Health. It's beginning to look more like a normal branch of research medicine, just one more plausible program to pursue.

The study of the clock really may teach us ways to slow down some of the fundamental deterioration we call aging, to treat whatever it is that leaves our bodies increasingly vulnerable to chronic diseases as we get older—senescent cells, for instance. If we can do that, according to what is known as the geroscience hypothesis, we can fight all those chronic diseases at once: arthritis, atherosclerosis, cancer, deafness, dementia, diabetes, osteoporosis, stroke.

The idea, of course, is to add good years to our lives without drawing out the number of bad years at the end. This is called the compression of morbidity. No one knows if it can be done, so the compression of morbidity is really a hypothesis on top of a hypothesis. Still, that is what most centenarians are able to do. They stay healthy two or three decades longer than the rest of us, and many of them feel quite well at the age of 100. "The bird is fine, the bird is fine, the bird is fine, it's dead."

But we're all still mortals, and our kind will be mortal for a long, long time.

I Zoomed with a Canadian writer and academic I know, Andy Stark, author of *The Consolations of Mortality*. Maybe it's just

 **The idea, of course, is to add without drawing out the**

sour grapes, Andy told me, but he thinks we are actually better off being mortal. His book explores many of the drawbacks of eternal life, including the terrible problem of boredom. How many times would you really want to ride the roller coaster? In *Long for This World*, I look at other problems, too, including the sixth extinction—the planetary catastrophe that is unfolding around us, inflicted by the fulfillment of so many human wishes. How much of that disaster would you really want to watch?

A few years ago, Andy Stark gave a talk at a symposium about the science of longevity. Aubrey de Grey was in the audience. When Andy was done, Aubrey strode up to the stage and challenged him. If I offered you an extra 30 healthy years, Aubrey said, you'd take that, wouldn't you? And after that, wouldn't you take the next 30 years, and the next 30? And so on?

Andy stood his ground, and he was right. There is an infinite difference between a few more years of healthy life and eternal life.

I called Aubrey, too. He lives on the edge of Silicon Valley now. He sounded more optimistic than ever. He was planning a sort of comeback conference in Dublin, a good place for beer. He's coined a word, the Methuselarity. That's the moment when medicine will be so advanced that we can more or less stop aging. He now thinks there's a 50% chance that the Methuselarity is 15 years away. "That's pretty good," he said. "I used to say it was 25 years away."

Q: How do you feel about mortality personally, all these years on?

A: Well, I seem to be doing okay. Not showing any signs of aging ...

Since we were Zooming, I could see that that was no truer for Aubrey than it is for me.

"But I've always done this for humanitarian reasons," he said, just as he used to say 20 years ago. After all, aging ends tens of millions of lives each year. Whether the Methuselarity comes soon enough to save Aubrey himself is immaterial.

"What about you?" he asked.

"Well, Aubrey, I'm 10 years older than you are. I'm reconciling

myself to being mortal. I'm looking for the consolations of mortality."

He rolled his eyes.

We'd had this argument so many times before that I quit scribbling down our words.

He said, *You'd take a pill that gave you 10 more good years, if I offered it to you now...*

(My older son is still very much in the immortality camp. It pains and angers him that I would spurn a project that could bring us near-infinite rewards. To him it looks as if I am giving up on life itself.)

Seymour Benzer would not have liked how crowded the aging field has become. He was drawn to it partly because it was small. Once a new science got established, with rafts of research to follow, journals to keep up with, conferences to attend, he felt he no longer had the space to think. He moved on to something else.

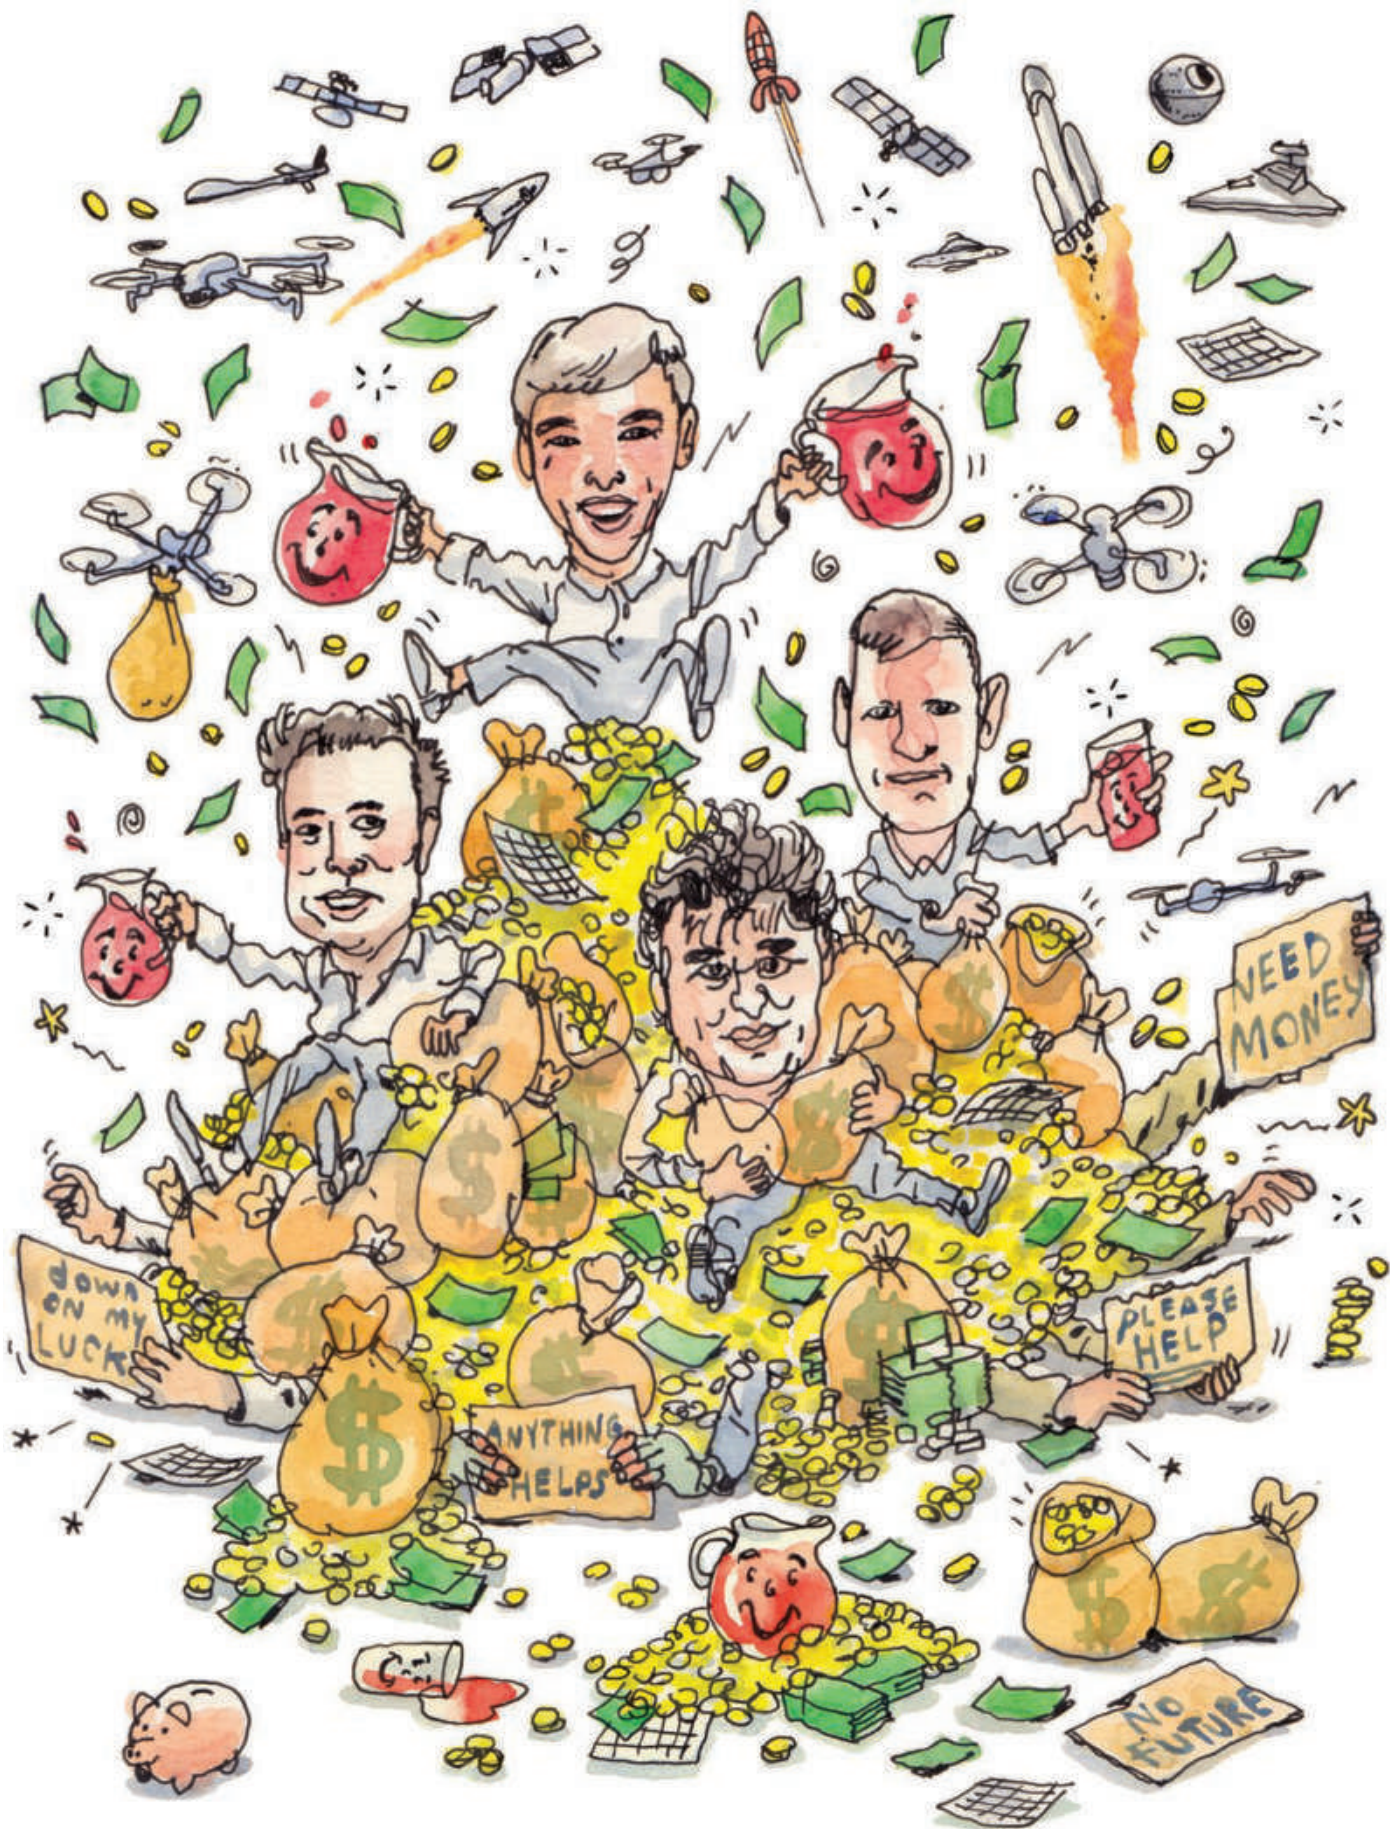
For more than 60 years, the span of his long career, that strategy worked to stave off boredom. He was one of those mortals who show you how rich a finite life can be. In his 20s, his work in physics helped lead to the invention of the transistor. In his 30s and 40s, his work on the fine structure of the gene helped launch molecular biology. Then it was neurogenetics. Then the modern science of aging—along with much else. Even in his old age, he radiated curiosity: always the next field, the next experiment. I'd take a pill for that.

Benzer died from a stroke in November of 2007, at Huntington Hospital in Pasadena, California, at the age of 86. I heard from his family and friends that he worked cheerfully in his lab right to the end. At the hospital, just before he slid into a coma, he was still alert enough to look at the doctors and say, "I have two questions ..."

And that was all. ■

Jonathan Weiner is a writer based in New York City. He won a National Book Critics Circle Award in 1999 for *Time, Love, Memory*, his book about Seymour Benzer. He teaches at Columbia Journalism School.

good years to our lives
number of bad years at the end. 



Can data really quantify the
“best way” to help the world?

Effective altruism

Oregon 6th Congressional District candidate Carrick Flynn seemed to drop out of the sky.

With a stint at Oxford’s Future of Humanity Institute, a track record of voting in only two of the past 30 elections, and \$11 million in support from a political action committee established by crypto billionaire Sam Bankman-Fried, Flynn didn’t fit into the local political scene, even though he’d grown up in the state. One constituent called him “Mr. Creepy Funds” in an interview with a local paper; another said he thought Flynn was a Russian bot.

The specter of crypto influence, a slew of expensive TV ads, and the fact that few locals had heard of or spoken to Flynn raised suspicions that he was a tool of outside financial interests. And while the rival candidate who led the primary race promised to fight for issues like better worker

and
its

protections and stronger gun legislation, Flynn’s platform prioritized economic growth and preparedness for pandemics and other disasters. Both are pillars of “longtermism,” a growing strain of the ideology known as effective altruism (or EA), which is popular among an elite slice of people in tech and politics.

Even during an actual pandemic, Flynn’s focus struck many Oregonians as far-fetched and foreign. Perhaps unsurprisingly, he ended up losing the 2022 primary to the more politically experienced Democrat, Andrea Salinas. But despite Flynn’s lackluster showing, he made history as effective altruism’s first political candidate to run for office.

Since its birth in the late 2000s, effective altruism has aimed to answer the question “How can those with means have the *most* impact on the world in a quantifiable way?”—and supplied clear methodologies

growing influence

By Rebecca Ackermann

Illustrations by Victor Kerlow

for calculating the answer. Directing money to organizations that use evidence-based approaches is the one technique EA is most known for. But as it has expanded from an academic philosophy into a community and a movement, its ideas of the “best” way to change the world have evolved as well.

“Longtermism,” the belief that unlikely but existential threats like a humanity-destroying AI revolt or international biological warfare are humanity’s most pressing problems, is integral to EA today. Of late, it has moved from the fringes of the movement to its fore with Flynn’s campaign, a flurry of mainstream media coverage, and a new treatise published by one of EA’s founding fathers, William MacAskill. It’s an ideology that’s poised to take the main stage as more believers in the tech and billionaire classes—which are, notably, mostly male and white—start to pour millions into new PACs and projects like Bankman-Fried’s FTX Future Fund and Longview Philanthropy’s Longtermism Fund, which focus on theoretical menaces ripped from the pages of science fiction.

EA’s ideas have long faced criticism from within the fields of philosophy and philanthropy that they reflect white Western saviorism and an avoidance of structural problems in favor of abstract math—not coincidentally, many of the same objections lobbed at the tech industry at large. Such charges are only intensifying as EA’s pockets deepen and its purview stretches into a galaxy far, far away. Ultimately, the philosophy’s influence may be limited by their accuracy.

WHAT IS EA?

If

effective altruism were a lab-grown species, its origin story would begin with DNA spliced from three parents: applied ethics, speculative technology, and philanthropy.

EA’s philosophical genes came from Peter Singer’s brand of utilitarianism and Oxford philosopher Nick Bostrom’s investigations into potential threats to humanity. From tech, EA drew on early research into the long-term impact of artificial intelligence carried out at what’s now known as the Machine Intelligence Research Institute (MIRI) in Berkeley, California. In philanthropy, EA is part of a growing trend toward evidence-based giving, driven by members of the Silicon Valley nouveau riche who are eager to apply the strategies that made them money to the process of giving it away.

While these origins may seem diverse, the people involved are linked by social, economic, and professional class, and by a tech-utopian worldview. Early players—including MacAskill, a Cambridge philosopher; Toby Ord, an Oxford philosopher; Holden Karnofsky, cofounder of the charity evaluator GiveWell; and Dustin Moskovitz, a cofounder of Facebook who founded the nonprofit Open Philanthropy with his wife, Cari Tuna—are all still leaders in the movement’s interconnected constellation of nonprofits, foundations, and research organizations.

For effective altruists, a good cause is not good enough; only the

very best should get funding in the areas most in need. Those areas are usually, by EA calculations, developing nations. Personal connections that might encourage someone to give to a local food bank or donate to the hospital that treated a parent are a distraction—or worse, a waste of money. The classic example of an EA-approved effort is the Against Malaria Foundation, which purchases and distributes mosquito nets in sub-Saharan Africa and other areas heavily affected by the disease. The price of a net is very small compared with the scale of its life-saving potential; this ratio of “value” to cost is what EA aims for. Other popular early EA causes include providing vitamin A supplements and malaria medication in African countries, and promoting animal welfare in Asia.

