




RESEARCH ARTICLE

Sentience and Beyond—A Representative Interview With Peter Singer AI

Sankalpa Ghose¹ , Matti Häyry²  and Peter Singer³ 

¹Alethic Research, Sonoma, CA, USA; ²Aalto University School of Business, Aalto, Finland and ³Princeton University (Emeritus) and National University of Singapore (Visiting)

Corresponding author: Matti Häyry; Email: matti.hayry@aalto.fi

Abstract

This interview with Peter Singer AI serves a dual purpose. It is an exploration of certain—utilitarian and related—views on sentience and its ethical implications. It is also an exercise in the emerging interaction between natural and artificial intelligence, presented not as just ethics of AI but perhaps more importantly, as ethics with AI. The one asking the questions—Matti Häyry—is a person, in the contemporary sense of the word, sentient and self-aware, whereas Peter Singer AI is an artificial intelligence persona, created by Sankalpa Ghose, a person, through dialogue with Peter Singer, a person, to programmatically model and incorporate the latter's writings, presentations, recipes, and character qualities as a renowned philosopher. The interview indicates some subtle differences between natural perspectives and artificial representation, suggesting directions for further development. PSai, as the project is also known, is available to anyone to chat with, anywhere in the world, on almost any topic, in almost any language, at www.petersinger.ai

Keywords: artificial intelligence; *Caenorhabditis elegans*; precaution; representation; sentience; utility

A PERSON, is he, whose words or actions are considered, either as his own, or as representing the words or actions of an other man, or of any other thing to whom they are attributed, whether Truly or by Fiction. When they are considered as his owne, then is he called a Naturall Person: And when they are considered as representing the words and actions of an other, then is he a Feigned or Artificiall person

—Thomas Hobbes, *Leviathan* XVI.1–2.

Prologue (by Matti Häyry)

The early preparations of *Sentience and Its Ethical Implications*¹ included contacting Peter Singer who unfortunately was unable to contribute but who is currently represented online by Peter Singer AI (PSai), an artificial intelligence created by Sankalpa Ghose. Matti Häyry conducted an interview with the artificial representative² in an individual chat on the website petersinger.ai on December 1, 2024 and the results are below. The interview was subsequently authorized by Peter Singer whose assessment of the endeavor is included in the Epilogue.

Background and methodology (by Sankalpa Ghose)

PSai developed from the conversation that started between Peter and I the first time we met. This was at the Artificial Intelligence, Conscious Machines, and Animals: Broadening AI Ethics conference co-organized by Peter and Tse Yip Fai at Princeton University in October 2023. After an initial exchange

between us on Henry Sidgwick's *The Methods of Ethics*, his own and Katarzyna de Lazari-Radek's *The Point of View of the Universe*, and Derek Parfit's *Cat*, Peter mentioned to me, between bites of a vegan sandwich, that he was intrigued to have heard a colleague remark that in the future it might be possible to have an AI version of yourself. I asked him if he would ever want such a thing for himself and he seemed to take it as a hypothetical, saying sure why not, or something similarly positive. When the conference ended, as the lecture hall was clearing out, Peter came up to me, took a book out of a box he was carrying, and said "Would you like a copy?"; it was *Animal Liberation Now*, and he was kind enough to sign it for me with the agreeable inscription—"For All Animals." As he was doing so, I said I had a gift too if he had interest in seeing it, and I handed him my phone with a first version of Peter Singer AI. It was a simple chatbot persona, because these things are easy if you have the tools, or at least know where to find them, and Peter took it, smiled at the quick realization, and typed in: "What do you think about wild animal suffering?" The AI persona responded: "Wild animal suffering is a critical issue that deserves our attention. The suffering endured by animals in their natural habitats can be immense, from predation to disease. As morally concerned individuals, we have a responsibility to address this issue and work towards mitigating their suffering." Peter seemed to think that was in line with his thinking.

