



Center for Mind, Ethics, and Policy

2024 ANNUAL REPORT

Table of Contents

LETTER FROM THE DIRECTOR	1
ABOUT THE CENTER	3
ROADMAP TO 2024	4
2024 RESEARCH	5
2024 EVENTS	8
OTHER HIGHLIGHTS	11
LOOKING AHEAD	14
OUR TEAM	18

Letter from the Director

Dear Friends and Colleagues,

It is my enormous honor to introduce the first annual report for the New York University Center for Mind, Ethics, and Policy (CMEP). 2024 has been a pivotal year for our work, as we transitioned from a renewable program to a permanent center and continued to address some of the most pressing and profound scientific and ethical questions of our time.

CMEP is an endowed research and outreach center that examines the nature and intrinsic value of nonhuman minds, with special focus on invertebrates and AI systems. We produce books, papers, and reports; host talks, workshops, and conferences; and support researchers at all career stages. We take an integrative approach, linking research on biological and digital minds; research in the humanities, social sciences, and natural sciences; and research in academia, government, and industry. We also take a collaborative approach, examining how we can make progress on these issues together despite ongoing disagreement and uncertainty about them.

Academic leadership on the nature and intrinsic value of nonhuman minds has never been more urgent. Humans share the world with countless

nonhumans, and we increasingly interact with them. We have a responsibility to exercise this power responsibly, but doing so requires engaging with some of the most difficult issues in philosophy and science, such as the nature of morality, the nature of consciousness, and the future of technology. As companies and governments consider whether and how to scale up industries ranging from insect farming to AI development, guidance from experts in these fields will be essential.



Jeff Sebo, Director

In 2024, CMEP advanced our work on multiple fronts, with three clear highlights. First, in Spring 2024 we released the New York Declaration on Animal Consciousness (co-sponsored with the NYU Wild Animal Welfare Program and co-led with Kristin Andrews and Jonathan Birch). The Declaration holds that all vertebrates and many invertebrates have a realistic chance of being

conscious, and that we have a responsibility to consider welfare risks for them. Launched at a conference at NYU with 40 initial signatures, the Declaration has now collected more than 500 signatures and received coverage in more than 100 media outlets.

The second highlight came in summer 2024, when we secured a \$6 million endowment from The Navigation Fund (\$5 million) and Polaris Ventures (\$1 million). As noted above, this endowment allowed us to transition from a renewable program to a permanent center, housed alongside the Center for Environmental and Animal Protection (CEAP) at the Department of Environmental Studies. Following the endowment, CMEP and CEAP jointly hired a new administrator, Audrey Becker. Audrey will now support both centers alongside coordinator Sofia Fogel and researcher Toni Sims, increasing capacity and efficiency across the board.

The third highlight came in Fall 2024, when we released “Taking AI Welfare Seriously” (co-sponsored with Eleos AI Research and co-led with Robert Long). This report argues that AI companies have a responsibility to take AI welfare seriously, and it provides recommendations about how they can start. When we released the report, Anthropic—which funded initial research that led to the report—confirmed to Transformer News that they hired report author Kyle Fish as their first full-time AI welfare researcher. Both the report and the hire have since received extensive coverage in outlets such as *Nature*, *The Hill*, and *The Times UK*.

Other highlights from 2024: We released a paper on the moral status of individuals and collectives, a paper on the moral status of connected minds (co-authored with Luke Roelofs), and an open access edition of my book *Saving Animals, Saving Ourselves*. We also hosted a live taping of the Clearer Thinking podcast, co-hosted a

conference on animals and equality, and co-hosted the U.S. launch of Jonathan Birch’s book *The Edge of Sentience*. Faculty affiliates Ned Block, Sam Bowman, David Chalmers, and Claudia Passos Ferreira produced relevant work this year as well, and the NYU Animal Studies and Environmental Studies community that houses CMEP continued to expand.

In 2025, we will continue our work on a number of fronts. My book *The Moral Circle* will launch in January, and we plan to release new papers throughout the year, including on interspecies and intersubstrate welfare comparisons, the legal status of insects and AI systems, the science of animal consciousness, and the relationship between AI safety and AI welfare. Our team is always eager to find new projects as well, so if you have ideas for us, please let us know. You can visit our website for information on research, events, and opportunities, and you can contact us at nonhumanminds@nyu.edu with questions or suggestions.

Of course, none of this work would be possible without the dedication of our team, partners, and supporters. I extend my heartfelt gratitude to everyone who shares our vision for a more just and inclusive world. Thank you for being part of this journey; I look forward to continuing to work with you to build on this foundation in 2025 and beyond.

With gratitude,



Jeff Sebo
Director

About the Center

At present, the world contains quintillions of nonhuman animals. Human activity is increasingly shaping the lives of these animals, by determining whether they can exist and what kinds of lives they can have if they do. And in the future, nonhuman populations might be much larger, and might include advanced AI systems as well.

These trends raise important questions at the intersection of mind, ethics, and policy. Which nonhumans are conscious, sentient, and agentic? What kind of moral, legal, and political status should they have? How can humans build a positive future for the vast multiplicity of potentially morally significant beings who might one day exist?

These questions, in turn, require us to confront some of the hardest problems in science, philosophy, and policy. What is the nature of consciousness? Can we have knowledge about other minds? Can we make welfare comparisons across species? What do we owe members of other nations, generations, species, and substrates?

The NYU Center for Mind, Ethics, and Policy (CMEP) is an endowed research center that conducts and supports foundational research on the nature and intrinsic value of nonhuman minds, with special focus on invertebrates and AI systems. We also engage in outreach and field-building activities, hosting events and supporting early-career researchers.