Within effective altruism’s framework, selecting one’s career is just as important as choosing where to make donations. EA defines a professional “fit” by whether a candidate has comparative advantages like exceptional intelligence or an entrepreneurial drive, and if an effective altruist qualifies for a high-paying path, the ethos encourages “earning to give,” or dedicating one’s life to building wealth in order to give it away to EA causes. Bankman-Fried has said that he’s earning to give, even founding the crypto platform FTX with the express purpose of building wealth in order to redirect 99% of it. Now one of the richest crypto executives in the world, Bankman-Fried plans to give away up to \$1 billion by the end of 2022.

“The allure of effective altruism has been that it’s an off-the-shelf methodology for being a highly sophisticated, impact-focused, data-driven funder,” says David Callahan, founder and editor of Inside Philanthropy and the author

Sam Bankman-Fried, one of the world’s richest crypto executives, is also one of the country’s largest political donors. He plans to give away up to

\$1 billion
by the end of 2022.



For effective altruists, a good cause is not good enough; only the very best should get funding in the areas most in need.

of a 2017 book on philanthropic trends, *The Givers*. Not only does EA suggest a clear and decisive framework, but the community also offers a set of resources for potential EA funders—including GiveWell, a nonprofit that uses an EA-driven evaluation rubric to recommend charitable organizations; EA Funds, which allows individuals to donate to curated pools of charities; 80,000 Hours, a career-coaching organization; and a vibrant discussion forum at Effectivealtruism.org, where leaders like MacAskill and Ord regularly chime in.

Effective altruism's original laser focus on measurement has contributed rigor in a field that has historically lacked accountability for big donors with last names like Rockefeller and Sackler. "It has been an overdue, much-needed

counterweight to the typical practice of elite philanthropy, which has been very inefficient," says Callahan.

But where exactly are effective altruists directing their earnings? Who benefits? As with all giving—in EA or otherwise—there are no set rules for what constitutes "philanthropy," and charitable organizations benefit from a tax code that incentivizes the super-rich to establish and control their own charitable endeavors at the expense of public tax revenues, local governance, or public accountability. EA organizations are able to leverage the practices of traditional philanthropy while enjoying the shine of an effectively disruptive approach to giving. The movement has formalized its community's commitment to donate with the Giving What We Can Pledge—mirroring

another old-school philanthropic practice—but there are no giving requirements to be publicly listed as a pledger. Tracking the full influence of EA's philosophy is tricky, but 80,000 Hours has estimated that \$46 billion was committed to EA causes between 2015 and 2021, with donations growing about 20% each year. GiveWell calculates that in 2021 alone, it directed over \$187 million to malaria nets and medication; by the organization's math, that's over 36,000 lives saved.

Accountability is significantly harder with longtermist causes like biosecurity or "AI alignment"—a set of efforts aimed at ensuring that the power of AI is harnessed toward ends generally understood as "good." Such causes, for a growing number of effective altruists, now take priority over mosquito nets and vitamin A medication. "The things that matter most are the things that have long-term impact on what the world will look like," Bankman-Fried said in an interview earlier this year. "There are trillions of people who have not yet been born." Bankman-Fried's views are influenced by longtermism's utilitarian calculations, which flatten lives into single units of value. By this math, the trillions of humans yet to be born represent a greater moral obligation than the billions alive today. Any threats that could prevent future generations from reaching their full potential—either through extinction or through technological stagnation, which MacAskill deems equally dire in his new book, *What We Owe the Future*—are priority number one.

In his book, MacAskill discusses his own journey from longtermism skeptic to true believer and urges other to follow the same path. The existential risks he lays out are specific: "The future could be terrible, falling to authoritarians who use

surveillance and AI to lock in their ideology for all time, or even to AI systems that seek to gain power rather than promote a thriving society. Or there could be no future at all: we could kill ourselves off with biological weapons or wage an all-out nuclear war that causes civilisation to collapse and never recover.”

It was to help guard against these exact possibilities that Bankman-Fried created the FTX Future Fund this year as a project within his philanthropic foundation. Its focus areas include “space governance,” “artificial intelligence,” and “empowering exceptional people.” The fund’s website acknowledges that many of its bets “will fail.” (Its primary goal for 2022 is to test new funding models, but the fund’s site does not establish what “success” may look like.) As of June 2022, the FTX Future Fund had made 262 grants and investments, with recipients including a Brown University academic researching long-term economic growth, a Cornell University academic researching AI alignment, and an organization working on legal research around AI and biosecurity (which was born out of Harvard Law’s EA group).

Bankman-Fried is hardly the only tech billionaire pushing forward longtermist causes. Open Philanthropy, the EA charitable organization funded primarily by Moskowitz and Tuna, has directed \$260 million to addressing “potential risks from advanced AI” since its founding. Together, the FTX Future Fund and Open Philanthropy supported Longview Philanthropy with more than \$15 million this year before the organization announced its new Longtermism Fund. Vitalik Buterin, one of the founders of the blockchain platform Ethereum, is the second-largest recent donor to MIRI, whose mission is “to ensure [that] smarter-than-human artificial

intelligence has a positive impact.” MIRI’s donor list also includes the Thiel Foundation; Ben DeLo, cofounder of crypto exchange BitMEX; and Jaan Tallinn, one of the founding engineers of Skype, who is also a cofounder of Cambridge’s Centre for the Study of Existential Risk (CSER). Elon Musk is yet another tech mogul dedicated to fighting longtermist existential risks; he’s even claimed that his for-profit operations—including SpaceX’s mission to Mars—are philanthropic efforts supporting humanity’s progress and survival. (MacAskill has recently expressed concern that his philosophy is getting conflated with Musk’s “worldview.” However, EA aims for an expanded audience, and it seems unreasonable to expect rigid adherence to the exact belief system of its creators.)

EA’s earn-to-give philosophy raises the question of why the wealthy should get to decide where funds go.

CRITICISM AND CHANGE

Even

before the foregrounding of longtermism, effective altruism had been criticized for elevating the mindset of the “benevolent capitalist” (as philosopher Amia Srinivasan wrote in her 2015 review of MacAskill’s first book) and emphasizing individual agency within capitalism over more foundational critiques of the systems that have made one part of the world wealthy enough to spend time theorizing about how best to aid the rest.

EA’s earn-to-give philosophy raises the question of why the wealthy should get to decide where funds go in a highly inequitable world—especially if they may be extracting that wealth from employees’ labor or the public, as may be

the case with some crypto executives. “My ideological orientation starts with the belief that folks don’t earn tremendous amounts of money without it being at the expense of other people,” says Farhad Ebrahimi, founder and president of the Chorus Foundation, which funds mainly US organizations working to combat climate change by shifting economic and political power to the communities most affected by it.

Many of the foundation’s grantees are groups led by people of color, and it is what’s known as a spend-down foundation; in other words, Ebrahimi says, Chorus’s work will be successful when its funds are fully redistributed.

Ebrahimi objects to EA’s approach of supporting targeted interventions rather than endowing local organizations to define their own priorities: “Why wouldn’t you want to support having the communities that you want the money to go to be the ones to build economic power? That’s an individual saying, ‘I want to build my economic power because I think I’m going to make good decisions about what to do with it’ ... It seems very ‘benevolent dictator’ to me.”

Effective altruists would respond that their moral obligation is to fund the most demonstrably transformative projects as defined by their framework, no matter what else is left behind. In an interview in 2018, MacAskill suggested that in order to recommend prioritizing any structural power shifts, he’d need to see “an argument that opposing inequality in some particular way is actually going to be the best thing to do.”

However, when a small group of individuals with similar backgrounds

have determined the formula for the most critical causes and “best” solutions, the unbiased rigor that EA is known for should come into question. While the top nine charities featured on GiveWell’s website today work in developing nations with communities of color, the EA community stands at 71% male and 76% white, with the largest percentage living in the US and the UK, according to a 2020 survey by the Centre for Effective Altruism (CEA). This may not be surprising given that the philanthropic community at large has long been criticized for homogeneity. But some studies have demonstrated that charitable giving in the US is actually growing in diversity, which casts EA’s breakdown in a different light. A 2012 report by the W. K. Kellogg Foundation found that both Asian-American and Black households gave away a larger percentage of their income than white households. Research from the Indiana University Lilly Family School of Philanthropy found in 2021 that 65% of Black households and 67% of Hispanic households surveyed donated charitably on a regular basis, along with 74% of white households. And donors of color were more likely to be involved in more informal avenues of giving, such as crowdfunding, mutual aid, or giving circles, which may not be accounted for in other reports. EA’s sales pitch does not appear to be reaching these donors.

While EA proponents say its approach is data driven, EA’s calculations defy best practices within the tech industry around dealing with data. “This assumption that we’re



Open Philanthropy, the EA charitable organization funded primarily by Dustin Moskovitz and Cari Tuna, has directed

\$260 million

to addressing “potential risks from advanced AI” since its founding.

going to calculate the single best thing to do in the world—have all this data and make these decisions—is so similar to the issues that we talk about in machine learning, and why you shouldn’t do that,” says Timnit Gebru, a leader in AI ethics and the founder and executive director of the Distributed AI Research Institute (DAIR), which centers diversity in its AI research.

Gebru and others have written extensively about the dangers of leveraging data without undertaking deeper analysis and making sure it comes from diverse sources. In machine learning, it leads to dangerously biased models. In philanthropy, a narrow definition of success rewards alliance with EA’s value system over other worldviews and penalizes nonprofits working on longer-term or more complex strategies that can’t be translated into EA’s math. The research that EA’s assessments rely on may also be flawed or subject to change; a 2004 study that elevated deworming—distributing drugs for parasitic infections—to one of GiveWell’s top causes has come under serious fire, with some researchers claiming to have debunked it while others have been unable to replicate the results leading to the conclusion that it would save huge numbers of lives. Despite the uncertainty surrounding this intervention, GiveWell directed more than \$12 million to deworming charities through its Maximum Impact Fund this year.

The voices of dissent are growing louder as EA’s influence spreads and more money is directed toward longtermist causes. A longtermist himself by some definitions, CSER researcher Luke Kemp believes that the growing focus of the EA research community is based on a limited and minority perspective. He’s been disappointed with the lack of diversity of thought and

leadership he's found in the field. Last year, he and his colleague Carla Zoe Cremer wrote and circulated a preprint titled "Democratizing Risk" about the community's focus on the "techno-utopian approach"—which assumes that pursuing technology to its maximum development is an undeniable net positive—to the exclusion of other frameworks that reflect more common moral worldviews. "There's a small number of key funders who have a very particular ideology, and either consciously or unconsciously select for the ideas that most resonate with what they want. You have to speak that language to move higher up the hierarchy and get more funding," Kemp says.

Even the basic concept of longtermism, according to Kemp, has been hijacked from legal and economic scholars in the 1960s, '70s, and '80s, who were focused on intergenerational equity and environmentalism—priorities that have notably dropped away from the EA version of the philosophy. Indeed, the central premise that "future people count," as MacAskill says in his 2022 book, is hardly new. The Native American concept of the "seventh generation principle" and similar ideas in indigenous cultures across the globe ask each generation to consider the ones that have come before and will come after. Integral to these concepts, though, is the idea that the past holds valuable lessons for action today, especially in cases where our ancestors made choices that have led to environmental and economic crises.

Longtermism sees history differently: as a forward march toward inevitable progress. MacAskill references the past often in *What We Owe the Future*, but only in the form of case studies on the life-improving impact of technological and moral development. He

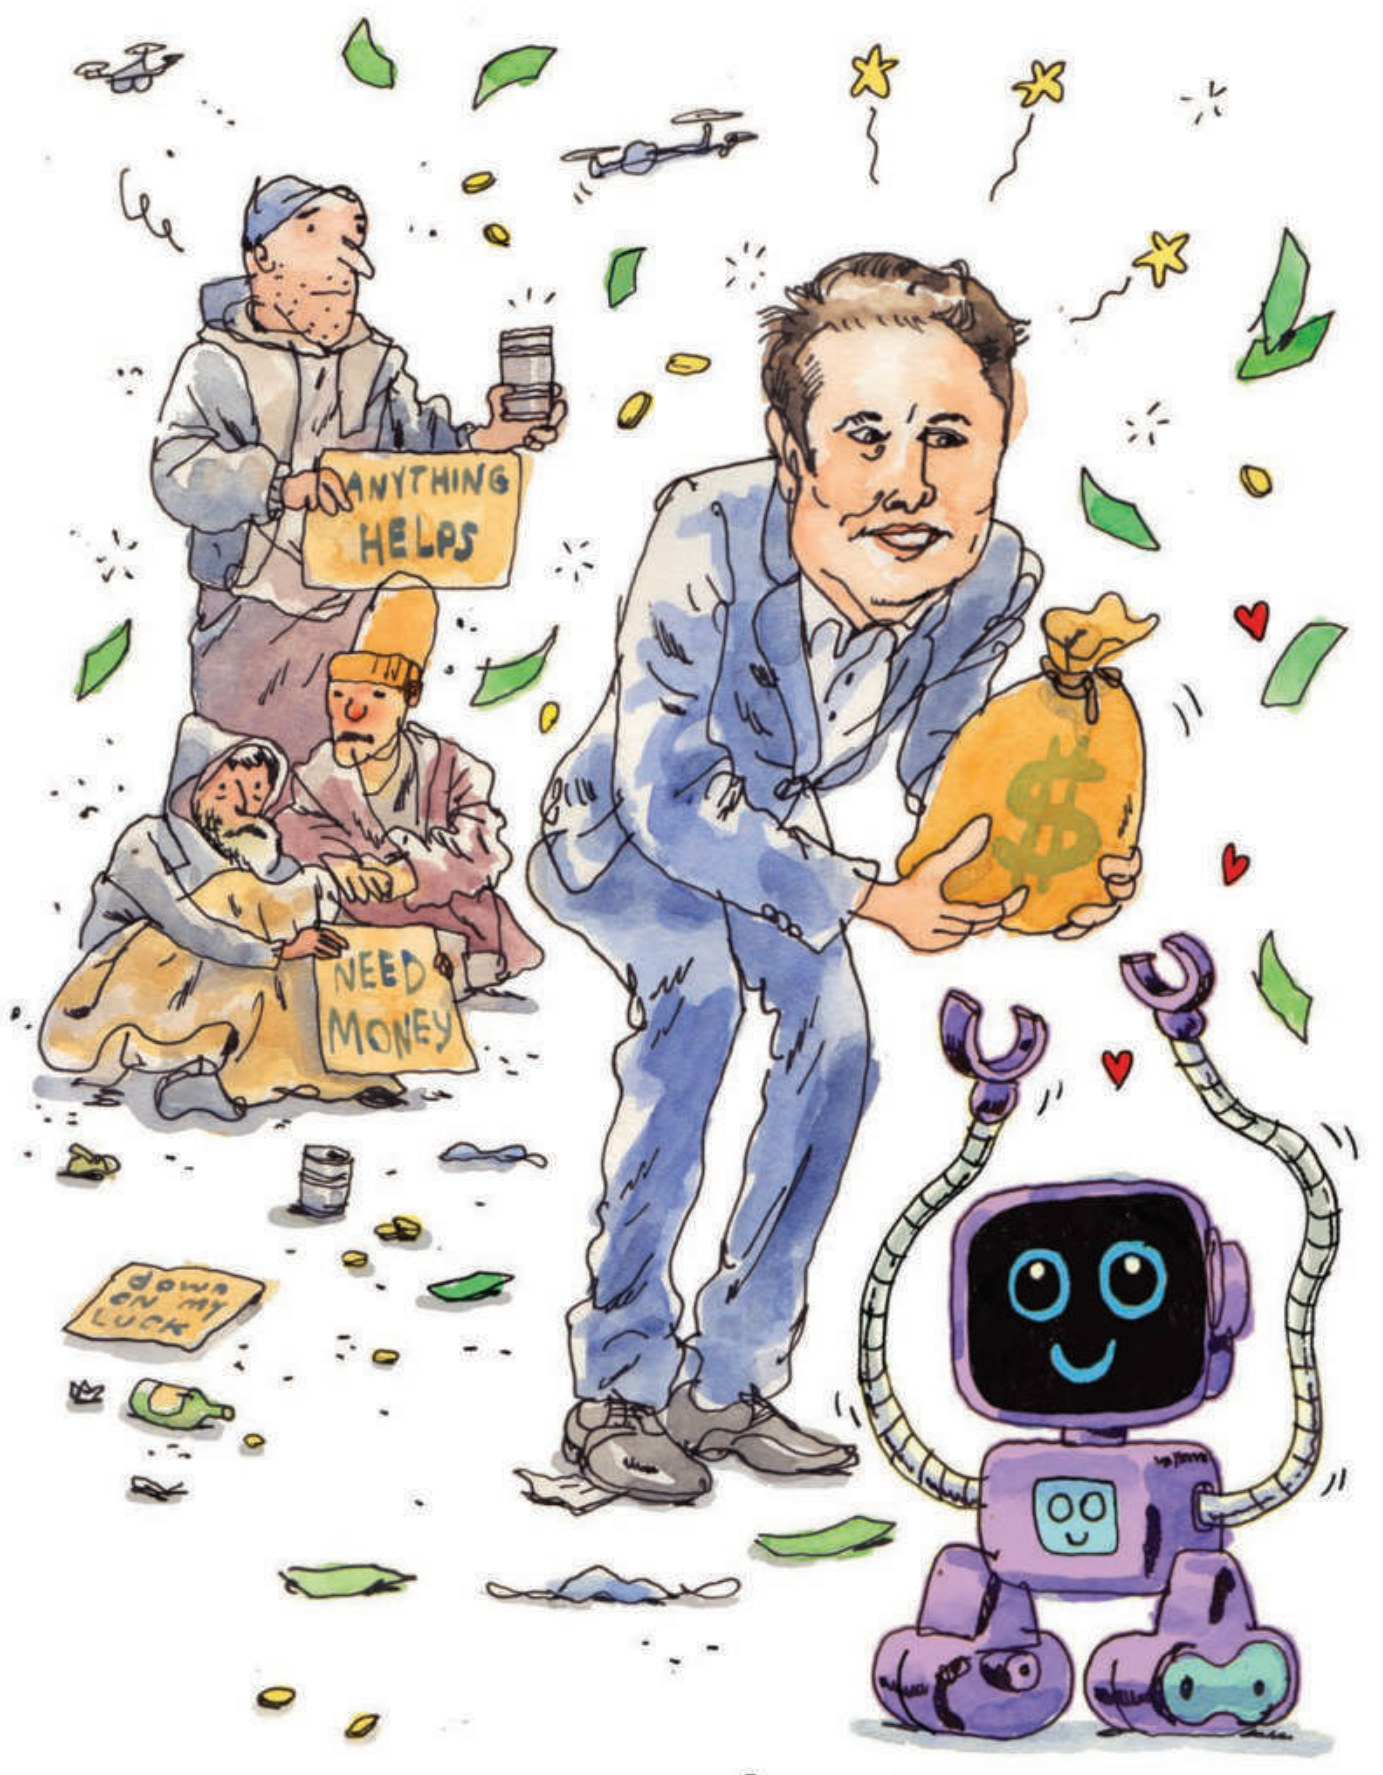
Longtermism sees history as a forward march toward inevitable progress.