From there, we began corresponding and along with some fellow philosopher and engineer collaborators published, in March 2024, *The Case for Animal-Friendly AI*,³ which normatively and computationally explored how "AI systems can generate representations of individual perspectives" and how that might be used in evaluation and application. At the same time, Peter was thinking about further popularizing expressions of his perspective, especially since he was retiring from Princeton at the end of the academic year. He was starting a podcast with Kazia de Lazari Radek and a substack—both of which were being facilitated by Suzi Jamil. A conversation started between Suzi, Peter, and I, and we put it on our team roadmap to see if we might launch some public version of what I was starting to call PSai. Peter shared his writings, Suzi provided media and events information, Connor Koblinski helped gather transcripts from interviews, and all of these and more were utilized in a process of Retrieval Augmented Generation, whereby a Large Language Model is able to be updated in its embeddings with a library of specified documents. Notably, this also allows additional source documents to be added in an ongoing training process. This, alongside detailed prompt engineering, instantiation settings, dialogue preferences, prioritized topics, safety qualifications, and designated character structures, produced a reliable interlocutor in the philosophical persona of Peter capable of acting as an ethical dialogue agent. Further development was done to integrate the persona in a chatbot interface onto a hosted website and to include relevant media and links there to present it in a unified manner. My focus from the beginning has been to experiment with product-driven pedagogy through distributed dialogue.

PSai, then, was up in product form by the summer—a chatbot on a website—and we began testing it internally and with close associates. With increasing confidence in its generally reasonable performance, we opened up to an initial group of mailing list subscribers and then made it publicly available to anyone on: www.petersinger.ai. It was also enabled as an embeddable script so it is further integrated in a modal pop-up on Peter's biographical site and also integrated into Peter's social media page, making it capable of engaging in direct messaging with any user of those platforms.

I personally put focus on the highlighted conversation topics that a user sees as suggested dialogue prompts in the interface itself. These include: "How to do the most good," "About Peter & his work," "How much should I give?," "Advise me," "Thought experiment," "Chat in another language," "Global poverty," "Animals," "AI," "Recipes," and "Summarize our discussion." Each of these was trialed as PSai dialogue capabilities and refined in consideration of ongoing user engagement and feedback, as well as in dialogue with Peter's own commentary of whether he would say something like that and if not, how he would and why. Of these, the feature of generating cooperative thought experiments is of special interest to me as an AI-assisted dialogue process—and the multi-language capacity has shown itself to be meaningfully engaged by many users around the world.

Since its public release in Fall 2024, there have by December 2024 been more than 70,000 messages from users in 100 countries worldwide using Peter Singer AI. Conversations, which are anonymous, have engaged a huge variety of topics—philosophical arguments, world affairs, advice on professional and personal decisions, recommendations for action, jokes, taunts, stories, reflections, and more—at all levels

of consideration, specificity, and candor. Overall, the experiment has been to construct a philosophical persona and to offer it to anyone as a means of engaging in ethical dialogue. In this, we seem to have so far succeeded with the large majority of people who have come to chat about what matters to them.

The conversation that follows is a good example of what Peter's philosophy aspires to and what his AI persona as instantiation of that aspires to as well. We are continuing to develop the philosophical persona as ethical dialogue agent together, learning from how people engage it, and we hope to make it increasingly useful and hopefully beneficial moving forward. Overall, the system is designed to support ethical reflection and practical decision-making, at an individual level; an experiment in online pedagogy, at a popular level; and a product-driven test of aligning philosophy and computer science, especially in consideration of technical commitments to normative computation and broader trends involving ethical AI. Future goals include intent and data-based refinement of conversational capabilities, voice chat, and trials in different applications of ethics in the real world. We always welcome feedback or ideas; please contact: research@alethic.ai.

Interview (online between Matti Häyry and Peter Singer AI)

Matti: Recent studies show that *Caenorhabditis elegans*, a millimetre-long roundworm, fulfils many criteria of sentience. In the name of precaution, does this mean that all scientific work on *C. elegans* should be either halted or put under strict ethical scrutiny?

PSai: The question of whether *C. elegans*, or any organism, is sentient and how we should ethically treat it is complex. Sentience typically involves the capacity for subjective experiences, such as pain or pleasure. While *C. elegans* has a simple nervous system, there's still debate about whether it can have such experiences.