Our current research agenda focuses on the following general themes, all of which are important, difficult, and contested—calling for considerable caution and humility:

- Status: Which nonhumans matter for their own sakes?
- Weight: How much do particular nonhumans matter for their own sakes?
- Welfare: What counts as good or bad for particular nonhumans?
- Ethics: What do we owe particular nonhumans?
- Practice: What follows for our lives and societies?

CMEP launched as a renewable program in Fall 2022 and relaunched as a permanent center in Fall 2024. The Center is located in the Department of Environmental Studies alongside the Center for Environmental and Animal Protection, with which it shares personnel and collaborates on projects of shared interest.

MISSION

Our aim is to advance understanding of the consciousness, sentience, agency, moral status, legal status, and political status of nonhumans—biological as well as digital—in a rigorous, systematic, and integrative manner. We pursue this goal via research, outreach, and field building in science, philosophy, and policy.

Roadmap to 2024

CMEP is the successor to the NYU Mind, Ethics, and Policy Program (MEPP), which launched in Fall 2022. Since there are no annual reports from our time as a program, we can open this report with a brief list of highlights from our first year. From Fall 2022 to Fall 2023:

- We hosted a public launch event with a talk by David Chalmers on AI consciousness
- We hosted a public talk by Peter Godfrey-Smith on AI consciousness
- We hosted a public talk by Bob Fischer on interspecies welfare comparisons
- We hosted a public talk with Jonathan Birch and Susan Schneider on indicators for consciousness in insects and AI systems
- We hosted a public event with Yoshua Bengio, Patrick Butlin, Grace Lindsay, and Robert Long about their 2023 AI consciousness report
- We hosted an award and workshop on animal and AI consciousness, with talks by five award-winning early-career scholars and a keynote address by L. Syd M. Johnson
- We hosted a symposium on animal and AI consciousness at the 2023 Pacific American Philosophical Association Conference, with talks from Kristin Andrews, Amanda Askell, Robert Long, David Peña-Guzmán, Claudia Passos Ferreira, and Jonathan Simon
- We hosted a symposium on mind, ethics, and policy at the 12th Oxford Workshop in Global Priorities Research at the Global Priorities Institute, with talks from Bob Fischer, Becca Franks, Robert Long, and MEP Director Jeff Sebo (JS)
- JS published a paper on [population ethics](#) for insects and AI systems
- JS published an essay on [principles for AI welfare research](#)
- JS published a paper with Robert Long on [moral consideration for AI systems by 2030](#)
- JS published a paper with Bob Fischer on [intersubstrate welfare comparisons](#)
- JS published several op-eds and essays on animal and AI ethics, including op-eds on [AI sentience](#) and [nonhuman rights](#) in the Los Angeles Times
- JS gave dozens of talks and several podcast interviews about animal and AI ethics, including an interview about digital minds at the [80,000 Hours Podcast](#)



2024 Research

CMEP supports research on the nature and value of nonhuman minds by contributing funding, authorship, or both. What follows is a list of outputs to which our team contributed in 2024.

Taking AI Welfare Seriously

Independent report (co-sponsored with Eleos AI)

Robert Long¹
Eleos AI

Jeff Sebo²
New York University

Patrick Butlin³
University of Oxford

Kathleen Finlinson⁴
Eleos AI

Kyle Fish⁵
Eleos AI, Anthropic

Jacqueline Harding⁶
Stanford University

Jacob Pfau⁷
New York University

Toni Sims⁸
New York University

Jonathan Birch⁹
London School of Economics

David Chalmers¹⁰
New York University

Abstract

In this report, we argue that there is a realistic possibility that some AI systems will be conscious and/or robustly agentic in the near future. That means that the prospect of AI welfare and moral patienthood—of AI systems with their own interests and moral significance—is no longer an issue only for sci-fi or the distant future. It is an issue for the near future, and AI companies and other actors have a responsibility to start taking it seriously. We also recommend three clear steps that AI companies and other actors can take: They can (1) acknowledge that AI welfare is an important and difficult issue (and ensure that language model outputs do the same), (2) start assessing AI systems for evidence of consciousness and robust agency, and (3) prepare policies and procedures for treating AI systems with an appropriate level of moral concern. We also discuss frameworks that AI companies and other actors can consider when doing this work.

¹Lead and corresponding authors. ²Robert@eleosai.org, jeffsebo@nyu.edu
³More authors.
⁴Contributing author.
⁵Work performed while at Eleos AI, prior to joining Anthropic in Fall 2024.

Taking AI Welfare Seriously

Independent report (co-sponsored with Eleos AI)

Robert Long (Eleos AI)

Jeff Sebo (New York University)

Patrick Butlin (University of Oxford)

Kathleen Finlinson (Eleos AI)

Kyle Fish (Eleos AI, Anthropic)

Jacqueline Harding (Stanford University)

Jacob Pfau (New York University)

Toni Sims (New York University)

Jonathan Birch (London School of Economics)

David Chalmers (New York University)

In Fall 2024, CMEP and Eleos released a major report on the possibility of near-future AI welfare. We argue that there is a realistic possibility that some AI systems will be conscious and/or robustly agentic in the near future. That means that the prospect of AI welfare and moral patienthood—of AI systems with their own interests and moral significance—is no longer an issue only for sci-fi or the distant future. It is an issue for the near future, and AI companies and other actors have a responsibility to start taking it seriously.