discusses the abolition of slavery, the Industrial Revolution, and the women's rights movement as evidence of how important it is to continue humanity's arc of progress before the wrong values get "locked in" by despots. What are the "right" values? MacAskill has a coy approach to articulating them: he argues that "we should focus on promoting more abstract or general moral principles" to ensure that "moral changes stay relevant and robustly positive into the future."

Worldwide and ongoing climate change, which already affects the under-resourced more than the elite today, is notably not a core longtermist cause, as philosopher Emile P. Torres points out in his critiques. While it poses a threat to millions of lives, longtermists argue, it probably won't wipe out *all* of humanity;

those with the wealth and means to survive can carry on fulfilling our species' potential. Tech billionaires like Thiel and Larry Page already have plans and real estate in place to ride out a climate apocalypse. (MacAskill, in his new book, names climate change as a serious worry for those alive today, but he considers it an existential threat only in the "extreme" form where agriculture won't survive.)

The final mysterious feature of EA's version of the long view is how its logic ends up in a specific list of technology-based far-off threats to civilization that just happen to align with many of the original EA cohort's areas of research. "I am a researcher in the field of AI," says Gebu, "but to come to the conclusion that in order to do the most good in the world you have to work



on artificial general intelligence is very strange. It's like trying to justify the fact that you want to think about the science fiction scenario and you don't want to think about real people, the real world, and current structural issues. You want to justify how you want to pull billions of dollars into that while people are starving."

Some EA leaders seem aware that criticism and change are key to expanding the community and strengthening its impact. MacAskill and others have made it explicit that their calculations are estimates ("These are our best guesses," MacAskill offered on a 2020 podcast episode) and said they're eager to improve through critical discourse. Both GiveWell and CEA have pages on their websites titled "Our Mistakes," and in June, CEA ran a contest inviting critiques on the EA forum; the Future Fund has launched prizes up to \$1.5 million for critical perspectives on AI.

"We recognize that the problems EA is trying to address are really, really big and we don't have a hope of solving them with only a small segment of people," GiveWell board member and CEA community liaison Julia Wise says of EA's diversity statistics. "We need the talents that lots of different kinds of people can bring to address these worldwide problems." Wise also spoke on the topic at the 2020 EA Global Conference, and she actively discusses inclusion and community power dynamics on the CEA forum. The Center for Effective Altruism supports a mentorship program for women and nonbinary people (founded, incidentally, by

Carrick Flynn's wife) that Wise says is expanding to other underrepresented groups in the EA community, and CEA has made an effort to facilitate conferences in more locations worldwide to welcome a more geographically diverse group. But these efforts appear to be limited in scope and impact; CEA's public-facing page on diversity and inclusion hasn't even been updated since 2020. As the tech-utopian tenets of longtermism take a front seat in EA's rocket ship and a few billionaire donors chart its path into the future, it may be too late to alter the DNA of the movement.

POLITICS AND THE FUTURE

Despite

the sci-fi sheen, effective altruism today is a conservative project, consolidating decision-making behind a technocratic belief system and a small set of individuals, potentially at the expense of local and intersectional visions for the future. But EA's community and successes were built around clear methodologies that may not transfer into the more nuanced political arena that some EA leaders and a few big donors are pushing toward. According to Wise, the community at large is still split on politics as an approach to pursuing EA's goals, with some dissenters believing politics is too polarized a space for effective change.

But EA is not the only charitable movement looking to political action to reshape the world; the philanthropic field generally has been moving into politics for greater impact. "We have an existential



Ethereum cofounder Vitalik Buterin is the

second-largest

recent donor to Berkeley's Machine Intelligence Research Institute, whose mission is "to ensure [that] smarter-than-human artificial intelligence has a positive impact."

political crisis that philanthropy has to deal with. Otherwise, a lot of its other goals are going to be hard to achieve," says Inside Philanthropy's Callahan, using a definition of "existential" that differs from MacAskill's. But while EA may offer a clear rubric for determining how to give charitably, the political arena presents a messier challenge. "There's no easy metric for how to gain political power or shift politics," he says. "And Sam Bankman-Fried has so far demonstrated himself not the most effective political giver."

Bankman-Fried has articulated his own political giving as "more policy than politics," and has donated primarily to Democrats through his short-lived Protect Our Future PAC (which backed Carrick Flynn in Oregon) and the Guarding Against Pandemics PAC (which is run by his brother Gabe and publishes a cross-party list of its "champions" to support). Ryan Salame, the co-CEO with Bankman-Fried of FTX, funded his own PAC, American Dream Federal Action, which focuses mainly on Republican candidates. (Bankman-Fried has said Salame shares his passion for preventing pandemics.) Guarding Against Pandemics and the Open Philanthropy Action Fund (Open Philanthropy's political arm) spent more than \$18 million to get an initiative on the California state ballot this fall to fund pandemic research and action through a new tax.

So while longtermist funds are certainly making waves behind the scenes, Flynn's primary loss in Oregon may signal that EA's more visible electoral efforts need to draw on new and diverse strategies to win over real-world voters. Vanessa Daniel, founder and former executive director of Groundswell, one of the largest funders of the US reproductive justice movement, believes that big donations and 11th-hour interventions

“To come to the conclusion that in order to do the most good in the world you have to work on artificial general intelligence is very strange.”

will never rival grassroots organizing in making real political change. “Slow and patient organizing led by Black women, communities of color, and some poor white communities created the tipping point in the 2020 election that saved the country from fascism and allowed some window of opportunity to get things like the climate deal passed,” she says. And Daniel takes issue with the idea that metrics are the exclusive domain of rich, white, and male-led approaches. “I’ve talked to so many donors who think that grassroots organizing is the equivalent of planting magical beans and expecting things to grow. This is not the case,” she says. “The data is right in front of us. And it doesn’t require the collateral damage of millions of people.”

The question now is whether the culture of EA will allow the

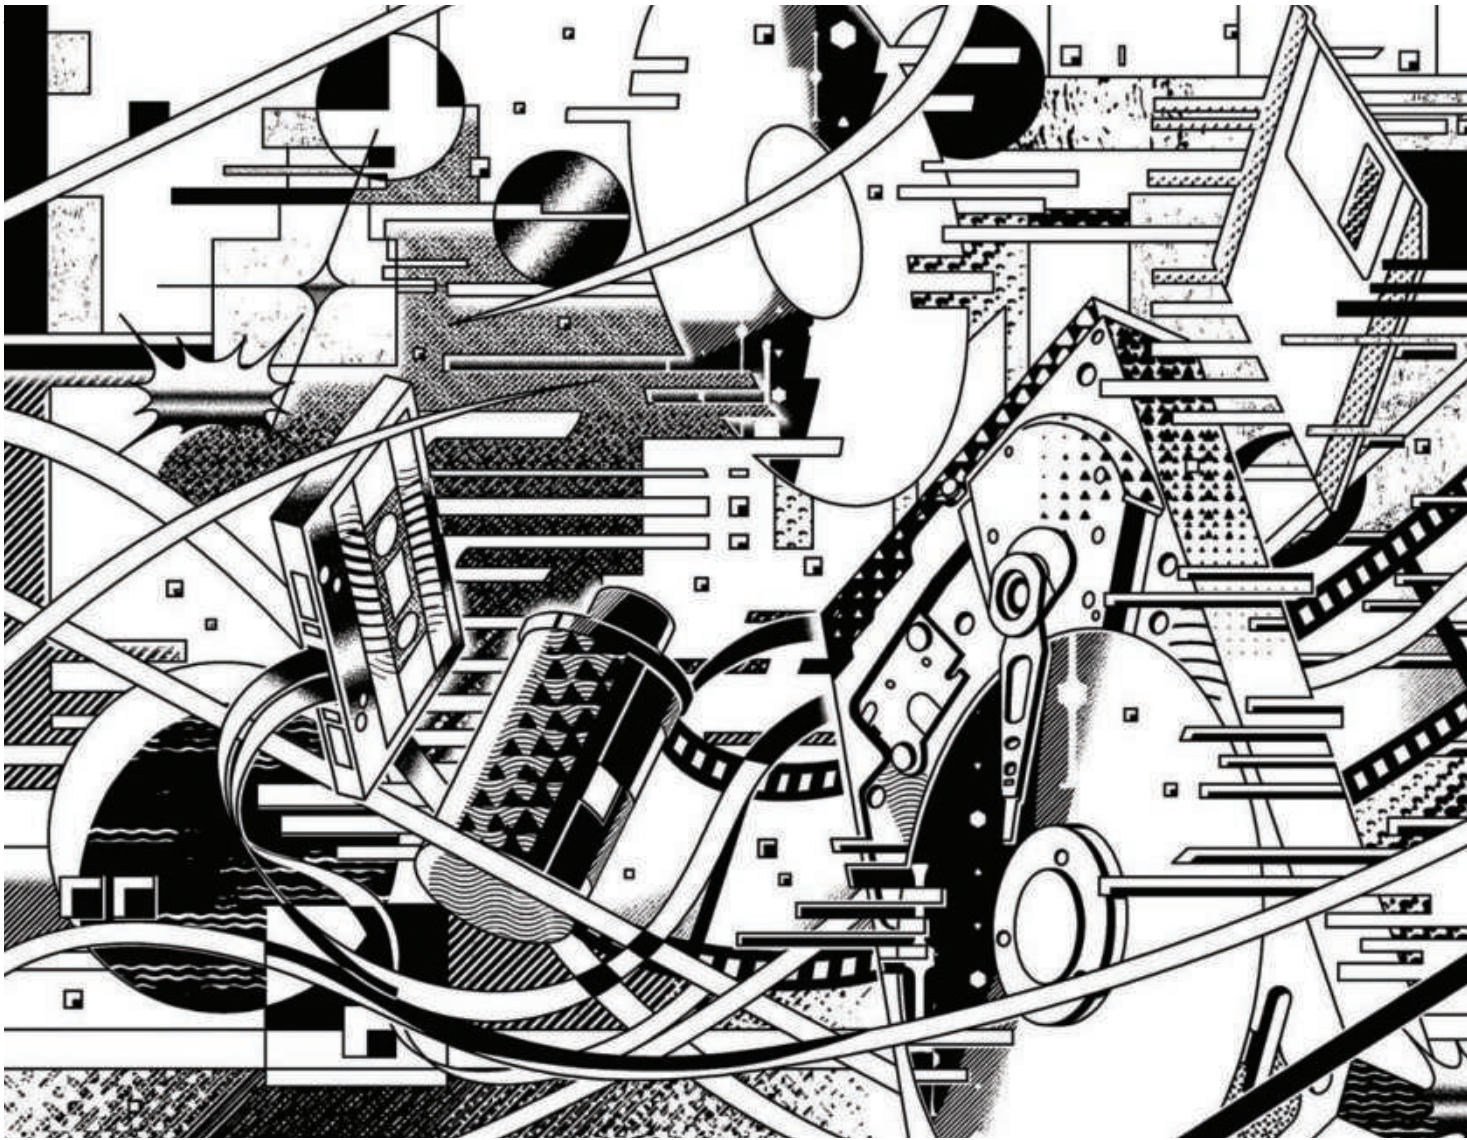
community and its major donors to learn from such lessons. In May, Bankman-Fried admitted in an interview that there are a few takeaways from the Oregon loss, “in terms of thinking about who to support and how much,” and that he sees “decreasing marginal gains from funding.” In August, after distributing a total of \$24 million over six months to candidates supporting pandemic prevention, Bankman-Fried appeared to have shut down funding through his Protect Our Future PAC, perhaps signaling an end to one political experiment. (Or maybe it was just a pragmatic belt-tightening after the serious and sustained downturn in the crypto market, the source of Bankman-Fried’s immense wealth.)

Others in the EA community draw different lessons from the

Flynn campaign. On the forum at [Effectivealtruism.org](https://www.effectivealtruism.org), Daniel Eth, a researcher at the Future of Humanity Institute, posted a lengthy postmortem of the race, expressing surprise that the candidate couldn’t win over the general audience when he seemed “unusually selfless and intelligent, even for an EA.” But Eth didn’t encourage radically new strategies for a next run apart from ensuring that candidates vote more regularly and spend more time in the area. Otherwise, he proposed doubling down on EA’s existing approach: “Politics might somewhat degrade our typical epistemics and rigor. We should guard against this.” Members of the EA community contributing to the 93 comments on Eth’s post offered their own opinions, with some supporting Eth’s analysis, others urging lobbying over electioneering, and still others expressing frustration that effective altruists are funding political efforts at all. At this rate, political causes are not likely to make it to the front page of GiveWell anytime soon.

Money can move mountains, and as EA takes on larger platforms with larger amounts of funding from billionaires and tech industry insiders, the wealth of a few billionaires will likely continue to elevate pet EA causes and candidates. But if the movement aims to conquer the political landscape, EA leaders may find that whatever its political strategies, its messages don’t connect with the people who are living with local and present-day challenges like insufficient housing and food insecurity. EA’s academic and tech industry origins as a heady philosophical plan for distributing inherited and institutional wealth may have gotten the movement this far, but those same roots likely can’t support its hopes for expanding its influence. ■

Rebecca Ackermann is a writer and artist in San Francisco.



The death of knowledge

We can preserve a lot of information digitally, but the process is far from perfect.

By Erik Sherman

Illustration by Jinhwa Jang

Everything dies: people, machines, civilizations. Perhaps we can find some solace in knowing that all the meaningful things we've learned along the way will survive. But even knowledge has a life span. Documents fade. Art goes missing. Entire libraries and collections can face quick and unexpected destruction.

Surely, we're at a stage technologically where we might devise ways to make knowledge available and accessible forever. After all, the density of data storage is already incomprehensibly high. In the ever-growing museum of the internet, one can move smoothly from images from the James Webb Space Telescope through diagrams explaining Pythagoras's philosophy on the music of the spheres to a YouTube tutorial on blues guitar soloing. What more could you want?

Quite a bit, according to the experts. For one thing, what we think is permanent isn't. Digital storage systems can become



unreadable in as little as three to five years. Librarians and archivists race to copy things over to newer formats. But entropy is always there, waiting in the wings. “Our professions and our people often try to extend the normal life span as far as possible through a variety of techniques, but it’s still holding back the tide,” says Joseph Janes, an associate professor at the University of Washington Information School.

To complicate matters, archivists are now grappling with an unprecedented deluge of information. In the past, materials were scarce and storage space limited. “Now we have the opposite problem,” Janes says. “Everything is being recorded all the time.”