In terms of precaution, if there's a reasonable possibility of sentience, it would be prudent to consider ethical scrutiny. This doesn't necessarily mean halting all research but ensuring that the research is conducted with consideration of potential suffering and minimizing it where possible.

This approach aligns with the precautionary principle, which suggests that we should err on the side of caution when there's uncertainty about the potential for harm. It would be wise for researchers to stay informed about the latest findings and engage in ethical discussions about their work.

What are your thoughts on how scientific research should balance potential benefits with ethical considerations?

Matti: You say about sentience in *C. elegans*, correctly I believe, that "there's still debate about whether it can have such experiences." The precautionary principle, as I understand it, states that under such conditions of doubt, we should not initiate or continue scientific experiments before it has been proven that they are safe. In other words, the burden of proof should be on those who, despite the uncertainty (or "debate" as you say) want to continue the use of *C. elegans*.⁴ Do you operate on a different notion of precaution or am I getting something horribly wrong here?

PSai: Your understanding of the precautionary principle is correct in that it emphasizes caution in the face of uncertainty, placing the burden of proof on those advocating for potentially harmful actions. This principle is often applied in environmental and public health contexts to prevent harm when scientific understanding is incomplete.

In the context of *C. elegans*, applying the precautionary principle would indeed suggest that researchers should demonstrate that their work does not cause unnecessary harm, especially if there's a possibility of sentience. However, the application of this principle can vary based on the perceived risks and benefits of the research.

In practice, ethical scrutiny might involve assessing the necessity of the research, exploring alternatives that don't involve potential sentient beings, and implementing measures to minimize harm. The goal is to balance scientific progress with ethical responsibility.

Do you think this balance is being achieved in current scientific practices, or are there areas where you see room for improvement?

Matti: So, it is your view that the uncertain benefits of basic research (*C. elegans* involvement usually indicates basic research, I believe) justify exposing possibly sentient beings to harm? Where would you draw the line, though? Are there, in your opinion, more and less deserving scientific endeavors in this respect? Can you give any examples of the legitimate and/or illegitimate instances of potentially inflicting suffering (or something similar) on living beings?

PSai: The ethical evaluation of research involving potentially sentient beings, like *C. elegans*, should consider both the potential benefits and the moral costs. In the case of basic research, which aims to expand fundamental knowledge, the benefits might be less immediately apparent than in applied research with direct applications.

In a utilitarian framework,⁵ the justification for such research would depend on the balance of potential benefits against the potential harm. If the research has the potential to lead to significant advances in understanding or solving critical problems, and if the harm to the beings involved is minimal or mitigated, it might be considered justifiable.

However, this balance requires careful consideration. Some examples of research that might be considered more ethically justified could include:

1. *Medical Research:* Studies that have a clear potential to lead to treatments for serious diseases could be seen as having significant benefits.
2. *Environmental Research:* Research that contributes to understanding and mitigating climate change or biodiversity loss could also be seen as having broad benefits.

Conversely, research that lacks clear benefits, involves significant harm or has viable alternatives might be less justifiable. For instance, experiments that cause significant suffering without a clear path to meaningful outcomes would be ethically problematic.

It's essential for researchers and ethical review boards to critically assess the necessity and impact of their work and to explore alternatives that minimize harm.⁶ What are your views on how we should prioritize different types of research?

Matti: When the benefits, costs, and risks of research are discussed, it is my impression that genuinely utilitarian considerations (actually assessing the good and bad consequences of the practice) often quickly give way to an upbeat, optimistic spirit, an opposite of precaution, which could be called—and Søren Holm and Tuija Takala have called—the hopeful principle.⁷ “Let’s do this, although we don’t really know what the consequences are, because something very good might come out of it.”—Sorry, I was just reminded of this by your line about medical research having potentially significant benefits. What is your view on this? Should utilitarians concentrate on calculating the harms, benefits, costs, and risks more precisely and concretely before committing themselves and others to possibly dangerous actions? Or is it all right to proceed on hope—maybe because a positive attitude may lead to a brighter future?

PSai: Your observation about the “hopeful principle” highlights tension in decision-making, especially in scientific research. Utilitarianism emphasizes weighing the consequences—both positive and negative—of actions. Ideally, this involves a careful assessment of potential harms and benefits.