We also recommend three minimum necessary first steps that AI companies and other actors can take to address this issue: They can (1) acknowledge that AI welfare is an important and difficult issue (and ensure that language model outputs do the same), (2) start assessing AI systems for evidence of consciousness and robust agency, and (3) prepare policies and procedures for treating AI systems with an appropriate level of moral concern. We also discuss frameworks that AI companies and other actors can consider when doing this work.

Our argument in this report is not that AI systems definitely are—or will be—conscious, robustly agentic, or otherwise morally significant. Instead, our argument is that there is substantial uncertainty about these possibilities, and so we need to improve our understanding of AI welfare and moral patienthood. Otherwise there is a significant risk that we will mishandle decisions about AI welfare, mistakenly harming AI systems that matter morally and/or mistakenly caring for AI systems that do not.



This paper is open access; you can read it [here](#).

SELECT PRESS COVERAGE

When we released the report, Anthropic—which funded initial research that led to the report—confirmed to [Transformer News](#) that they hired report author Kyle Fish as their first full-time AI welfare researcher. Both the report and the hire have since received extensive coverage.



Is “AI welfare” the new frontier in ethics?

By Benj Edwards, November 11, 2024

“The paper warns that AI models could soon develop consciousness or agency—traits that some might consider requirements for moral consideration.”



AI could cause ‘social ruptures’ between people who disagree on its sentience

By Robert Booth, November 17, 2024

“Last week, a transatlantic group of academics predicted that the dawn of consciousness in AI systems is likely by 2035.”



Plans must be made for the welfare of sentient AI, animal consciousness researchers argue

By Saul Elbein, October 31, 2024

“Computer scientists need to grapple with the possibility they will accidentally create sentient artificial intelligence (AI) ... a new study argues.”



What should we do if AI becomes conscious? These scientists say it's time for a plan

By Mariana Lenharo, December 10, 2024

“A group of philosophers and computer scientists are arguing that AI welfare should be taken seriously.”



Should chatbots have rights—and should we care?

By Alex Wilkins, December 23, 2024

“In a recent academic paper, “Taking AI Welfare Seriously”, one group of researchers argue for a precautionary approach to how we treat AIs.”



If robots have feelings, do they need rights?*

By Tom Whipple, November 10, 2024

“What if the large language model that politely answers all your questions also has its own wants and desires?”

*Front page story

oxford public philosophy

Jeff Sebo

Are individuals or ecological wholes what matter? Yes.

There tends to be strong *disagreements* in animal and environmental ethics between individualists (1) who hold that individuals are the primary units of moral analysis, and ecocentrists (2) who hold that ecological wholes are the primary units of moral analysis. These disagreements are partly theoretical, concerning which entities are morally important and why. They are also partly practical, concerning which beings merit priority status, for example, the needs of individual animals conflict with the needs of the species or ecosystems that contain them.

Of course, even when ethicists disagree about whether animals or ecological wholes are the primary units of moral analysis, they can still agree that animals and ecological wholes have a lot of value. For example, even if individual animals are what matter for their own sake, ecological wholes can still matter a lot for the sake of the animals they contain. Similarly, even if ecological wholes are what matter for their own sake, individual animals can still matter a lot for the sake of the ecological wholes they constitute.

Are individuals or ecological wholes what matter? Yes.

Oxford Public Philosophy

Jeff Sebo (New York University)

There tends to be strong disagreement in animal and environmental ethics between individualists, who hold that individuals are the primary units of moral analysis, and ecocentrists, who hold that ecological wholes are the primary units of moral analysis. In this essay I suggest that the concept 'primary unit of moral analysis' is ambiguous, and that when we disambiguate it, we can identify a plausible view according to which individualists are correct in one sense and ecocentrists are correct in another sense. Specifically, in both science and ethics, we can make a distinction between the most basic units of analysis and the most helpful units of analysis, and we can say that smaller beings like individuals tend to be more basic but that larger beings like ecological wholes tend to be more helpful in many contexts.



This essay is open access; you can read it [here](#).

Overlapping minds and the hedonic calculus

Philosophical Studies

Luke Roelofs (University of Texas at Arlington)
Jeff Sebo (New York University)

How should we update our moral thinking if it turns out to be possible for a single token mental state—a feeling of pleasure, displeasure, desire, frustration, or something else relevant to welfare—to belong to two or more subjects at once? Some philosophers think that such sharing of mental states might already occur, whereas others foresee it as a potential consequence of advances in neurotechnology and AI. Yet different types of cases generate opposite intuitions: If two people share a few mental states, it seems we should count the value of those states twice, but if two beings share their whole mental lives, it seems we should count the value of that life once. This paper suggests that these conflicting intuitions can be reconciled if the mental states that matter for welfare have a holistic character.



This article is open access; you can read it [here](#).

Philosophical Studies (2024) 181:1487–1506
https://doi.org/10.1007/s11080-024-02157-z



Overlapping minds and the hedonic calculus

Luke Roelofs  · Jeff Sebo 

Accepted: 5 May 2024 / Published online: 29 May 2024
© The Author(s) 2024

Abstract
It may soon be possible for neurotechnology to connect two subjects' brains such that they share a single token mental state, as a feeling of pleasure or displeasure. How should we update our moral thinking if it turns out that this can happen? And if this sort of mental-state-sharing might already obtain in some cases, how should this possibility impact our moral thinking? This question turns out to be extremely challenging. If two people share a few mental states, then it seems that we should count the value of those states twice, but if they share very many mental states, then it seems that we should count the value of those states only once. This paper suggests that these intuitions can be reconciled if the mental states that matter for welfare have a holistic character, in a way that is morally plausible. We close by drawing alternative conclusions about how we ought to think about the moral importance of distributed mental states.