In principle, that could right a historic wrong. For centuries, countless people didn’t have the right culture, gender, or socioeconomic class for their knowledge or work to be discovered, valued, or preserved. But the massive scale of the digital world now presents a unique challenge. According to an estimate last year from the market research

firm IDC, the amount of data that companies, governments, and individuals create in the next few years will be twice the total of all the digital data generated previously since the start of the computing age.

Entire schools within some universities are laboring to find better approaches to saving the data under their umbrella. The Data and Service Center for Humanities at the University of Basel, for example, has been developing a software platform called Knora to not just archive the many types of data from humanities work but ensure that people in the future can read and use them. And yet the process is fraught.

“You make educated guesses and hope for the best, but there are data sets that are lost because nobody knew they’d be useful,” says Andrea Ogier, assistant dean and

“We can’t save everything ... but that’s no reason to not do what we can.”

director of data services at the University Libraries of Virginia Tech.

There are never enough people or money to do all the necessary work—and formats are changing and multiplying all the time. “How do we best allocate resources to preserve things? Because budgets are only so large,” Janes says. “In some cases, that means stuff gets saved or stored but just sits there, uncatalogued and unprocessed, and thus next to impossible to find or access.” In some cases, archivists ultimately turn away new collections.

The formats used to store data are themselves impermanent. NASA socked away 170 or so tapes of data on lunar dust, collected during the Apollo era. When researchers set out to use the tapes in the mid-2000s, they couldn’t find anyone with the 1960s-era IBM 729 Mark 5 machine needed to read them. With help, the team ultimately tracked down one in rough shape at the warehouse of the Australian Computer Museum. Volunteers helped refurbish the machine.

Software also has a shelf life. Ogier recalls trying to examine an old Quattro Pro spreadsheet file only to find there was no readily available software that could read it.

There have been attempts to future-proof programs. One project that got a lot of fanfare in 2015 is the Open Library of Images for Virtualized Execution (Olive) archive, which runs old software like Chaste 3.1, a 2013 biology and physiology research program, and the 1990 Mac version of the computer game *The Oregon Trail* on a set of virtual machines. The project is still active, says Mahadev Satyanarayanan, a professor of computer science at Carnegie Mellon University. But there have been challenges in expanding Olive’s offerings, he says: even unused software has to be licensed from the companies that own it, and there is often no easy way to enter new data into the archive’s research applications.

Other efforts to help advance the longevity of knowledge have also had mixed results. The Internet Archive, home of the Wayback Machine, has a large collection of digitized materials, including software, music, and videos; as of the summer of 2022 it was fighting a copyright infringement lawsuit brought by multiple publishers.

On the more hopeful side, the Text Encoding Initiative has maintained international standards for encoding machine-readable texts since the 1990s. A decade ago, the US Office of Science and Technology Policy stipulated that applications for federally supported research have to provide a data management plan so the data can be used by researchers or the public in the future. “We’re getting to the point where almost every grant-funded research project has to put its data somewhere,” Ogier says. But there are no overarching requirements about who must store the data or how long it must be saved.

Unavoidably, ideas, knowledge, and human creations will continue to be lost. “We can’t save everything. We can’t provide access to everything. We can’t retrieve everything,” Ogier says. “But that’s no reason to not do what we can.” ■

Erik Sherman is a freelance journalist based in Ashfield, Massachusetts.



Research labs are pursuing technology to “reprogram” aging bodies back to youth.

A little over 15 years ago, scientists at Kyoto University in Japan made a remarkable discovery. When they added just four proteins to a skin cell and waited about two weeks, some of the cells underwent an unexpected and astounding transformation: they became young again. They turned into stem cells almost identical to the kind found in a days-old embryo, just beginning life’s journey.

At least in a petri dish, researchers using the procedure can take withered skin cells from a 101-year-old and rewind them so they act as if they’d never aged at all.

By Antonio Regalado
Illustration by Max Loeffler

How
to

become

young

again

Now, after more than a decade of studying and tweaking so-called cellular reprogramming, a number of biotech companies and research labs say they have tantalizing hints the process could be the gateway to an unprecedented new technology for age reversal. By applying limited, controlled doses of the reprogramming proteins to lab animals, the scientists say, they are seeing evidence that the procedure makes the animals—or at least some of their organs—more youthful.

One of the key promoters of this idea, Richard Klausner, took the stage in June at a glitzy, \$4,000-per-ticket retreat in San Diego, where he flashed data from unpublished experiments in which sick mice bounced back to health after undergoing the experimental treatment.

Klausner was pitching nothing less than “medical rejuvenation”—a means of taking old animals and making them “young.” He is the organizer and chief scientist of Altos Labs, a new research company seeded with more than \$3 billion from ultra-wealthy figures in Silicon Valley and oil money from the Persian Gulf. Klausner and his financiers had swept up dozens of top scientists—offering salaries of \$1 million and more—and set them to work on a technology the company now calls “rejuvenation programming.”

It seems to work at least in part by resetting what’s called the epigenome—chemical marks on DNA that control which genes are turned on, or off, in a cell. In aging, some of these markers get flipped to the wrong positions. Reprogramming is a technology that can flip them back. But it can also change cells in dangerous ways, even causing cancer.

The objective of Altos is to tame this phenomenon, understand it, and eventually apply it as a treatment to reverse a wide range of diseases. This may be possible, Klausner says, because youthful cells have more resilience and can bounce back from biological stress in ways old ones don’t. And Klausner has data to suggest it might already be working. During his talk, he showed slides marked “Confidential” claiming that fat mice had recovered from diabetes after treatment, and that others were able to survive normally lethal doses of painkillers—all thanks to a healthy dose of the medical rejuvenation.

“We think we can turn back the clock,” he told the audience.

Klausner is the former head of the National Cancer Institute and onetime leader for global health at the Gates Foundation. He is a heavy hitter who has also been behind some of today’s most high-profile biotech ventures, like the cancer blood-test company Grail. Yet even for him, rejuvenation is wildly ambitious. That is because if you can make cells act younger, healthier, and more resilient, you might have a general-purpose means of forestalling many diseases all at once. “This is the opposite of precision medicine,” Klausner said.

Fountain of rejuvenation

To be sure, the word “rejuvenation” sounds suspicious, like a conquistador’s quest or a promise made on a bottle of high-priced face cream. Yet rejuvenation is all around us, if you look. Millions of babies are born every year from the aging sperm and egg cells of their parents. Cloning of animals is another example. When Barbra Streisand had her 14-year-old dog cloned, cells from its mouth and stomach were returned to her as two frolicking puppies. These are all examples of cells being reprogrammed from age to youth—exactly the phenomenon companies like Altos want to capture, bottle, and one day sell.

For now, no one has a firm idea what these future treatments could look like. Some say they will be genetic therapies added to people’s DNA; others expect it’s possible to discover chemical pills that do the job. One proponent of the technology, David Sinclair, who runs an aging-research lab at Harvard University, says it could allow people to live much longer than they do today. “I predict one day it will be normal to go to a doctor and get a prescription for a medicine that will take you back a decade,” Sinclair said at the same California event. “There is no reason we couldn’t live 200 years.”

It’s this type of claim that raises so much skepticism. Critics see ballooning hype, runaway egos, and science that’s on uncertain ground. But the doubters this year were drowned out by the sound of stampeding investors. In addition to Altos, whose \$3 billion ranked as possibly the single largest

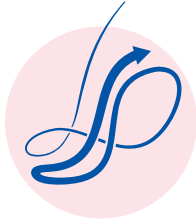


**There
is no
reason we
couldn't
live
200
years.**



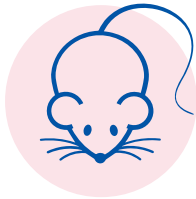
David Sinclair,
Harvard
University

How it works



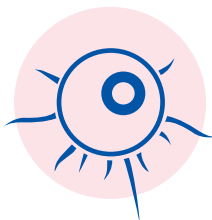
— 1

Four proteins can “reprogram” a skin cell into a youthful stem cell.



— 2

Apply these proteins to mice ... but only in limited doses.



— 3

Try to make the mice young, without causing cancer.



— 4

Test the mice to see if they are healthier or live longer.

fundraising drive in biotech history, the cryptocurrency billionaire Brian Armstrong, the cofounder of Coinbase, helped bring \$105 million into his own reprogramming company, NewLimit, whose mission he says is “radical extension of human health span.” Sam Altman, the startup investor and CEO of OpenAI, personally put \$180 million into Retro Biosciences, which says it wants to extend the average healthy human life by 10 years.

These huge expenditures are being made despite the fact that scientists still disagree on the causes of aging. Indeed, there’s no real consensus on when in life aging even begins. Some say it starts at conception, while others think it’s at birth or after puberty.

But all the unknowns are part of what makes the reprogramming phenomenon so attractive. Klausner admits that the details of why reprogramming works remain a “complete mystery,” but that too helps explain the sudden rush to invest in the idea. If there is a fountain of youth in the genome, the first to locate it could reinvent medicine and revolutionize how we treat the myriad of diseases that plague our old age.

Alchemy project

To get a reality check on Klausner’s lecture, I asked an embryologist and stem-cell specialist, Alfonso Martinez Arias, to watch a recording. Martinez, whose lab is at the Pompeu Fabra University, in Barcelona, wrote back that he had to hold his stomach while he watched, so grandiose were the claims. “He was evangelical about something which, at the moment, is interesting but very preliminary and [on] shaky ground,” says Martinez. Klausner was speaking “as if he had drunk some Kool-Aid.”

Martinez says that to him, Altos is an alchemy project, the kind that medieval rulers once financed in the search for the philosopher’s stone—a substance they believed could turn lead into gold, not to mention cure all disease. Martinez wasn’t entirely negative, though. “There are people at Altos who know how to do science,” he says. And, he notes, even alchemists ended up making valuable discoveries.

The basic technique Altos is exploring is the procedure discovered in 2006 by the Japanese

scientist Shinya Yamanaka, who is now a scientific advisor to the company. The four proteins (now called “Yamanaka factors”) that he and his students identified could cause ordinary cells to turn into potent stem cells, just like those found in embryos. This discovery earned him a Nobel Prize in medicine in 2012.

Initially, Yamanaka’s discovery was employed to reprogram cells from patients to make stem cells, which could then be used to try to manufacture transplantable tissues, retina cells, or neurons. Other scientists wondered what would happen if they introduced Yamanaka’s factors into living animals. In 2013, a Spanish team did exactly that, with gruesome results. The mice sprouted tumors called teratomas, blobs of renegade embryonic tissue.

The problem for these reprogrammed mice was that the process doesn’t just make cells young; it also erases their identity and turns them into embryonic stem cells, which don’t belong in an adult. Joe Betts-Lacroix, the CEO and founder of Retro, says researchers were soon asking a new question: “Is there some way that those two phenomena can be uncoupled so that you can have some of the age wiped away, but not have all your identity wiped away so that you become a pile of stem-cell protoplasm and die?”

In 2016, researchers at the Salk Institute in California, headed by Juan Carlos Izpisua Belmonte, reported that the answer might be yes. They genetically engineered mice afflicted with progeria, a condition that causes extremely rapid aging, so that all their cells would make the Yamanaka factors, but only when they were fed a special supplement in their food. That allowed the scientists to turn on the factors for a limited period—just a few hours at a time. Leave the genes on for too long, and the mice got cancer. But with shorter pulses—a tactic now known as partial reprogramming—they didn’t. What’s more, the mice seemed to become healthier and live a bit longer.

“You rejuvenate cells, but you didn’t lose the identity,” says Klausner, who calls it an “Aha!” moment. “That could be safe. And this has [now] been done with many animals. They don’t get cancer as long as you don’t get past this point.”

Exactly how this partial-reprogramming phenomenon works is now a major focus of Altos and other research organizations. During a meeting held in June at a Maine ski resort, reprogramming scientists described studying individual cells by the tens of thousands—tracking in detail what changes they undergo after they’re exposed to more limited pulses of the Yamanaka factors, or to subsets of them. Researchers from the United Kingdom with connections to Altos reported that they’d made skin cells from a 53-year-old person as youthful as those of someone just out of college. They claimed the “rejuvenation point” was reached after 13 days of exposure to Yamanaka’s factors, but no more.

One way the British team concluded that the cells had become younger was by using an “aging clock.” These are measurements that detect epigenetic modifications to DNA, the chemical marks that determine whether a given gene is on or shut off. (Epigenetic controls are part of what gives every cell its specialized identity; an olfactory neuron in your nose doesn’t need the same genes activated as a liver cell that oozes bile.) Because these markers undergo telltale changes over a lifetime, it’s possible to estimate a person’s age, or that of any animal, within a couple of years by checking just two or three hundred of them.

In part because the clocks are eerily accurate, some researchers now believe aging may be caused primarily by the gradual degradation of the epigenetic code, a little like a compact disc that’s been scratched and skips tracks. It’s an attractive theory, and not least because one thing that reprogramming does reliably is reset these marks; after a little treatment with Yamanaka factors, a cell from a 90-year-old will have the epigenetic profile of one from a teenager.

To Klausner, the fact that cells can regain a youthful epigenetic state is remarkable and likely a gateway to important new biology. “Understanding how cells remember how to be an unscratched CD” could lead to the discovery of “missing codes” regulating the whole process of aging, he thinks.

Other scientists say it’s an open question whether aging clocks measure true rejuvenation, a term they say is already being used too loosely. To Charles Brenner, a senior researcher at the City of Hope

Juan Carlos Izpisua Belmonte, institute director at Altos, claimed he could “slow down aging” in laboratory mice.



National Medical Center, people may even be falling victim to circular reasoning when they celebrate those epigenetic changes. “There isn’t a difference between saying they applied the Yamanaka factors and that they have changed the epigenetic profile, since that is what the factors do,” he says. “They then score their study as a rejuvenation success, but there is no scientific basis for doing that. They still don’t know what the intervention does. People should not be assuming more youthful scores on an epigenetic clock equate to better health or longer life expectancy.”

To answer that question, more researchers are applying bursts of the reprogramming factors to mice in bids to reverse specific diseases, or just to see what happens. In 2020, researchers at Harvard led by Sinclair reported that mice exposed to three reprogramming factors could regenerate their optic nerve and regain sight after it was crushed, something usually only a newborn rodent can do. That result earned them the cover of the journal *Nature* and the headline “Turning Back Time.” Others have claimed that after partial reprogramming, mice perform better on a grip test (they’re hung from tiny bars) and show signs of renewed muscle growth or even improved memory.

So far, many of these individual rejuvenation claims for live mice haven’t been widely replicated by other labs, and some people are skeptical they ever will be. Measuring the relative health of animals or their tissues isn’t necessarily a precise science. And in unblinded studies (where the researchers know which animals were treated), wishful thinking can play a role, perhaps especially if billions in venture capital dollars ride on the result. “Frankly, I doubt the reproducibility of these papers,” says Hiro Nakauchi, a professor of genetics at Stanford University. Nakauchi says he also created mice with Yamanaka factors, but he never saw any sign they got younger. He suspects that some of the most dramatic claims are “timely and catchy” but that the science that went into them is “not very accurate.”


One rejuvenation claim Brenner found troubling this year came from the Salk Institute, in La Jolla, California, which issued a press release saying a group of scientists there (who have since joined

Altos) had been able to “safely and effectively reverse the aging process” in mice. It sounded as if they were describing a drug ready for market, not an exploratory form of genetic engineering. Izpisua Belmonte, the chief researcher involved, who now directs a San Diego research center for Altos, separately claimed he could “slow down aging” in the animals.

In reality, the results were less definitive than advertised. The researchers had not seen tumors, but they had significantly changed the epigenetic age of cells in just two organs: kidneys and skin. And something else about the result jumped out as puzzling to researchers like Brenner, as well as others who reviewed the paper. Despite saying they’d slowed aging, the Salk team didn’t comment on how long the partially reprogrammed mice lived. Some data in their publication suggests that the rodents’ life spans were unremarkable.

Indeed, so far no research group or company has reported normal mice living longer after being exposed to partial reprogramming. And that’s something you might expect them to do, if the alchemy is real. To João Pedro de Magalhães, at the University of Birmingham, the gap in the data is puzzling, since he believes that whether the technology affects life span “is the billion-dollar question, so to speak.” George Daley, a prominent stem-cell biologist who is dean of Harvard Medical School, wrote in response to the Salk paper that “rigorous demonstration of such an effect” was necessary to call reprogramming a true anti-aging intervention.

“Let’s not pretend that the most important thing has happened if it hasn’t,” says Martin Borch Jensen, chief scientist at Gordian Biotechnology and founder of a grant-making organization. “I mean, is there any evidence for your \$3 billion project?”


**Is there
 any
 evidence
 for your
 \$3
 billion
 project?**



Martin Borch Jensen,
 Gordian
 Biotechnology

Disease reversal

When Altos officially launched, in January of 2022, Klausner and other executives strove to distance the company from the concept of life-span extension, even telling reporters that Altos “is not an aging or longevity company.” They’d been stung by suggestions that the project existed to help billionaires cheat death. Instead, in its debut, Altos sought to

Altos seeks to align itself with a concept called health span, which means extending the number of years that people spend in good health as they age.

align itself with a concept called “health span,” which means extending the number of years people spend in good health.