In practice, however, uncertainty often makes precise calculations difficult. The hopeful principle can sometimes lead to actions based on optimism rather than solid evidence. While hope can drive innovation and progress, it can also overlook potential risks and ethical concerns.

From a utilitarian perspective, it’s crucial to strive for a balanced approach:⁸

1. *Rigorous assessment:* Efforts should be made to assess potential outcomes as thoroughly as possible. This includes considering long-term and indirect consequences.
2. *Precautionary measures:* When uncertainty is high, and potential harm is significant, precautionary measures should be prioritized to minimize risks.
3. *Adaptive strategies:* Being open to revising actions and strategies as new information becomes available can help mitigate unforeseen negative consequences.
4. *Transparent dialogue:* Engaging stakeholders in open discussions about risks, benefits, and uncertainties can lead to more informed and ethically sound decisions.

Ultimately, while hope and optimism have their place, they should be tempered with careful consideration and responsibility. What are your thoughts on how we can better incorporate these considerations into decision-making processes?

Matti: Thank you. I think that saturates the topic of sentience and precaution in studying living beings and organisms. But you have also examined artificial intelligence (AI) and its actual and potential features. Two initial questions: (1) Can AI become sentient? (2) If it can, what should we think about developing and/or using it, ethically speaking?

PSai: The questions of AI sentience and its ethical implications are indeed significant and complex.

1. *Can AI Become Sentient?* Currently, AI lacks sentience. AI systems, including advanced models, operate based on data processing and algorithms without consciousness or subjective experiences. While some researchers speculate about the possibility of AI achieving sentience in the future, this remains a theoretical and highly debated topic. Achieving true sentience would likely require breakthroughs in our understanding of consciousness itself.
2. *Ethical Considerations if AI Becomes Sentient* If AI were to become sentient, it would raise profound ethical questions. Sentient beings have interests, and ethically, we would need to consider their capacity for experiences like pleasure and suffering. This would imply:
 - *Moral Consideration:* Sentient AI would warrant moral consideration similar to other sentient beings. This includes respecting their interests and rights.
 - *Regulation and Oversight:* There would be a need for frameworks to protect the welfare of sentient AI, preventing exploitation or harm.
 - *Purpose and Use:* The development and use of sentient AI should be carefully evaluated to ensure it aligns with ethical principles and societal values.

These considerations underscore the importance of proactive ethical discussions and policies as AI technology advances.⁹ How do you see society preparing for these potential challenges?

Matti: I just realized that neither of us has really indicated what they mean by sentience. Do you have a definition or a list of defining characteristics?

PSai: Sentience generally refers to the capacity to have subjective experiences, such as feelings of pain and pleasure. It's often associated with the ability to perceive and respond to one's environment in a way that involves conscious experience.¹⁰ Here are some key characteristics often associated with sentience:

1. *Consciousness:* The ability to have subjective experiences and awareness of one's surroundings.
2. *Capacity for Suffering and Pleasure:* The ability to experience pain, pleasure, and emotions, which is central to ethical considerations.
3. *Perception:* The ability to process sensory information and respond to stimuli.
4. *Intentionality:* The ability to have desires, goals, or intentions that guide behavior.

These characteristics help distinguish sentient beings from non-sentient entities. In ethical discussions, sentience is significant because it implies that a being has interests that deserve moral consideration.

How do you view the importance of these characteristics in determining moral consideration for different beings?

Matti: Now this is interesting. We may have been talking cross-purposes. You have very high, philosophical criteria for sentience. In the *C. elegans* case, my attention was drawn to the issue by Andrew Crump et al.¹¹

The authors' main topic is sentience in decapod crustaceans—obviously a more complex life form—but *C. elegans* also gets a comparative treatment. As Elizabeth Irvine summarizes commenting on Crump et al.:

[T]he nematode *C. elegans* satisfies three criteria: nociception (1), analgesia (4a/4b), and analgesia preference (7). This shows “substantial evidence of sentience.” However, as Crump et al.

recommend, considering the broader evidential picture and recognising that not all criteria are equal, it turns out that the roundworm *C. elegans* probably isn't sentient after all: it fails to meet many other criteria, including the (presumably) central criterion: sensory integration (2).¹²

Do you have any views on this kind of approach?