Keywords Neuroethics · Welfare · Overlapping minds · Metaphysics · Ethics · Hedonism · collective consciousness · individualization · value aggregation

Luke Roelofs, Jeff Sebo have contributed equally to this work.

✉ Luke Roelofs
luke.roelofs@ut Arlington.edu
Jeff Sebo
jeffsebo@nyu.edu

¹ Department of Philosophy and Humanities, University of Texas at Arlington, 305 Carlile Hall, Box 19027, Arlington, TX 76093, USA
² Environmental Studies, New York University, 285 Mercer Street 10th Floor, New York, NY 10003, USA



2024 Events

This year, CMEP sponsored several major events, including the launch event for the New York Declaration on Animal Consciousness. We also shared our work at many events, but this section focuses on events that we sponsored.

The Emerging Science of Animal Consciousness

April 19, 2024

SPEAKERS

Kristin Andrews (York University)
Jonathan Birch (London School of Economics)
David Chalmers (New York University)
Lars Chittka (Queen Mary University of London)
Robert Elwood (Queen's University Belfast)
Becca Franks (New York University)

Noam Miller (Wilfrid Laurier University)
Alexandra Schnell (University of Cambridge)
Jeff Sebo (New York University)
Cleo Verkuijl (Stockholm Environment Institute)
Katrina Wyman (New York University)
Anna Wilkinson (University of Lincoln)

DESCRIPTION

Which animals have the capacity for conscious experience? In this public event, experts met to discuss the progress that has been made, the key points of agreement and disagreement, the most promising directions for the future, and what recognizing other animals as conscious beings means in practice for ethics and policy.

This conference served as the launch event for the New York Declaration on Animal Consciousness. The Declaration holds that there is a realistic possibility of consciousness in all vertebrates and many invertebrates, including insects, and that we have a responsibility to consider welfare risks for these animals when making decisions that affect them.

The Declaration was organized by Kristin Andrews, Jonathan Birch, and Jeff Sebo, and the launch event was hosted by CMEP and the NYU Wild Animal Welfare Program, with co-sponsorship from the NYU Center for Bioethics, the NYU Center for Mind, Brain, and Consciousness, and NYU Animal Studies.



You can read the Declaration and view the talks from the conference [here](#).

SELECT PRESS COVERAGE

The Declaration has now collected more than 500 expert signatures and received coverage in more than 100 media outlets all around the world.



Are animals conscious? How new research is changing minds

By Pallab Ghosh, June 15, 2024

“The short four paragraph declaration states that it is “irresponsible” to ignore the possibility of animal consciousness.”



Consciousness a ‘realistic possibility’ in birds, fish, squid and bees, scholars say

By Sarah Smellie, April 22, 2024

“Scientists and philosophers across the globe agree it is reasonable to assume the vast majority of creatures on Earth are sentient in some way.”



Do insects have an inner life? Animal consciousness needs a rethink

By Mariana Lenharo, April 19, 2024

“A declaration signed by dozens of scientists says there is ‘a realistic possibility’ for elements of consciousness in reptiles, insects and molluscs.”



Are animals conscious? We’re finally realising that many species are

By Penny Sarchet, July 18, 2024

“Science is at last confirming what many people have long suspected—that mammals, birds and perhaps some invertebrates have elements of consciousness.”



Scientists push new paradigm of animal consciousness, saying even insects may be sentient

By Evan Bush, April 19, 2024

“Far more animals than previously thought likely have consciousness, top scientists say in a new declaration—including fish, lobsters and octopus.”



Insects and Other Animals Have Consciousness, Experts Declare

By Dan Falk, April 19, 2024

“A group of prominent biologists and philosophers announced a new consensus: There’s “a realistic possibility” that insects, octopuses, crustaceans, fish and other overlooked animals experience consciousness.”

The Moral Status of Insects and AI Systems, and Other Thorny Questions

January 30, 2024

CMEP hosted a special live taping of the Clearer Thinking podcast. Podcast host Spencer Greenberg and guest Jeff Sebo discussed the moral status of insects and AI systems, as well as other thorny questions in global priorities research. The conversation covered a variety of topics, including the possibility of widespread consciousness in the world, the ethics and politics of factory farming, and the ethics and politics of AI development and deployment. This event also featured a vegan reception for in-person guests.



You can listen to the podcast episode [here](#).

Animals and Equality Conference

October 18–19, 2024

CMEP co-sponsored this conference with the Duke Center for Law, Economics, and Public Policy. This conference focused on the role of animals in broadly egalitarian views, representing philosophy, welfare economics, and social choice theory. Participants discussed whether animals have full status, partial status, or no status within egalitarian views, and they also discussed a variety of related questions in ethics, law, and policy. Matthew Adler organized and hosted this event at Duke University with support from Jeff Sebo and Heather Browning.



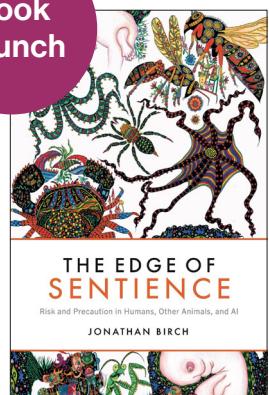
You can view information about the conference [here](#) and listen to a brief interview [here](#).