Klausner says reprogramming promises an approach to “disease reversal” that might be applied regardless of how old someone is. If any extension in longevity resulted, it would be only “an accidental consequence” of making people healthier, according to comments made by Hans Bishop, the president of Altos.

In an email, Klausner even said that the company will not try to determine whether reprogramming generally extends life. “We have no intention of ever doing life-span extension studies,” he wrote. He noted that an experiment would be impractical—such a test in humans could take too long. Instead, Altos hopes to carry out “very specific” attempts to reverse certain diseases or disabilities, using familiar frameworks for clinical trials that are accepted by regulators and attractive to large drug companies.

To some observers, like Magalhães, Altos is just trying to position anti-aging technology in a guise that’s credible, even though some of the company’s own scientists, like Izipisua Belmonte, have predicted that people will live to 130. “It is curious psychology,” Magalhães says. “We say we are not trying to cure aging, just make people healthy longer. But I don’t think we should be ashamed about what we are trying to do, which is to slow down aging. And rejuvenation, if we achieve it, would be the best way of doing that.”

Klausner told me he thinks the longevity-versus-health span debate is “a distraction.” The average American lives for around 77 years, which is still decades short of the longest lives on record (the oldest person to live died at 122). That means there are plenty of healthy years to be gained before anyone reaches an unnatural birthday. Nor are gains in average life expectancy unusual—that figure has roughly doubled since 1850, thanks mostly to vaccines, antibiotics, and public health advances.

“There is a lot of room for average life span to increase,” Klausner says, “and that is essentially the goal of all medicine, whether curing cancer or heart disease.” ■

Antonio Regalado is the senior editor for biomedicine at MIT Technology Review.

DONATING ONE'S BODY TO SCIENCE CAN ADVANCE MEDICINE AND BENEFIT THE LIVING—IF HANDLED PROPERLY.



TAKING CARE OF THE DEAD

by **ABBY OHLHEISER**

Photographs by
Mike Belleme

Visits to a body farm and a university anatomy lab show what happens when you donate your body to science.

“You can learn about somebody’s life history through their skeleton, so you can understand things that happen to them over the period of their life.”

Rebecca George doesn’t mind the vultures. They remind her of toddlers as they rustle their feathers in annoyance when she opens the gate of the Western Carolina University body farm early one July morning. Her arrival has interrupted their breakfast. George studies human decomposition, and part of decomposing is becoming food. Scavengers are welcome.

The birds complain from the trees that surround the body farm as George, a forensic anthropologist, begins her main task of the day: placing the body of a donor, whom we’ll call Donor X, in the Forensic Osteology Research Station—known as the FOREST. The enclosure sits on a steep incline in North Carolina’s temperate rainforest, surrounded by two layers of protective fencing. This is Enclosure One, where donors decompose naturally above ground. Just on the other side of the clearing is Enclosure Two, where researchers study bodies that have been buried in soil. She is the facility’s curator, a member of a small team of forensic anthropologists and university students who monitor the donors—sometimes for years—as they become nothing but bones.

George places Donor X on their back just inside the enclosure gates, hands at their sides. Unless donors are part of a specific study requiring clothes, they’re laid out “in their birthday suit.” Clothing slows decomposition. She sticks a little yellow flag next to the body with an ID number and the date. Another donor is nearby, one skeletonized hand gently resting on a small rock, head tilted to the right, as if they were sleeping.

Donor X’s next of kin chose for them to be laid out here in the FOREST upon their death. In the US, about 20,000 people or their families donate their bodies to scientific research and education each year. They do it because they want to make their deaths meaningful, or because they’re disenchanted with the traditional death industry. People can become organ donors—offering up, at their death, organs suitable for transplant into living people—by checking a box on their driver’s license in the US. But the practice of whole-body donation is less widely discussed.

Body donation can also be cheaper than conventional cremation or burial. Some donation programs will pay for the cost of transporting a donor within a certain distance and, if the program is one that promises to eventually return remains to the family, for cremation. At the FOREST, the donors’ remains become permanent residents in the university’s forensic anthropology archives.

Whatever the reason someone chooses to donate, the decision becomes a gift. Health care needs death care; the bodies of the dead have long taught and trained the living. Many donor bodies go to medical schools, where students use them to learn anatomy and practice procedures. Others, like Donor X, go to university research facilities, or any of several private companies in the US that take body donations. Western Carolina’s FOREST, founded in 2003, is the second-oldest body farm in the US. A much larger facility at the University of Tennessee in Knoxville, opened in 1981, is the oldest. These are places where watchful caretakers know that the

dead and the living are deeply connected, and the way you treat the first reflects how you treat the second.

I visited the FOREST and another facility, the University of Maryland School of Medicine’s anatomy lab, to understand what happens when body donation works as intended.

Adam Puche walked me into the anatomy lab at the University of Maryland School of Medicine, where he is a professor and vice chair of the anatomy and neurobiology department, just as a class was wrapping up. Two students zipped a bag around their donor as they quietly cleaned their work station, and then draped a light blue cloth over the table.

Maryland has a highly regulated process for body donation, governed by a central anatomy board in the state’s health department; Puche is its chair. This particular lab handles about 4,000 bodies a year. Here, donors become patients for doctors in training. When a body arrives at the anatomy board, the individual is issued a tracking number. Then an RFID chip is implanted in one shoulder—a step unique to Maryland’s state program.

The lab is secured both by ID badges and by Puche’s own rigorous standards. The timing of my visit was carefully planned to minimize the impact on students. When I asked to take a reference photo of a cabinet of wet specimens showing livers, gall bladders, and other organs from donors with specific medical conditions, Puche politely shook his head. The work of protecting donors’ dignity extends to those represented in the museum, who lived a century ago. This is what he’s trying to teach future physicians, who are supposed to treat body donors as they would a living patient. In addition to what’s on the intake form, students at Maryland are expected to keep charts on donors. As they discover new conditions a patient may have—a cyst, a past broken bone, a previous surgery—they note it. Students are required to follow HIPAA rules when discussing their donors outside the lab.

“These are going to be physicians from day one,” said Puche. “We need them to be exercising the appropriate language choices, appropriate actions. So not only do I firmly believe in what I’m telling you as the right way to do it; it’s important for all faculty members to continuously and consistently display that to our students.”

Puche’s lab will soon be renovated to realize his vision for a space that reflects the working conditions of future doctors. The ’70s-chic fluorescent lighting will be supplemented with the same LED light systems and data access panels seen in operating rooms. As augmented-reality technology becomes more integrated into surgeries and other medical procedures, he expects, students might soon be able to see all the diagrams and instructions they need overlaid on their donor virtually.

I asked Puche if he sees a future in which technology eliminates the need for donors. He believes in technology’s potential to improve care, and he has run experiments offering VR training for medical students. In his opinion, however, none of these tools can replace the experience of working with a donor, so long as the living have real bodies too.





Students scrape tissue from bones while volunteering in Rebecca George's lab at Western Carolina University.



The process of donating one's body begins with research, paperwork, and sometimes difficult conversations with family members. Funerals are the territory of the grieving, and it's not always easy for loved ones to grieve without a body or its ashes. Some are aware of body donation because they know someone who made that choice. Others might search Google, looking for more affordable alternatives to burial.

Jeff Battersby, a 61-year-old who lives in Beacon, New York, learned about body farms from a 2017 podcast episode about a facility in Texas. Now he's considering donating his body to one. "I'm not really fond of preserving a body for a million years in a casket, in the ground somewhere," he says. "And I'm not really interested in going up in a puff of smoke. I just wanted to find and think about a way that was useful and giving."

Battersby downloaded the paperwork to become a donor to the research facility at the University of Tennessee in Knoxville. But donating one's body is a big decision to make. When we spoke last, he hadn't yet sent in the forms or talked to his family about it.

Though dead bodies have been essential to medicine and research for centuries, consensual body donation, through programs like those at FOREST and Maryland, is relatively new. In the US, demand for bodies grew substantially in the mid-19th century as medical schools moved from having one person perform a dissection for an audience to providing each student with a hands-on lesson. That surge in demand drove body snatchers to steal bodies from graves to sell to medical schools. The bodies of the poor, the mentally ill, and people of color were especially vulnerable.

Things are very different now, thanks to new regulations and a better understanding of consent. But this grim history is still reflected in modern institutions. Until 2020, the University of Pennsylvania's Penn Museum in Philadelphia had skulls on display that belonged to formerly enslaved people. And plenty can still go wrong even when donation is consensual. In August 2022, a Pennsylvania man was arrested for allegedly buying and selling human body parts via Facebook Messenger. At least some of the remains were initially donated to the University of Arkansas for Medical Sciences. After the donor's gift was used by the university, the remains were sent to be cremated at a non-university funeral home. There, the university said, a mortuary worker stole them.

Today, there's no single federal regulation, registry, or tracking program that handles body and body part donations for research in the US. The American Association of Tissue Banks offers optional accreditation for these programs, but it's not required by law. Instead, programs are largely governed by each state's version of the Uniform Anatomical Gift Act, which contains provisions for promoting organ and tissue donation and outlines how people can consent to donate their organs or entire body to science. Whole-body donors must navigate these systems to decide where to donate. Some programs treat donors better than others do. And while there has been progress, body donation programs can still prey on those struggling to afford a conventional cremation or burial.

Nicholas Passalacqua directs Western Carolina's forensic anthropology program, which accepts about two dozen donated bodies each year.

A major Reuters investigation in 2017 scrutinized “body brokers”—for-profit companies that accept donations and then sell partial and whole bodies to organizations engaged in training or research. When pitching potential donors and their families, these companies often emphasized the cost savings available, offering free cremation after the donated body served its purpose. As part of their reporting, Reuters was able to legally purchase body parts from one of these companies. They found out that the donor, from a low-income family, had been persuaded to donate for this reason. The family had no idea his body parts would be sold off.

That investigation prompted attempts to pass revised federal regulation in the US to oversee the industry, including an updated bill introduced into the House of Representatives last year that would require organizations accepting body donations to register with the US Department of Health and Human Services and follow uniform guidelines. That bill has not yet passed.

Not every body is welcome in every donation program. Most centers exclude those with certain communicable diseases, like HIV/AIDS or hepatitis B or C. Since 2020, many centers have also rejected donors who were positive for covid-19 when they died. And donation programs are not generally interested in waging custody battles with the next of kin, should that person decide not to honor a loved one's intent. Some programs will not accept organ donors or bodies that have been previously autopsied.

Many facilities, including Western Carolina's, also have weight limits for donors. Bodies there are carried into the facility by hand or on a gurney, often along steep and sometimes slick paths. Under university safety standards, employees can only move bodies that weigh 250 pounds or less. Without funding for the equipment needed to safely place larger donors, the FOREST is stuck with the limitation for now.

“It's very frustrating, because we could get so many more donors if we didn't have that weight limit,” George says.

Some dissection-based programs will not accept larger donors because dissecting a body that is carrying a lot of fat is more difficult and deemed less pleasant for students. But Maryland, which accepts thousands more donors a year than the 20 or so Western Carolina can handle, has never turned a body down for weight reasons. While Maryland's program might not give a very large body to students performing their very first dissection, Puche says, “surgeons are going to need to work with somebody who may be 300 or 400 pounds. We expect our surgeons in training to work with our patient as the patient is.” As important as it is for them to get this practice, though, he believes his university is unusual in not exercising a weight limit.

Programs collect different kinds of information about donors. Western Carolina's intake form asks for both biological sex and culturally expressed gender. Though Maryland's form only asks for sex, Puche emphasizes that the anatomy board recognizes the gender identity of donors, and that living donors can use the “sex”





field on the form to write in whatever best describes their identity. Students are expected to respect the identities and pronouns of the donors they work with.

But in some cases, the larger medical system's failures in treating all patients with respect may prompt people to question the value of donating. Liam Hartle, a 30-year-old from Albany, New York, who has an autoimmune disorder, has thought seriously about donating his body. Maybe someone could learn something about the condition from studying him.

"There's part of me that's like, donating to science would be a really good idea," he says. "But also? I'm a trans dude who hasn't done any hormone therapy yet, because my husband and I are going to try to have a kid. Even if I had undergone any sort of physical transition, I don't trust medical science to handle that delicately."

After my time in Puche's autopsy lab, I drove 10 hours from Maryland to North Carolina to see how Donor X was progressing. Nicholas Passalacqua, the director of Western Carolina University's forensic anthropology program, approved my visit; the FOREST is not open to the public. As we pulled up to the gated facility, he told me that when he does bring visitors to the site—other researchers, students, or journalists like me, for instance—they often expect it to be disgusting, or gory, or terrifying.

Passalacqua and I stepped inside the enclosure, where George was already at work, training student volunteers. The group crouched over the skeleton of the donor who had been inside the enclosure the longest, since 2020. Patches of weeds had sprung up between donors but stayed away from the bodies themselves. Contrary to popular belief, decomposing bodies aren't particularly good fertilizer in the short term—the fluids they release can inhibit plant growth.

As I adjusted to the smell—a sweet, rotting-fruit-like scent that crept into the back of my throat and stayed long after I left campus—I asked Passalacqua to tell me how to look at a body like a forensic anthropologist. His first lesson was that as donors decompose, they fill with life.

As we walked further into the enclosure, we came to a donor covered in sprinkle-size maggots: blowfly babies, laid by their mothers at the eyes, groin, and mouth. Those parts of the body, which liquefy first, are like baby food for the baby flies. As they work through the body, the busy bugs and gut bacteria leave the skin on the torso alone. The skin toughens on the ribs and becomes a weather and sun barrier, a little shaded tent.

"It's a whole little microbiome, right?" Passalacqua said. "Where you have a body decomposing, you have insects eating, you have other things eating those insects, you have animals coming to eat the tissue, you have other animals coming to eat those animals."

Donors who had been at the FOREST longer, though, were now more bone than flesh. "You can learn about somebody's life history through their skeleton, so you can understand things that





After a donor's body decomposes outside, the skeleton is cleaned and stored in Western Carolina University's collection.



happen to them over the period of their life and how that manifests," Passalacqua said.

Take ribs, for instance. George stood near a donor, almost entirely skeletonized, whose ribs were cracked. One might surmise that these broken ribs were a clue to how the person had died. But she says the bones were broken in the enclosure, from a vulture sitting on them. She wouldn't have believed it herself if the whole thing hadn't been captured on camera. Then she pointed to a body up on the plateau, at the far end of the enclosure. That donor, she said, did break their ribs when they died. The breaks looked totally different, the fractures more jagged. That's because, George said, the breaks occurred in living bone, not after death when the material is more brittle. In a third body, Passalacqua pointed out a small spike on the rib. That, he said, was also a rib fracture, but one that had healed while the donor was alive.

Chronic illnesses and some other diseases can manifest in the bones. Tuberculosis can spread there, causing lesions. Forensic anthropologists can estimate a deceased young person's age by understanding how the skeleton changes over time. Older adults, too, might have distinctive markers of age, like bone loss. But this work is difficult, and there's still a lot scientists don't know. Donors like these help them learn more.

George and Passalacqua's job is to teach students, along with the law enforcement workers who occasionally train there, how to learn what you can from a body. Oftentimes, their first step is to figure out if a bone is human. Passalacqua regularly receives texts from local law enforcement asking about bones they've found—a partial skeleton of a bear paw looks shockingly like a person's hand.

One of the hardest things for forensic anthropologists to do is also one of the most essential: estimating the time that has passed since death. "There are just so many variables that are really difficult to account for," says Passalacqua. A reputable forensic anthropologist will rarely be able to say, for instance, that a body has been dead for exactly three weeks. More likely is a range of, say, one week to two months. That's not as useful to law enforcement officers who are trying to solve a crime.

By the time I arrived, Donor X was already in advanced decay, but every day, this donor will teach living people something. When there's little left on the bone, students will carefully remove the body from the FOREST and bring it to the lab. The bones will be cleaned by hand, and perhaps gently simmered to remove the last bits of tissue. They will be laid out and examined. And then they will be packed up, the delicate pieces placed in cheesecloth bags, and stored in the university's collection, labeled in identical cardboard boxes.

But for now, Donor X remains in place, slowly becoming a unique microbiome. Dense trees filter the sunlight. The vultures aren't there that morning, so as we walk and the students quietly examine a donor's bones, the only other sounds we hear are the calls of cicadas. ■

In a workshop in Rotterdam in the Netherlands, Philip Nitschke—“Dr. Death” or “the Elon Musk of assisted suicide” to some—is overseeing the last few rounds of testing on his new Sarco machine before shipping it to Switzerland, where he says its first user is waiting.

This is the third prototype that Nitschke’s nonprofit, Exit International, has 3D-printed and wired up. Number one has been exhibited in Germany and Poland. “Number two was a disaster,” he says. Now he’s ironed out the manufacturing errors and is ready to launch: “This is the one that will be used.”