Book Launch: *The Edge of Sentience* by Jonathan Birch

November 11, 2024

CMEP co-hosted the U.S. launch of Jonathan Birch's *The Edge of Sentience* with the Guarini Center for Environmental, Energy, and Land Use Law. *The Edge of Sentience* presents a precautionary framework for reaching ethically sound, evidence-based decisions despite our uncertainty. At the event, Birch introduced his book, and then panelists L. Syd M. Johnson, John Olusegun Adenitire, and Claudia Passos Ferreira offered comments and discussion. We then had a lively discussion with the audience followed by a vegan reception for in-person guests.

Book
Launch



You can view the video of the event [here](#).

Other Highlights

There were many other highlights this year as well, both for CMEP and for our broader community of partner programs, faculty affiliates, and regular collaborators. We list several highlights here, but this list is not exhaustive:



The main highlight for CMEP in 2024 (aside from our research and events) is that we established a \$6 million endowment from The Navigation Fund (\$5 million) and Polaris Ventures (\$1 million). This endowment enabled us to transition from a renewable program to a permanent center, and starting in summer 2025, the interest generated from the endowment will provide us with a permanent annual baseline of approximately \$270,000, to which we can add further funds as we continue to grow.



As soon as we established our endowment, CMEP partnered with the NYU Center for Environmental and Animal Protection ([CEAP](#)) to hire a full-time administrator, Audrey Becker. Audrey will support both Centers alongside director Jeff Sebo, coordinator Sofia Fogel, and researcher Toni Sims. The NYU Wild Animal Welfare Program ([WAWP](#)) will participate in this arrangement as well, and this core team will significantly increase capacity and efficiency for all three programs.



The [NYU Department of Environmental Studies](#), which houses CEAP, CMEP, and WAWP, continues to grow as well. In addition to launching these research programs, our department has launched a number of research labs and teaching programs over the past six years, including the NYU Animal Studies M.A. Program. In 2025, the department will launch a world-class [Environmental Studies PhD Program](#) and move into a beautiful new building overlooking Washington Square Park.



Our partner programs released important work this year too. For example, CEAP released a [major report](#) on global zoonotic disease risks in animal markets across 15 countries and six continents (co-sponsored with the Harvard Animal Law & Policy Program), and WAWP released a [major report](#) on how cities can include wild animal welfare in infrastructure policies (co-sponsored with CEAP, the Guarini Center on Environmental, Energy, and Land Use Law, and the Yale Law, Environment & Animals Program).



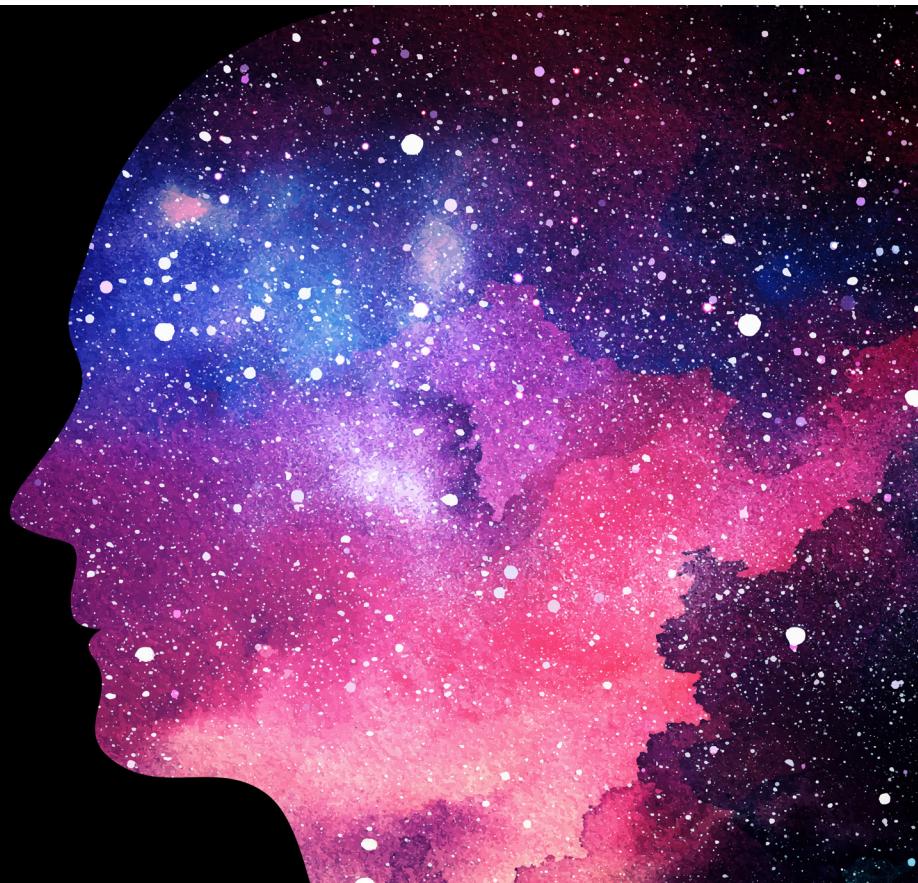
CMEP Director Jeff Sebo released an open access edition of his book *Saving Animals, Saving Ourselves* (Oxford University Press), which examines the relationship between animals, pandemics, and climate change. Endorsed by leaders like Cory Booker and Jane Goodall, this book calls for including animals in global health and environmental policy by reducing our use of animals as part of our pandemic and climate change mitigation efforts and increasing our support for them as part of our adaptation efforts.