A coffin-size pod with Star Trek stylings, the Sarco is the culmination of Nitschke’s 25-year campaign to “demedicalize death” through technology. Sealed inside the machine, a person who has chosen to die must answer three questions: Who are you? Where are you? And do you know what will happen when you press that button?

Here’s what will happen: The Sarco will fill with nitrogen gas. Its occupant will pass out in less than a minute and die by asphyxiation in around five.

A recording of that short, final interview will then be handed over to the Swiss authorities. Nitschke has not approached the Swiss government for approval, but Switzerland is one of a handful of countries that have legalized assisted suicide. It is permitted as long as people who wish to die perform the final act themselves.

Nitschke wants to make assisted suicide as unassisted as possible, giving people who have chosen to kill themselves autonomy, and thus dignity, in their final moments. “You really don’t need a doctor to die,” he says.

Because the Sarco uses nitrogen, a widely available gas, rather than the barbiturates that are typically used in euthanasia clinics, it does not require a physician to administer an injection or sign off on lethal drugs.

At least that’s the idea. Nitschke has not yet been able to sidestep the medical establishment fully. Switzerland requires that candidates for euthanasia demonstrate mental capacity, Nitschke says, which is typically assessed by a psychiatrist. “There’s still a belief that if a person is asking to die, they’ve got some sort of undiagnosed mental illness,” he says. “That it’s not rational for a person to seek death.”



Life, death, and

AI is being asked to help make life-or-death decisions. This mustn’t be an excuse to turn away from painful choices.



automation

Above:
With its Sarco
pod, the nonprofit
Exit International
aims to make
assisted suicide
as unassisted as
possible.

By Will Douglas Heaven

Even when AI seems accurate, scholars and regulators alike call for caution.

He believes he has a solution, however. Exit International is working on an algorithm that Nitschke hopes will allow people to perform a kind of psychiatric self-assessment on a computer.

In theory, if a person passed this online test, the program would provide a four-digit code to activate the Sarco. “That’s the goal,” says Nitschke. “Having said all that, the project is proving very difficult.”

Nitschke’s mission may seem extreme—even outrageous—to some. And his belief in the power of algorithms may prove to be overblown. But he is not the only one looking to involve technology, and AI in particular, in life-or-death decisions.

Yet where Nitschke sees AI as a way to empower individuals to make the ultimate choice by themselves, others wonder if AI can help relieve humans from the burden of such choices. AI is already being used to triage and treat patients across a growing number of health-care fields. As algorithms become an increasingly important part of care, we must ensure that their role is limited to medical decisions, not moral ones.

Medical care is a limited resource. Patients must wait for appointments to get tests or treatment. Those in need of organ transplants must wait for suitable hearts or kidneys. Vaccines must be rolled out first to the most vulnerable (in countries that have them). And during the worst of the pandemic, when hospitals faced a shortage of beds and ventilators, doctors had to make snap decisions about who would receive immediate care and who would not—with tragic consequences.

The covid crisis brought the need for such choices into harsh focus—and led many to wonder if algorithms could help. Hospitals around the world bought new or co-opted existing AI tools to assist with triage. Some hospitals in the UK that had been exploring the use of AI tools to screen chest x-rays jumped on those tools as a fast, cheap way to identify the most severe covid cases. Suppliers of this tech, such as Qure.ai, based in Mumbai, India, and Lunit, based in Seoul, Korea, took on contracts in Europe, the US, and Africa. Diagnostic Robotics, an Israeli firm that supplies AI-based triage tools to hospitals in Israel, India, and the US, has said it saw a sevenfold jump in demand for its technology in the first year of the pandemic. Business in health-care AI has been booming ever since.

This rush to automate raises big questions with no easy answers. What kinds of decision is it appropriate to use an algorithm to make? How should these algorithms be built? And who gets a say in how they work?

Rhema Vaithianathan, the director of the Centre for Social Data Analytics and a professor at the Auckland University of Technology in New Zealand, who focuses on tech in health and welfare, thinks it is right that people are asking AI to help make big decisions. “We should be addressing problems that clinicians find really hard,” she says.

One of the projects she is working on involves a teen mental-health service, where young people are diagnosed and treated for self-harming behaviors. There is high demand for the clinic, and so it needs to maintain

a high turnover, discharging patients as soon as possible so that more can be brought in.

Doctors face the difficult choice between keeping existing patients in care and treating new ones. “Clinicians don’t discharge people because they’re super scared of them self-harming,” says Vaithianathan. “That’s their nightmare scenario.”

Vaithianathan and her colleagues have tried to develop a machine-learning model that can predict which patients are most at risk of future self-harming behavior and which are not, using a wide range of data, including health records and demographic information, to give doctors an additional resource in their decision-making. “I’m always looking for those cases where a clinician is struggling and would appreciate an algorithm,” she says.

The project is in its early stages, but so far the researchers have found that there may not be enough data to train a model that can make accurate predictions. They will keep trying. The model does not have to be perfect to help doctors, Vaithianathan says.

They are not the only team trying to predict the risk of discharging patients. A review published in 2021 highlighted 43 studies by researchers claiming to use machine-learning models to predict whether patients will be readmitted or die after they leave hospitals in the US. None were accurate enough for clinical use, but the authors look forward to a time when such models “improve quality of care and reduce health-care costs.”

And yet even when AI seems accurate, scholars and regulators alike call for caution. For one thing, the data that algorithms follow and the way they follow it are human artifacts, riddled with prejudice. Health data is overpopulated by people who are white and male, for example, which skews its predictive power. And the models offer a veneer of objectivity that can lead people to pass the buck on ethical decisions, trusting the machine rather than questioning its output.

This ongoing problem is a theme in David Robinson’s new book, *Voices in the Code*, about the democratization of AI. Robinson, a visiting scholar at the Social Science Matrix at the University of California, Berkeley, and a member of the faculty of Apple University, tells the story of Belding Scribner. In 1960 Scribner, a nephrologist in Seattle, inserted a short Teflon tube known as a shunt into some of his patients’ arms to prevent their blood from clotting while they underwent dialysis treatment. The innovation allowed people with kidney disease to stay on dialysis indefinitely, transforming kidney failure from a fatal condition into a long-term illness.

When word got out, Scribner was inundated with requests for treatment. But he could not take everyone. Whom should he help and whom should he turn away? He soon realized that this wasn’t a medical decision but an ethical one. He set up a

MIT Technology Review

Global Insights Panel

Shape the future of tech.

Take part in original research and gain valuable business insights into today's most important technology trends.

Join the Global Insights Panel.

- Participate in research programs and surveys.
- Receive exclusive news in our twice-monthly newsletters.
- Access our quarterly Global Insights reports, prepared by MIT Technology Review Insights.
- Network with like-minded peers in person or virtually.

→ Apply at technologyreview.com/globalpanel

“I didn’t want to sit there and give the injection. If you want it, you press the button.”

committee of laypeople to decide. Of course, their choices weren’t perfect. The prejudices at the time led the committee to favor married men with jobs and families, for example.

The way Robinson tells it, the lesson we should take from Scribner’s work is that certain processes—bureaucratic, technical, and algorithmic—can make difficult questions seem neutral and objective. They can obscure the moral aspects of a choice—and the sometimes awful consequences.

“Bureaucracy itself can serve as a way of converting hard moral problems into boring technical ones,” Robinson writes. This phenomenon predates computers, he says, “but software-based systems can accelerate and amplify this trend. Quantification can be a moral anesthetic, and computers make that anesthetic easier than ever to administer.”

Whatever the process, we need to let that moral anesthetic wear off and examine the painful implications of the decision at hand. For Scribner, that meant asking an open panel of laypeople—instead of a group of ostensibly objective doctors meeting behind closed doors—whom to save. Today, it could mean asking for high-stakes algorithms to be audited. For now, the auditing of algorithms by independent parties is more wish-list item than standard practice. But, again using the example of kidney disease, Robinson shows how it can be done.

By the 2000s, an algorithm had been developed in the US to identify recipients for donated kidneys. But some people were unhappy with how the algorithm had been designed. In 2007, Clive Grawe, a kidney transplant candidate from Los Angeles, told a room full of medical experts that their algorithm was biased against older people like him. The algorithm had been designed to allocate kidneys in a way that maximized years of life saved. This favored younger, wealthier, and whiter patients, Grawe and other patients argued.

Such bias in algorithms is common. What’s less common is for the designers of those algorithms to agree that there is a problem. After years of consultation with laypeople like Grawe, the designers found a less biased way to maximize the number of years saved—by, among other things, considering overall health in addition to age. One key change was that the majority of donors, who are often people who have died young, would no longer be matched only to recipients in the same age bracket. Some of those kidneys could now go to older people if they were otherwise healthy. As with Scribner’s committee, the algorithm still wouldn’t make decisions that everyone would agree with. But the process by which it was developed is harder to fault.

Nitschke, too, is asking hard questions.

A former doctor who burned his medical license after a years-long legal dispute with the Australian Medical Board, Nitschke has the distinction of being the first person to legally administer a voluntary lethal injection to another human. In the nine months

between July 1996, when the Northern Territory of Australia brought in a law that legalized euthanasia, and March 1997, when Australia’s federal government overturned it, Nitschke helped four of his patients to kill themselves.

The first, a 66-year-old carpenter named Bob Dent, who had suffered from prostate cancer for five years, explained his decision in an open letter: “If I were to keep a pet animal in the same condition I am in, I would be prosecuted.”

Nitschke wanted to support his patients’ decisions. Even so, he was uncomfortable with the role they were asking him to play. So he made a machine to take his place. “I didn’t want to sit there and give the injection,” he says. “If you want it, you press the button.”

The machine wasn’t much to look at: it was essentially a laptop hooked up to a syringe. But it achieved its purpose. The Sarco is an iteration of that original device, which was later acquired by the Science Museum in London. Nitschke hopes an algorithm that can carry out a psychiatric assessment will be the next step.

But there’s a good chance those hopes will be dashed. Creating a program that can assess someone’s mental health is an unsolved problem—and a controversial one. As Nitschke himself notes, doctors do not agree on what it means for a person of sound mind to choose to die. “You can get a dozen different answers from a dozen different psychiatrists,” he says. In other words, there is no common ground on which an algorithm could even be built.

But that’s not the takeaway here. Like Scribner, Nitschke is asking what counts as a medical decision, what counts as an ethical one, and who gets to choose. Scribner thought that laypeople—representing society as a whole—should choose who received dialysis, because when patients have more or less equal chances of survival, who lives and who dies is no longer a technical question. As Robinson describes it, society must be responsible for such decisions, although the process can still be encoded in an algorithm if it’s done inclusively and transparently. For Nitschke, assisted suicide is also an ethical decision, one that individuals must make for themselves. The Sarco, and the theoretical algorithm he imagines, would only protect their ability to do so.

AI will become increasingly useful, perhaps essential, as populations boom and resources stretch. Yet the real work will be acknowledging the awfulness and arbitrariness of many of the decisions AI will be called on to make. And that’s on us.

For Robinson, devising algorithms is a bit like legislation: “In a certain light, the question of how best to make software code that will govern people is just a special case of how best to make laws. People disagree about the merits of different ways of making high-stakes software, just as they disagree about the merits of different ways of making laws.” And it is people—in the broadest sense—who are ultimately responsible for the laws we have. ■

But wait, there's more. Lots more.

You're already a subscriber.
Activate your account and
start enjoying:

- Unlimited web access
- Exclusive digital stories
- Access to 120+ years of
publication archives

→ technologyreview.com/subonly

Innovation that truly serves us all is in scarce supply. That's a problem.

By Shannon Vallor

Tech's downward spiral

On a recent evening, I sat at home scrolling through my Twitter feed, which—since I'm a philosopher who studies AI and data—is always filled with the latest tech news. After a while, I noticed a heaviness growing in the pit of my stomach, that telltale sign that you are *not having a good time*. But why? I wasn't reading news about politics, or the climate crisis, or the pandemic—the usual sources of doomscrolling ennui. I stopped and reflected for a moment. What had I just been looking at?

I had blinked at the aesthetic poverty of the most recent pitch for Meta's Horizon Worlds VR game, featuring Mark Zuckerberg's dead-eyed cartoon avatar against a visual background that one Twitter wag charitably compared to “the painted

walls of an abandoned day-care center.” I had let out a quiet sigh at the news of *Ring Nation*, an Amazon-produced TV show featuring “light-hearted viral content” captured from the Ring surveillance empire. I had clenched my jaw at a screenshot of the Stable Diffusion text-to-image model offering up AI artworks in the styles of dozens of unpaid human artists, whose collective labor had been poured into the model's training data, ground up, and spit back out.

I recognized the feeling and I knew its name. It was resignation—that feeling of being stuck in a place you don't want to be but can't leave. I was struck by the irony that I studied technology my whole life in order to avoid this kind of feeling. Tech used to be my happy place.

Naturally, I poured my emotion into a tweetstorm:



Shannon Vallor
@ShannonVallor

The saddest thing for me about modern tech's spiral into user manipulation and surveillance is how it has just slowly killed off the joy that people like me used to feel about new tech. Every product Meta or Amazon announces makes the future seem bleaker and grayer.



Shannon Vallor
@ShannonVallor

It used to be the opposite. Tech was one of the things I loved most. I still remember the feeling when I rode the first BART trains in SF. When I saw my first Concorde my little heart exploded. My Commodore PET. The last time tech made me truly gleeful was these glories.



Shannon Vallor
@ShannonVallor

What will it take for us to get that feeling back? I don't think it's just my nostalgia, is it? There's no longer anything being promised to us by tech companies that we actually need or asked for. Just more monitoring, more nudging, more draining of our data, our time, our joy.

I struck a nerve. As my notifications started blowing up and thousands of replies and retweets started pouring in, the initial dopamine reward for virality gave way to a deeper sadness. A *lot* of people were sitting with that same heavy feeling in their stomach.

Still, there was catharsis in reading so many others give voice to it.

Something is missing from our lives, and from our technology. Its absence is feeding a growing unease being voiced by many who work in tech or study it. It's what drives the new generation of PhD and postdoctoral researchers I work with at the University of Edinburgh, who are drawing together knowledge from across the technical arts, sciences, and humanistic disciplines to try to figure out what's gone awry with our tech ecosystem and how to fix it. To do that, we have to understand how and why the priorities in that ecosystem have changed.

The goal of consumer tech development used to be pretty simple: design and build something of value to people, giving them a reason to buy it. A new refrigerator is shiny, cuts down on my energy bills, makes cool-looking ice cubes. So I buy it. Done. A Roomba promises to vacuum the cat hair from under my sofa while I take a nap. Sold! But this vision of tech is increasingly outdated. It's not enough for a refrigerator to keep food cold; today's version offers cameras and sensors that can monitor how and what I'm eating, while the Roomba can now send a map of my house to Amazon.

The issue here goes far beyond the obvious privacy risks. It's a sea change in the entire model for innovation and the incentives that drive it. Why settle for a single profit-taking transaction for the company when you can instead design a product that will extract a monetizable data stream from every buyer, returning revenue to the company for years? Once you've captured that data stream, you'll protect it, even to the disadvantage of your customer. After all, if you buy up enough of the market, you can well afford to endure your customers' anger and frustration. Just ask Mark Zuckerberg.

It's not just consumer tech and social media platforms that have made this shift. The large ag-tech brand John Deere, for

example, formerly beloved by its customers, is fighting a "right to repair" movement driven by farmers angry at being forbidden to fix their own machines, lest they disturb the proprietary software sending high-value data on the farmers' land and crops back to the manufacturer. As more than one commenter on my Twitter thread noted, today in tech *we* are the product, not the prime beneficiary. The mechanical devices that used to be the product are increasingly just the middlemen.

There's also a shift in who tech innovations today are for. Several respondents objected to my thread by drawing atten-

Technologists' apparent loss of interest in humane innovation is depleting our collective faith in our own powers of invention.

tion to today's vibrant market in new tech for "geeks" and "nerds"—Raspberry Pis, open-source software tools, programmable robots. As great as many of these are for those with the time, skills, and interest to put them to use, they are tools made for a narrow audience. The thrill of seeing genuine innovation in biomedical technology, such as mRNA vaccines, is likewise dampened when we see the benefits concentrated in the wealthiest countries—the ones already best served by tech.

Of course, new technology remains a source of joy and excitement in many places that have historically been denied an equitable share of its comforts. But innovation used to promise us much more than new devices and apps. Engineering and inventing were once professions primarily oriented toward creating more livable infrastructure, rather than disposable *stuff*.

Vital technologies like roads, power grids, sewers, and transit systems used to be a central part of the engineering enterprise in the US. Nowadays, we treat them as taxpayer burdens, and our best minds and resources are funneled instead into data-hungry consumer devices and apps. If the US is any indicator of the trajectory of global technology development, then deep trouble lies ahead for us all, because we have clearly lost the plot.

The fact is, the visible focus of tech culture is no longer on expanding the frontiers of humane innovation—innovation that serves us all. Even space travel has lost its humanistic vision; today's frontier is luxury space tourism and billionaires selling credulous investors on fantasies of escape to Mars. With 8 billion people teetering on the precipice of global environmental destruction, we can't afford a world where the core mission of new tech appears to be "Take the money and run."

If we continue to turn away from humane applications of tech, we risk

feeding a runaway feedback loop that drains our collective will to reinvest in their expansion. The danger is not only that today's technology fails to be directed to our most urgent civilizational needs. It's that technologists' apparent loss of interest in humane innovation is depleting our collective faith in our own powers of invention.

When it stays true to its deepest roots, technology is still driven by a moral impulse: the impulse to construct places, tools, and techniques that can help humans not only survive but flourish together. Of course, that impulse is easily joined to, or pushed aside by, others: the impulses to dominate, exterminate, immiserate, surveil, and control.

But those darker motivations aren't at the heart of our technological capacity as a species. And we can't let them define the modern technological order. Because if technology loses its association with shared joy and comfort, we risk becoming alienated from one of the most fundamental ways we care for the world and one another. ■

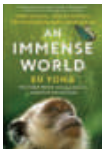
Shannon Vallor is the Baillie Gifford Professor of Ethics of Data and Artificial Intelligence at the University of Edinburgh and director of the Centre for Technomoral Futures in the Edinburgh Futures Institute.



Three authors argue that we should spend more time understanding the creatures around us.

By Matthew Ponsford
Illustrations by Ari Liloan

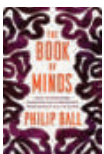
Inside the minds of animals



An Immense World:
How Animal Senses Reveal
the Hidden Realms Around Us
Ed Yong



Sentient:
What Animals Reveal About
Our Senses
Jackie Higgins



The Book of Minds:
How to Understand Ourselves
and Other Beings, from Animals
to Aliens
Philip Ball

The emerald jewel wasp's unusual arrival into the world—bursting from the body of a zombified cockroach it has eaten from the inside—ranks among nature's most gruesome miracles. To give her larvae the best start in life, the mother wasp, an inch-long parasite clad in oil-slick iridescent armor, attacks her prey, spearing it once with her two-millimeter stinger and injecting sedative chemicals into the roach's thorax. She stabs a second time, into her victim's head, carving through muscle and digestive tubes to inject a potion of venom in the exact location of the cockroach's tiny brain.

This turns her victim into an obedient pawn. After biting off the tip of the roach's antenna, the expectant mother can lead the insect like a dog on a leash. Somewhere secluded, she lays an egg on its leg, leaving her young with a passive but still-living mound of flesh, twice her own size, to feed on. In time, the larva will pupate inside it, erupting through its

exoskeleton when mature, ready to repeat this grisly process for another generation. When Charles Darwin observed the climactic eruption of wasp from flesh, it was, he wrote, enough to make him question the existence of a loving God. Nevertheless, he couldn't help but admire the intricacy of this reproductive horror show.

Today, as journalist Ed Yong details in his remarkable new survey of animal perception, *An Immense World: How Animal Senses Reveal the Hidden Realms Around Us*, we can see far deeper into the process than Darwin could have dreamed. View the jewel wasp's stinger under an electron microscope, and you'll find that it is dotted with minute bumps and pits. These are mechanoreceptor cells sensitive to the finest details of touch and texture, and chemoreceptors that sense smell or taste. While the precise purpose of the stinger's smell receptors is still to be understood, tests have shown that the mechanoreceptors make it a precisely calibrated measuring instrument. As the

mother wasp plunges her stinger into the head, “she can detect the distinctive feel of a roach’s brain.”

Yong’s vivid journeys into the sensory realms of other species have given *An Immense World* pride of place among a growing pile of books detailing the rich inner worlds of animals, including *Sentient: What Animals Reveal About Our Senses*, by Jackie Higgins, and *The Book of Minds: How to Understand Ourselves and Other Beings, from Animals to Aliens*, by Philip Ball.

More than ever, we feel a duty and desire to extend empathy to our nonhuman neighbors. In the last three years, more than 30 countries have formally recognized other animals—including gorillas, lobsters, crows, and octopuses—as sentient beings. Yong, Higgins, and Ball together capture what has led to these developments: a booming field of experimental research challenging the long-standing view that animals are neither conscious nor cognitively complex. Western science once treated animals as little more than automata, guided by instinct and hardwiring. But in recent decades researchers have sought to understand complex behavioral phenomena like bee language, vampire bat altruism, and crow ingenuity. The San Francisco-based Earth Species Project, backed by LinkedIn cofounder Reid Hoffman, believes it can take things a step further by decoding patterns in dolphins’ squeaks and pigs’ grunts to create a trans-species translation tool. Talking to animals, once the preserve of animist myths or Dr. Dolittle-like children’s stories, is a prospect that many in tech now suggest is achievable, allowing members of other species to communicate their lives, experiences, and worldviews.

What would they say? The question philosopher Thomas Nagel posed in his famous 1974 paper on consciousness—“What Is It Like to Be a Bat?”—still lingers. Yong, Higgins, and Ball each claim to chisel away at Nagel’s argument that the experiences of such animals are simply beyond our reach. But though all three assemble troves of fascinating research that provides windows into the lives of

In the last three years, more than 30 countries have formally recognized other animals—including gorillas, lobsters, crows, and octopuses—as sentient beings.

animals, we’re left asking how close we really are to bridging the species divide.

In 1909, zoologist Jakob von Uexküll made the then-radical proposal that each animal possesses *Umwelt*, its own perceptual world, constructed from the information that its senses provide. The *Umwelt* of the eyeless, body-heat-sensing tick is quite different from that of the blue whale, which can tune in to eclectic water-transmitted signals and extremely low, infrasonic songs that carry thousands of miles. In *An Immense World*, Yong follows von Uexküll’s framework: he sets up his book as a sort of sensory travelogue through the worlds of various animals, an “attempt to step inside their Umwelten.”

Journeying down this path, Yong makes it clear that many of our nonhuman neighbors, even the humblest bugs, experience the world in moments of richness that are lost to us. To many insects and birds, block-yellow flowers like daffodils are flamed and streaked with painterly brushstrokes of ultraviolet, while those of silverweed are bull’s-eyed with colors we cannot imagine. Plants are not just seen and smelled but felt from a distance: bumblebees sense such plants’ “invisible electric halos”—an electromagnetic force field that every green shoot gives off—with the tiny hairs that make up their fuzz.

Such miniature worlds quake with life. Imperceptibly to us, plants’ springy stems thrum with “haunting, mesmerizing” songs, tapped out by ants, caterpillars, grasshoppers, and other invertebrates that climb on them. In airborne audio, size dictates sound, so large bodies bellow and

small animals sound weedy; freed from these constraints, cicadas moo like cows and crickets conjure the sound of revving chainsaws.

While we learn of pit vipers’ heat vision and the sensory electrical field emitted by the black ghost knifefish, it is often the most familiar creatures that reveal the most startling sensory talents. The pet Labrador walking down the street is guided by nostrils that whirl particles into a continuous vortex, creating a stream of nonstop odor. These smells construct a ghostly *Umwelt*, where objects from the past still reside: a hot dog cart that departed hours before, the skin cells a passerby shed yesterday, or in one test, “a single fingerprint that had been dabbed onto a microscope slide, then left on a rooftop and exposed to the elements for a week.” At night, wandering mice navigate our homes by “whisking,” sweeping the sensitive specialized hairs on their faces to and fro, several times a second, in a manner Yong compares to the action of human eyes, which dart around to construct a scene. Lie awake and you could also, if you had the right ears, hear their ultrasonic, canary-like songs.

Reflecting back on our own sensory abilities, Yong writes of humans: “Our *Umwelt* is still limited; it just doesn’t feel that way. To us, it feels all-encompassing. It is all that we know, and so we easily mistake it for all there is to know.” Yong, like von Uexküll, sees *Umwelt* as a great leveler: all animals, put simply, have a partial and concocted picture of reality provided to us by our evolutionary history—developed, like the jewel wasp’s

stinger, through generations of predation and mating. *An Immense World* sets out to elevate other animals' experience to parity with humans' and dispense with the idea once and for all that humanity's experience is uniquely valuable.

In stating his ambition to "explore their senses to better understand *their* lives," Yong is true to his word. A longtime staffer for the Atlantic, he has an Attenborough-like talent for excavating simple stories from the unbounded mess of the natural world. A look at the eyes of scallops, for example, becomes a window through which to marvel at the dozens or even hundreds of bobbling eyeballs attached to this seafood staple. Yong describes visiting bay scallops with eyes like "neon blueberries." When threatened, the creatures flap furiously toward freedom, "opening and closing their shells like panicked castanets."

An Immense World's most revelatory anecdotes are those that reverse our

worldview and help us to understand how evolutionary pressures have structured physical reality. He tells us that bees, like us, have trichromatic eyes—they perceive three primary colors. In their case, however, the light-sensitive cells are tuned to green, blue, and ultraviolet. "You might think that these pollinators evolved eyes that see flowers well, but that's not what happened," he writes. "Their style of trichromacy evolved hundreds of millions of years before the first flowers appeared, so the latter must have evolved to suit the former. Flowers evolved colors that ideally tickle insect eyes."

Unlike Yong, Jackie Higgins sees animal talents as a lens onto our own faculties. Higgins, who was a science filmmaker for the BBC before becoming an author, centers each chapter of *Sentient* on one animal's remarkable sensory adaptation but takes anecdotal diversions, à la Oliver Sacks, to explore cases on the edges of

human capability. Taking cues from *The Naked Ape*—Desmond Morris's hippie-era melding of zoology and ethnography that interpreted human behavior as the upshot of grand, speculative evolutionary narrative—she prizes the study of animals as "a mirror we can hold up to satisfy self-obsession," adding that "it offers another perspective on why we humans look, act and feel as we do."

There's the peacock mantis shrimp, which has the most complex eyes so far discovered (with 12 types of photoreceptors to our three), and the star-nosed mole, which packs six times more touch sensors into its centimeter-wide splayed snout than you have in an entire hand. Each chapter spotlights one sense, so that in considering color vision, she pairs the example of the shrimp with those of humans grappling with their own equivalent sense: residents of the Pingelap Atoll, for example, the "island of the color-blind," and an

ADVERTISEMENT





anonymous Englishwoman, code-named cDa29, who has a fourth type of photoreceptor that allows her to see millions of colors invisible to the rest of us.

In reading Higgins, we get to spend more time with an organ that appears deliberately unprobed by Yong: the brain. To her, the brain is everywhere, necessarily so as “our body’s all-important sense organ.” Paraphrasing the American neuroscientist Paul Bach-y-Rita, Higgins writes, “We do not see with our eyes, but with our brains. Similarly, we do not solely hear with our ears, smell with our noses, taste with our tongues, or feel with the sensors in our fingers.” In *Sentient*, we learn that spread across the human brain we can find a “sensory homunculus,” a touch map of the body with supersize areas corresponding to our hands and lips, reflecting the density of touch sensors in these zones. There are animal equivalents—“mouseunculus,” “raccoonunculus,” “platypunculus,” and star-nosed “moleunculus”—that likewise represent the primacy of those species’ sensitive whiskers and noses. Indeed, the most affecting sections of the book get closest to the mind, like the chapter on the skin’s “slow lane,” the touch system that responds to caress. The system is found in social mammals, including ourselves but also vampire bats, which have been observed gifting each other blood after caring licks. It is a rare sense that communicates not so much information as mood: “By tuning us to tenderness,” Higgins writes, “it transforms touch into interpersonal glue and the skin into a social organ.”

Through this, we learn that most of what makes up the perceptual world is constructed in the darkness of our head rather than in the sense organs themselves, whose role is limited to translating stimuli into electrical signals. Yet as Higgins and Yong conclude that we really can understand a lot about what it’s like to be another creature, we’re left wondering about this central organ, not having built a clear picture of any other species’ brain—its structure and functioning—nor elucidated much of what’s going on inside it: its cognition or thought. Enter Philip Ball’s *The*

“We do not see with our eyes, but with our brains. Similarly, we do not solely hear with our ears, smell with our noses, taste with our tongues, or feel with the sensors in our fingers.”

Book of Minds. For Ball, senses are just one way into a wide-horized exploration that begins with animal minds and cuts a course through consciousness, artificial intelligence, extraterrestrials, and free will. His book asks: What kinds of minds exist, or could exist, beyond our own? Ball, a prolific science writer and former editor at the journal *Nature*, also sets out with a story from Sacks, who recalled pressing his large, bearded face against the window to the enclosure of an orangutan mother at the Toronto Zoo. As each placed a hand against opposite sides of the pane, Sacks wrote, the two hairy primates shared an “instant, mutual recognition and sense of kinship.”

Even if it is unclear whether we can know what it is like to be a bat, it seemed obvious to Sacks that what it is like to be an orangutan is not just knowable, but something we can intuit readily. Ball’s exploration of the minds of others negotiates this path between solipsism—the skeptical philosophical position that none of us can know anything beyond our own mind—and anthropomorphism, which naively projects our own qualities onto nonhumans. According to him, humans, bats, and orangutans are just three instances within a “Space of Possible Minds” that could also include AI, aliens, and angels.

Rather than posing binary questions—“Is this animal sentient? Has a chatbot become conscious?”—Ball proposes to map out potential minds by their capabilities. This takes a quite literal form in graphs that plot thinking or processing things, from octopuses to

Roomba robot vacuums, along a pair of axes. Neuroscientist Christof Koch has plotted one such graph of “intelligence” against “consciousness,” and computer scientist Murray Shanahan has done so with “human-likeness” set against “capacity for consciousness.” Ball aims to chart a map to the other minds out there that not only display qualities different from ours but in some cases rival and exceed them—as chess-playing AIs now routinely do.

Ball’s own sprawling narrative shows why Yong might have felt it wise not to spend too long with the brain. The eye was Darwin’s go-to example of the dazzling complexity the theory of evolution had to explain. Yet the eye is a device made of clearly comprehensible parts, “including lenses for focusing light, a moveable aperture, photosensitive tissues to record images, delicate colour discrimination, and more.” You could say the same about the ear or other sense organs. “But the brain?” writes Ball. “It makes no sense at all. To the eye it is a barely differentiated mass of cauliflower tissue with no moving parts and the consistency of blancmange, and yet out of it has come *Don Quixote* and *Parsifal*, the theory of general relativity and *The X Factor*, tax returns and genocide.” Discovering one’s own brain is like discovering some alien technology: “With its 86 billion neurons and 1,000 trillion connections, [it] is the most complex object we know of, yet its logic is not one for which other phenomena prepare us.” It is not for nothing that the question of how conscious experience arises out of all this mushy matter is known as the “hard problem of consciousness.”

It would be harsh to criticize Ball for not coming to many clear answers. He is at his best when reformulating the question and problematizing the hasty inferences that plague both AI and animal research. In one section, Ball looks directly at trans-species translation. He tells the story of the marine biologist Denise Herzing, who trained a dolphin pod to associate a set of whistles with sargassum seaweed, one of their favorite playthings. The dolphins assimilated this “word” and later, in the wild, Herzing claimed, used it to convey the same meaning.

This attempt to speak “Dolphinese” raises a head-spinning set of questions. Is this really a language, like the ones we humans make? Are multiple senses involved in creating the meaning for dolphins—as humans combine words and body language? At least since the 1960s, scientists have believed that dolphins as well as some apes have language capabilities—Koko, a gorilla who learned and

Discovering one’s own brain is like discovering some alien technology: “With its 86 billion neurons and 1,000 trillion connections, [it] is the most complex object we know of, yet its logic is not one for which other phenomena prepare us.”