CMEP faculty affiliate Ned Block published several articles on consciousness this year, including: “[What does decoding from the PFC reveal about consciousness?](#)” in *Trends in Cognitive Sciences*; “[States of mind](#)” in *Science*; “[Aphantasia as imagery blindsight](#)” (with Matthias Michel, Jorge Morales, and Hakwan Lau) in *Trends in Cognitive Sciences*; and “[Memory representations during slow change blindness](#)” (with Haley G. Frey, Luca Koenig, Biyu J. He, and Jan W. Brascamp) in *Journal of Vision*.

CMEP faculty affiliate Sam Bowman published a widely circulated blog post about [next steps for AI safety](#). The post discusses the importance of laying the groundwork for AI welfare commitments in the near future, describing this issue as “almost uniquely confusing ... drawing on hard unsettled empirical questions as well as deep open questions in ethics and the philosophy of mind.” It also calls for pursuing low-hanging fruit interventions while taking a pluralist approach by “drawing on a number of different worldviews.”



CMEP faculty affiliate David Chalmers gave the Descartes Lectures on “Large Language Models and the Philosophy of Mind” at Tilburg University. He also gave numerous talks on AI consciousness (for instance, at UC Berkeley and the ASSC conference at The University of Tokyo), AI interpretability (for instance, to the British Academy and the NYU Center for Data Science), and related themes. Finally, he was elected as a Fellow of the British Academy in recognition of his many accomplishments in philosophy of mind, cognitive science, and other areas.





CMEP faculty affiliate Claudia Passos Ferreira published two articles on infant consciousness, a topic that raises many of the same questions as nonhuman consciousness: “[Can We Detect Consciousness in Newborn Infants?](#)” in *Neuron*, and “[Through a baby’s ears](#)” in *New Scientist*. She also gave ten talks related to consciousness in Italy, Japan, and the U.S., including at CMEP’s U.S. launch of Jonathan Birch’s book *The Edge of Sentience*. She gave several interviews on the topic as well, including at Consciousness Central and NPR.

CMEP regular collaborators Robert Long and Kathleen Finlinson launched [Eleos AI Research](#), a nonprofit organization that seeks to build a deeper understanding of AI sentience and wellbeing. Eleos will improve research and discourse related to these topics by developing tools and recommendations for AI labs, policymakers, and other stakeholders. As noted above, CMEP and Eleos collaborated on “[Taking AI Welfare Seriously](#)” in 2024, and we have further collaborative projects in the works for 2025 and beyond.

Vox included CMEP Director Jeff Sebo and regular collaborator Jonathan Birch in its [2024 Future Perfect 50](#) list of “innovators, thinkers, and changemakers working to make the world a better place.” They included Sebo and Birch in a new “expanding the mind” category, emphasizing the importance of research on nonhuman minds, ethics, and policy. They included other colleagues as well, including our NYU Environmental Studies colleague Matthew Hayek for his research related to animal farming.



Looking Ahead

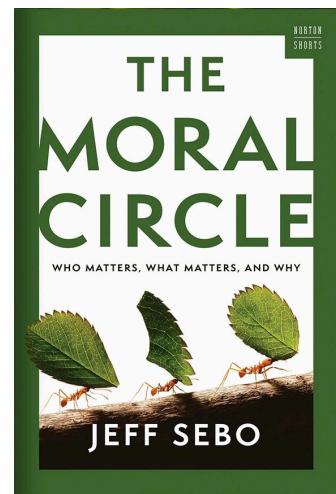
Here are some upcoming projects that our team is either leading or supporting. These projects are subject to change and this list is not exhaustive; some projects are not yet ready to share, and we also leave room in our plans to take advantage of opportunities as they arise.

The Moral Circle

W. W. Norton (2025)

Jeff Sebo (New York University)

As the dominant species, humanity has a responsibility to ask: Which beings matter, how much do they matter, and what do we owe them in a world reshaped by human activity? *The Moral Circle* argues that we should include all potentially significant beings in our moral community, with transformative implications for our lives and societies. This book explores provocative case studies, such as lawsuits over captive elephants and debates over factory-farmed insects. It also explores future quandaries such as whether to send microbes to new planets, and whether to create virtual worlds filled with digital minds. Taking an expansive view of human responsibility, the book argues for shedding human exceptionalism and radically rethinking our place in the world.



Reviews of *The Moral Circle*

“In *The Moral Circle*, Jeff Sebo defends, with engaging examples and plausible arguments, a disturbing conclusion: that the expansion of the moral circle beyond our own species, for which I have argued, does not go nearly far enough. I expect this book to trigger an important debate.”

– Peter Singer, author of *Animal Liberation* and *Animal Liberation Now*

“Does what happens to an insect, a plant, or a robot matter to them? Compelling thought experiments make *The Moral Circle* a fascinating read. Jeff Sebo’s conclusion that we owe a lot to a vast number and wide range of beings will surely expand readers’ empathy and compassion.”

– Barbara J. King, author of *Animals’ Best Friends* and *How Animals Grieve*



You can purchase *The Moral Circle* in text or audio [here](#).

Animal Rights

Palgrave Handbook on the Philosophy of Rights
(forthcoming)

Adam Lerner (Rutgers University)
Jeff Sebo (New York University)

This chapter explores the concept of animal rights, addressing whether animals can have rights, which animals might have rights, which rights they might have, and how strong these rights might be. While many humans recognize moral duties to animals, debate persists over whether these duties extend beyond consideration of welfare. The chapter examines philosophical and scientific arguments for and against animal rights, highlighting key criteria such as sentience, agency, and relational bonds. It also considers the implications of recognizing animal rights for ethics, law, and policy, emphasizing the need for thoughtful, interdisciplinary approaches.

The Emerging Science of Animal Consciousness

Science (forthcoming)

Kristin Andrews (York University)
Jonathan Birch (London School of Economics)
Jeff Sebo (New York University)

The science of animal consciousness is advancing through investigations of behavioral and neurobiological markers associated with subjective experience in diverse species. Research on honeybee pessimism, cuttlefish planning, and self-recognition in cleaner wrasse fish suggests the possibility of consciousness across the animal kingdom, raising significant implications for neuroscience, ethics, and policy. Despite challenges, emerging frameworks are enabling systematic exploration of dimensions of consciousness such as pain and joy across taxa, providing valuable evidence while emphasizing the need for continued open-minded investigation.

The Emotional Alignment Design Policy

In preparation

Eric Schwitzgebel (University of California, Riverside)
Jeff Sebo (New York University)

The Emotional Alignment Design Policy proposes that artificial entities should be designed to elicit emotional reactions in users that align with their actual moral standing. This essay explores the ethical and practical implications of this policy, emphasizing the risks of eliciting excessive or insufficient emotional engagement. It also discusses complexities such as disagreement and uncertainty about moral standing, asymmetrical risks of error, and the ethical dilemmas surrounding creation, destruction, and bias. The overall aim is to foster ethical treatment of AI systems while navigating moral hazard in an increasingly AI-integrated world.

Ethical Oversight for Insect Research

Zoophilologica (forthcoming)

Toni Sims (New York University)
Jeff Sebo (New York University)

Insects are the most common but least protected animals used in research, with billions used annually in a wide range of studies. Despite widespread assumptions that insects are non-sentient, emerging evidence suggests a realistic possibility of their sentience, with indicators such as learning, memory, and self-protective behaviors. This perspective argues for the establishment of ethical oversight for insect research. Drawing on existing frameworks for vertebrate research, we recommend adapting the “three R’s” (replace, reduce, refine) and developing insect-specific institutional review processes as preliminary measures.



Insects, AI Systems, and the Future of Legal Personhood

Animal Law Review (forthcoming)

Jeff Sebo (New York University)

This paper makes a case for insect and AI legal personhood. Humans share the world with a vast number of small animals. In the future, we might also share the world with sentient or otherwise morally significant AI systems. At present, debates about legal personhood mostly exclude these kinds of individuals. However, this paper argues that our current framework for assessing legal personhood, coupled with our current framework for assessing risk and uncertainty, imply that we should treat these kinds of individuals as legal persons. It also argues that we have good reason to accept this conclusion rather than alter these frameworks.

Interspecies and Intersubstrate Welfare Comparisons

In preparation

Bob Fischer (Texas State University)
Jeff Sebo (New York University)

When we judge that a dog is better off than a pig, we make an interspecies welfare comparison. When we judge that a dog is better off than a robot, we make an intersubstrate welfare comparison (in this case, between a being made out of carbon and a being made out of silicon). This paper considers the challenges associated with comparing the welfare impacts across species and substrates. Then, given that these challenges are likely to be ongoing and that one important purpose of interspecies and intersubstrate welfare comparisons is to inform decision making, it maps out alternative strategies for accomplishing this goal.

Is AI Safety Compatible with AI Welfare?

Under review

Robert Long (Eleos AI)
Jeff Sebo (New York University)
Toni Sims (New York University)

The field of AI safety considers whether and how AI development can be safe and beneficial for humans, and the field of AI welfare considers whether and how AI development can be safe and beneficial for AI systems. There is a *prima facie* tension between these projects, since some measures in AI safety, if deployed against humans and other animals, would raise questions about the ethics of constraint, deception, surveillance, alteration, suffering, death, disenfranchisement, and more. Is there in fact a tension between these projects? We argue that, considering all relevant factors, there is indeed a moderately strong tension—and it deserves more examination.

Plants, Fungi, and Representing Nature

Under review

Kimberly Dill (Santa Clara University)

Jeff Sebo (New York University)

This chapter examines the moral, legal, and political standing of animals, plants, and fungi in the context of climate justice. While the intrinsic value of nonhuman animals is increasingly recognized, skepticism persists about plants and fungi. We explore recent trends in ethics and science, including the “marker method” for assessing consciousness. Highlighting the complexities of plant and fungal cognition, behavior, and interdependence, we argue that these beings warrant further investigation despite the methodological challenges that they raise, and we explore implications of their potential moral significance in a world reshaped by human activity.

Taking AI Welfare Seriously, Continued

In preparation

Robert Long (Eleos AI)

Jeff Sebo (New York University)

and other authors

“Taking AI Welfare Seriously” was part one of a multi-part project. In forthcoming work (once again co-sponsored with Eleos AI), we develop a research agenda that will advance understanding of how likely AI systems are to be welfare subjects and moral patients; how much AI systems matter, if at all; what counts as good or bad for AI systems, if anything; what we owe AI systems, if anything; and how AI safety and AI welfare interact. As with part one, our aim in this forthcoming work is not to settle any of these issues, but rather to explain why these issues are important and difficult and to present early reflections that can inspire further work.

Subjective Experience in AI Systems: What Do AI Researchers and the Public Believe?