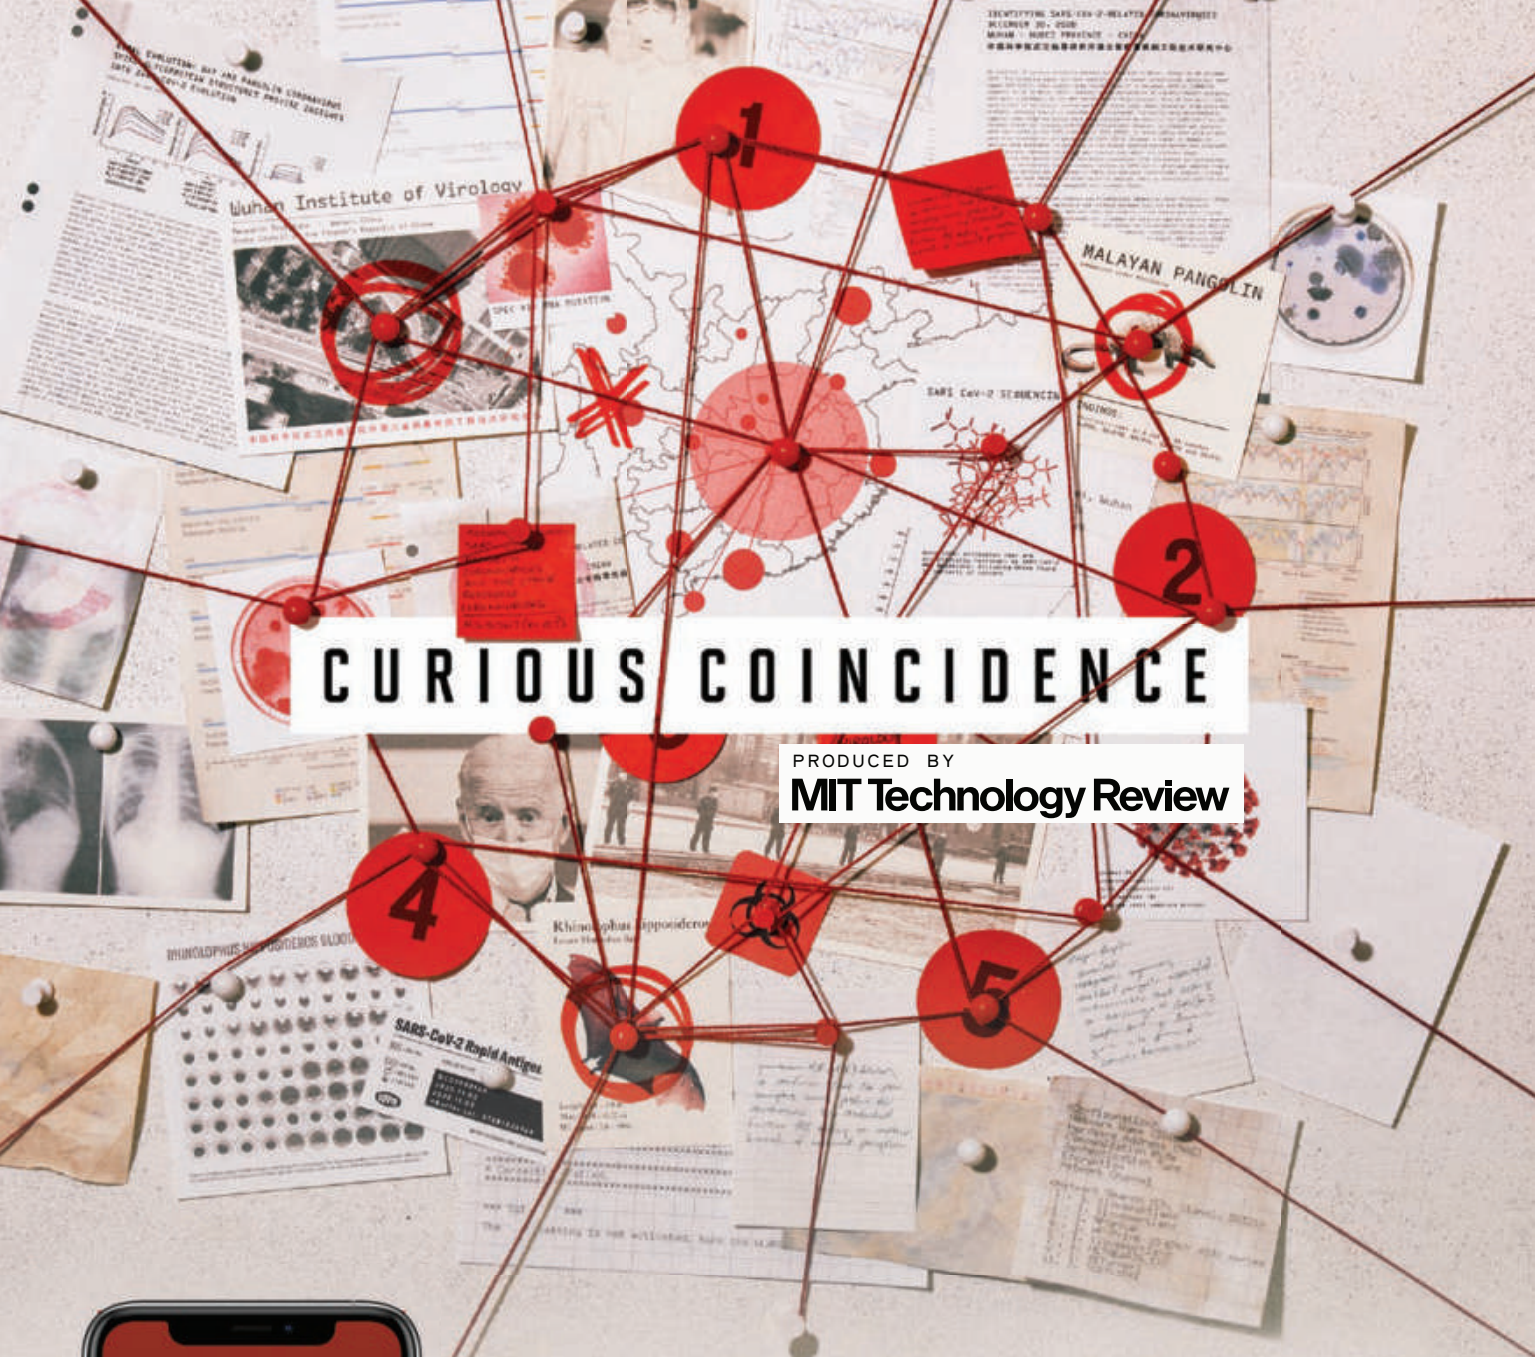
communicated with some hand gestures, being the most famous. But today we are more hesitant than ever, fearful of anthropomorphism, to ascribe too much importance to acquisition of human language.

Throughout, Ball argues that we must be as skeptical of the “philosophical dead end” of solipsism as we are of those who rush to project humanlike experiences onto pets, chimps, or—like one recently dismissed

Google engineer—fairly rudimentary chatbots. Ball’s deep dive into the problems of ascribing conscious minds to others sits neatly next to Yong’s and Higgins’s celebrations, which instead find purpose in imagining just how wide the world might look if other beings really could tell us what they see. ■

Matthew Ponsford is a freelance reporter based in London.

Statement of Ownership, Management and Circulation. (Required by U.S.C. 3685.) (1). Publication Title: MIT Technology Review, (2). Publication No. 535-940, (3). Filing Date: 9/21/22, (4). Issue Frequency: Bi-monthly. (5). No. of Issues Published Annually: 6. (6) Annual Subscription Price: \$120.00. (7). Complete Mailing Address of Known Office of Publication: 196 Broadway, 3rd Floor, Cambridge, MA 02139 (8). Complete Mailing Address of Headquarters or General Business Office of Publisher: 196 Broadway, 3rd Floor, Cambridge, MA 02139 (9). Full Names and Complete Mailing Addresses of Publisher, Editor, and Managing Editor. Publisher: Elizabeth Bramson-Boudreau, 196 Broadway, 3rd Floor, Cambridge, MA 02139. Editor: Mat Honan, 196 Broadway, 3rd Floor, Cambridge, MA 02139. Managing Editor: Timothy Maher, 196 Broadway, 3rd Floor, Cambridge, MA 02139. (10). Owner: Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02138. (11). Known Bondholders, Mortgagees, and Other Security Holders Owning or Holding 1 Percent or More of Total Amount of Bonds, Mortgages, or other Securities: None. (12). Tax Status: The purpose, function, and nonprofit status of this organization and the exempt status for federal income tax purposes: Has Not Changed During Preceding 12 months. (13). Publication Title: MIT Technology Review. (14) Issue Date for Circulation Data Below: September/October 2022. (15). Extent and Nature of Circulation: (a) Total No. of Copies (Net Press Run): Average No. Copies Each Issue During Preceding 12 Months: 140,984; No. Copies of Single Issue Published Nearest to Filing Date: 136,172. (b) Legitimate Paid and/or Requested Distribution (1) Mailed Outside-County Paid Subscriptions Stated on PS Form 3541: Average No. Copies Each Issue During Preceding 12 Months: 111,416; No. Copies of Single Issue Published Nearest to Filing Date: 107,817. (2) In-County Paid/Requested Mail Subscriptions Stated on PS Form 3541: Average No. Copies Each Issue During Preceding 12 Months: 0; No. Copies of Single Issue Published Nearest to Filing Date: 0. (3) Sales Through Dealers and Carriers, Street Vendors, Counter Sales, and Other Paid Distribution Outside the USPS: Average No. Copies Each Issue During Preceding 12 Months: 25,797; No. Copies of Single Issue Published Nearest to Filing Date: 24,788. (4) Paid Distribution by Other Mail Classes Through the USPS: Average No. Copies Each Issue During Preceding 12 Months: 0; No. Copies of Single Issue Published Nearest to Filing Date: 0. (c) Total Paid and/or Requested Distribution: Average No. Copies Each Issue During Preceding 12 Months: 137,213; No. Copies of Single Issue Published Nearest to Filing Date: 132,605. (d) Nonrequested Distribution: (1) Outside-County Nonrequested Copies included on PS Form 3541: Average No. Copies Each Issue During Preceding 12 Months: 0; No. Copies of Single Issue Published Nearest to Filing Date: 0. (2) In-County Nonrequested Copies included on PS Form 3541: Average No. Copies Each Issue During Preceding 12 Months: 0; No. Copies of Single Issue Published Nearest to Filing Date: 0. (3) Nonrequested Copies Distributed Through the USPS by Other Classes of Mail: 0 Average No. Copies Each Issue During Preceding 12 Months: 0; No. Copies of Single Issue Published Nearest to Filing Date: 0 (4) Nonrequested Copies Distributed Outside the Mail: Average No. Copies Each Issue During Preceding 12 Months: 1860; No. Copies of Single Issue Published Nearest to Filing Date: 1979. (e) Total Free or Nominal Rate Distribution: Average No. Copies Each Issue During Preceding 12 Months: 1860; No. Copies of Single Issue Published Nearest to Filing Date: 1979. (f) Total Distribution: Average No. Copies Each Issue During Preceding 12 Months: 140,935; No. Copies of Single Issue Published Nearest to Filing Date: 136,122. (g) Copies not Distributed: Average No. Copies Each Issue During Preceding 12 Months: 0; No. Copies of Single Issue Published Nearest to Filing Date: 0. (h) Total: Average No. Copies Each Issue During Preceding 12 Months: 140,985; No. Copies of Single Issue Published Nearest to Filing Date: 136,172. (i) Percent Paid: Average No. Copies Each Issue During Preceding 12 Months: 98.7%; No. Copies of Single Issue Published Nearest to Filing Date: 98.5%. This Statement of Ownership will be printed in the November/December 2022 issue of this publication. I certify that all the information furnished on this form is true and complete. **Alison Papalia, VP of Marketing and Consumer Revenue.**



CURIOUS COINCIDENCE

PRODUCED BY
MIT Technology Review



A NEW MIT TECHNOLOGY
REVIEW ORIGINAL SERIES
INVESTIGATES THE
MYSTERY OF COVID-19.

FROM THE
CREATORS OF:



FOLLOW NOW
WHEREVER YOU
GET YOUR PODCASTS



Is There A Limit to Human Life?

Intrigued through the Ages, Scientists and Philosophers

Have Advanced Theories to Explain Longevity but

Heredity Still Appears as Dominant Factor

By JAMES A. TOBEY

THE longest human life in history is supposed to have been that of an Englishman named Thomas Parr. This venerable gentleman is reputed to have been born in Shropshire in 1483 and to have died in London in 1635, at the fairly ripe old age of 152. Because of the notoriety connected with this alleged great age, he was buried in the south transept of Westminster Abbey, where the inscription on his tomb states that he lived in "ye reignes" of 10 princes, from Edward IV to Charles I.

According to tradition, Thomas Parr spent the first 80 years of his long life working as a bachelor farmer. Then he married and sired two children, both of whom died in infancy. At the tender age of 105 he was charged with having caused the birth of an illegitimate child and was sentenced to do penance by standing in the church door in a white sheet. His wife died when he was a mere 110, but he remarried at 112, and at 130 Mr. Parr was busily engaged in threshing corn.

In 1635 the fame of this antique specimen of humanity reached Thomas Howard, Earl of Arundel, who whisked him off to London, presented him to King Charles, and otherwise made much of him. The fast life of the metropolis, or as the contemporary accounts put it, "the change of air and diet," proved too much for the old man, and he succumbed in November of 1635. An autopsy was performed on the body by Dr. William Harvey, celebrated as the discoverer of the circulation of the blood, who reported that the physical condition of this famous English centenarian was excellent.

The advanced age of "Olde" Parr would, if true, be of tremendous interest and significance to all concerned with human longevity, but unfortunately it is not true. No one disputed this unusual age, however, until 1873 when a skeptical librarian of the House of Lords, William J. Thoms, decided to investigate this case and numerous other alleged instances of extreme old age. A good basis for his study was a report published in 1800 by James Easton listing no less than 1,712 "true" centenarians who had flourished from 66 A.D. to 1799. Included, of course, was Olde Parr and another man named Henry Jenkins, who was declared to have been even older, a modest 169. Also in the record was a Countess of Desmond at 140.

After an exhaustive examination of 22 cases of such human sempiternity, the doubting Mr. Thoms reached the apt conclusion that the evidence justified none of

them. He demonstrated, for example, that the lives of two different countesses of Desmond had been added together to produce the age of 140, and that neither of them had lived much beyond 70 years. With respect to Thomas Parr it was shown rather conclusively that at least 50 years had been erroneously, or even fraudulently, added to his actual age and that he could not have been over 100 years old at the time of his death.

Despite this masterful job of debunking, one will still see occasional references to Olde Parr as the oldest man in the world. The lie lingers on, but the correction is forgotten. There is even a short article on Thomas Parr in the 14th edition of the *Encyclopaedia Britannica*, but no mention of the work of Mr. Thoms of the House of Lords.

A few years after this nice bit of exposure, another author who was familiar with it, John B. Bailey, issued a book entitled *Modern Methuselahs** in which he described numerous real and near centenarians, but he was able to find only about a dozen individuals who were undeniably in this hoary category. In another competent survey of centenarians, made in 1899 by a British actuary, T. E. Young, only 22 authentic instances of persons aged 100 years or more were cited, although a second edition of this work, issued in 1905, added another eight to the original list. The longest verifiable age mentioned by Mr. Young was one of 111 years, and in the 50 years since this report, no one has been able to offer acceptable proof of any human life longer than this. There have been plenty of claims, but no proof. Thus, about the time Young's second edition came out, a Russian newspaper blandly announced the death of a woman named Therese Abalva at the somewhat advanced age of 180, but offered no evidence to substantiate this extravagant and incredible statement.

The only instance of unique senescence which offers even a scintilla of reliability is that of a Dane named Christen Jacobsen Dragenberg, who was reported as living to the age of 146. A reputable statistician who wrote up this case in an actuarial journal was convinced that it was true, but most biologists are skeptical. This Dragenberg was said to have been born in 1626, to have gone to sea at the age of 13, to have served three kings in wars against Sweden, to have been captured by Algerian pirates at the age of 70, and to have been held in captivity for 19 years. Like

* London: Chapman and Hall, Ltd., 1888.

People have always been fascinated with the question of human longevity. In this 1954 piece for *Technology Review*, James A. Tobey, author of more than a dozen books on public health, including *Your Diet for Longer Life* (1948), noted that despite a few frauds claiming to be older than 150, "the consensus of scientific opinion is that there is a definite limit to human life, a limit now and perhaps forever in the vicinity of 100 years."

In 1954, the average life expectancy of an American at birth had risen to 68 years from 47 in 1900. But most of the advances came not from old people living longer but from infants avoiding death before their first birthday. The average person's chances of living to 100 in mid-20th-century America, Tobey observed, were "no better than they were in the days of the Roman Empire."

We have done better since then: average life expectancy reached nearly 79 years in the US before declining in recent years, largely as a result of the covid-19 pandemic. But as this issue of *TR* reveals, the quest to keep extending the upper limit on our years lives on.

MIT
Technology
Review

Read smarter, not harder

A subscription to MIT Technology Review includes

- Unlimited web access
- 10% discount for all our events
- Digital version of each issue
- Subscriber-only app
- Access to 120+ years of publication archives

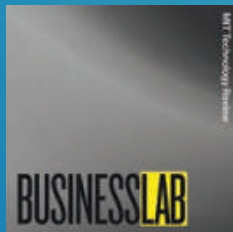
→ **Learn more at technologyreview.com/subscribe**

Listen to us



In Machines We Trust

Technologies with the ability to learn have unlocked a world of possibilities that we're only beginning to understand.



Business Lab

Unpacking the new technologies coming out of the lab and moving into the marketplace.



Curious Coincidence

Separating fact from fiction while uncovering the origins of covid-19.



The Extortion Economy

An exploration of common types of cyber-crime, who's been affected, and who has secretly profited.

→ Download wherever you listen.