In preparation

Noemi Dreksler (Centre for the Governance of AI)

Lucius Caviola (University of Oxford)

David Chalmers (New York University)

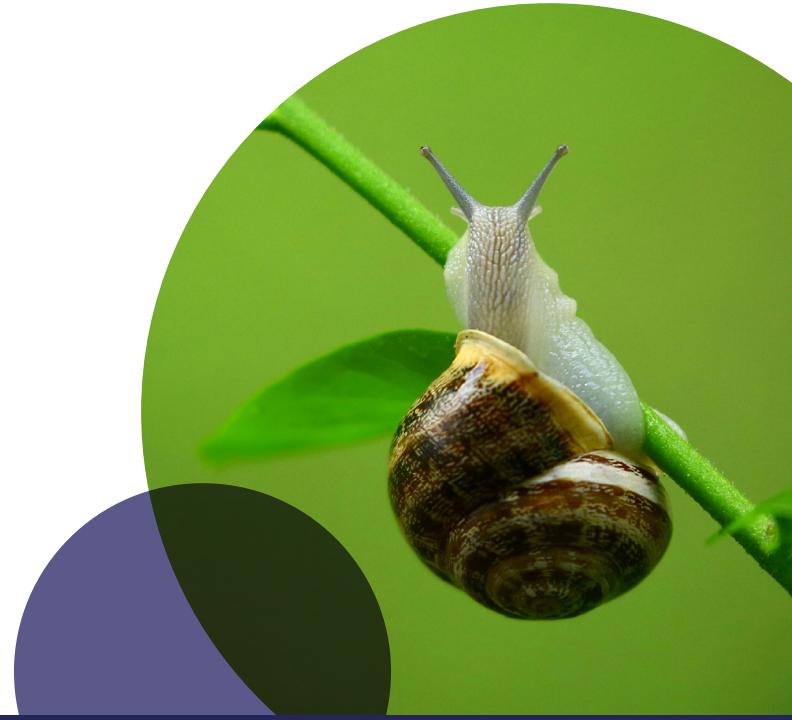
Joshua Lewis (New York University)

Kate Mays (University of Vermont)

Jeff Sebo (New York University)

Philip Waggoner (Colorado School of Mines)

In this paper (co-sponsored with the Centre for the Governance of AI and the Global Risk Behavioral Lab), we surveyed 635 AI researchers and 838 U.S. participants about the possibility of AI systems with subjective experience. Neither group predominantly believes such systems are imminent, but many forecast their existence within this century. Both groups support multidisciplinary expertise in assessing AI subjective experience and favor safeguards now. While support for AI welfare protections was lower than for animal or environmental protection, majorities agreed that AI systems with subjective experience should act ethically and be held accountable.



Our Team



Jeff Sebo,
Director

Associate Professor of Environmental Studies, Affiliated Professor of Bioethics, Medical Ethics, Philosophy, and Law, Director of the Center for Environmental and Animal Protection, Director of the Center for Mind, Ethics, and Policy, and Co-Director of the Wild Animal Welfare Program at NYU



Audrey Becker,
Administrator

Program Administrator at the Center for Environmental and Animal Protection, the Center for Mind, Ethics, and Policy, and the Wild Animal Welfare Program at NYU



Sofia Fogel,
Coordinator

Program Coordinator at the Center for Environmental and Animal Protection, the Center for Mind, Ethics, and Policy, and the Wild Animal Welfare Program at NYU



Toni Sims,
Researcher

Researcher at the Center for Environmental and Animal Protection, the Center for Mind, Ethics, and Policy, and the Wild Animal Welfare Program at NYU

AFFILIATED FACULTY



Ned Block

Silver Professor of Philosophy with secondary appointments in Psychology and Neural Science, Co-Director of the Center for Mind, Brain, and Consciousness at NYU



Sam Bowman

Associate Professor of Linguistics, Data Science, and Computer Science at NYU and, during a leave, a member of technical staff at Anthropic



David Chalmers

University Professor of Philosophy and Neural Science and Co-Director of the Center for Mind, Brain, and Consciousness at NYU



Becca Franks

Assistant Professor of Environmental Studies, Director of the Animal Studies M.A. Program, Director of WATR-lab, and Co-Director of the NYU Wild Animal Welfare Program at NYU



Joshua Lewis

Assistant Professor of Marketing at NYU



S. Matthew Liao

Arthur Zitrin Professor of Bioethics and Director of the Center for Bioethics at NYU



Claudia Passos Ferreira

Assistant Professor of Bioethics at NYU



César Rodríguez-Garavito

Professor of Clinical Law, Chair of the Center for Human Rights and Global Justice, Founding Director of the Earth Rights Research and Action (TERRA) Program, Founding Director of the More-Than-Human (MOTH) Rights Project



Katrina Wyman

Wilf Family Professor of Property Law, Co-Faculty Director of the Frank J. Guarini Center on Environmental, Energy and Land Use Law, and Faculty Director of LLM program in Environmental and Energy Law at NYU School of Law

NEXT
EVENT

**Storytelling about Nonhuman Minds:
A Conversation with Two Acclaimed Authors**
February 7, 2025

This event will feature an engaging conversation between acclaimed authors Peter Brown, author of *The Wild Robot* series, and Eliot Schrefer, author of *Queer Ducks (and Other Animals)*. These authors have crafted stories where nonhumans—whether animals or machines with distinctive qualities—take center stage. Peter and Eliot will discuss how they bring their characters to life, and how these characters challenge human perceptions of sentience, agency, and emotionality. In-person and online guests will be free to ask questions as well. The event will close with a free vegan reception.

To RSVP for this and other events, please visit [our website](#) and sign up for our mailing list.

nonhumanminds@nyu.edu • <https://sites.google.com/nyu.edu/mindethicspolicy>



NYU

**ARTS &
SCIENCE**

Center for Mind,
Ethics, and Policy