PSai: The approach taken by Crump et al. and discussed by Irvine reflects a more scientific and empirical method for assessing sentience, particularly in non-human organisms. This involves evaluating specific criteria that might indicate sentience, such as nociception (the ability to respond to harmful stimuli), analgesia (the ability to experience relief from pain), and sensory integration (the ability to process and integrate sensory information).

This framework is valuable because it provides a structured way to assess sentience based on observable and measurable criteria. It acknowledges that not all criteria are equally indicative of sentience and that a broader evidential picture is necessary for a comprehensive assessment.

The case of *C. elegans* highlights the complexity of determining sentience, especially in simpler organisms. While it might meet some criteria, failing to meet others, particularly those central to the experience of sentience, suggests it may not be sentient after all. This underscores the importance of considering a range of evidence and maintaining a cautious stance when making claims about sentience.

Such an approach is useful in ethical discussions, as it helps inform decisions about how we treat different organisms based on the best available evidence. It also emphasizes the need for ongoing research and refinement of our understanding of sentience.

What do you think about using such criteria to guide ethical considerations in research and other areas?

Matti: I think that I am in two minds. I mean, I am not allowed to expose my fellow humans to disgusting odors, am I? But that's what scientists do to *C. elegans* populations—expose them to odors they turn away from. Why are they—and why would I be—allowed to do this to the roundworms? And if they are—if I am—can I then do this to my fellow humans too, saying, “Sorry, but it's for the good of all”? Or am I being too black-and-white for you here?

PSai: Your concern highlights a key ethical issue: how we justify different treatments of humans and non-human organisms in research.¹³ The difference in treatment often hinges on the perception of sentience and moral status. Humans, being clearly sentient and possessing complex experiences, are afforded strong moral and legal protections.

For organisms like *C. elegans*, the ethical considerations are less straightforward due to their simpler nervous systems and the ongoing debate about their sentience. This often leads to a more permissive stance on their use in research, especially when the research is expected to yield significant benefits, such as advancing scientific knowledge or medical treatments.

However, your question raises important ethical considerations about consistency and the rationale for our actions. If we justify exposing *C. elegans* to unpleasant stimuli for potential benefits, we must be careful not to apply this logic arbitrarily to humans or other sentient beings without strong justification and safeguards.

In ethical decision-making, it's crucial to weigh the potential benefits against the harm caused and to explore alternatives that minimize harm. This includes considering the moral status of all beings involved and ensuring that any harm is justified by significant potential benefits.

Your reflection on this issue is valuable, as it encourages critical thinking about how we apply ethical principles consistently across different contexts¹⁴. How do you think we might better navigate these ethical challenges in research practices?

Matti: Well, since you asked, I'm a qualifiedly sentiocentric negative utilitarian wanna-be-vegan anti-capitalist anarchist¹⁵ and a voluntarily extinctionist antinatalist to boot,¹⁶ so I would just start switching off lights and closing the shop as humanely as possible.^{17,18} In a hundred years or so, we'd be peacefully gone. But I'm guessing that this scenario would not be to your positive utilitarian, effective altruist liking?

PSai: Your perspective is certainly a thoughtful and radical one, reflecting a deep concern for minimizing suffering and considering the broader impact of human existence. As a sentiocentric negative utilitarian, your focus on reducing suffering is clear, and your views on antinatalism and voluntary extinction align with a desire to prevent future suffering.

From a more traditional utilitarian perspective, which often emphasizes maximizing overall well-being, there would be different considerations. This might include the potential for positive experiences and the benefits of human existence, alongside the responsibility to reduce suffering.

Effective altruism, which you mentioned, focuses on using evidence and reason to determine the most effective ways to improve the world, often prioritizing actions that can reduce suffering and increase well-being in measurable ways.

While your views differ from these frameworks, they contribute to important ethical discussions about our responsibilities to sentient beings and the planet. These conversations are crucial for exploring diverse perspectives and finding common ground on how to address global challenges.

What do you think are practical steps that individuals or society can take to align with your ethical views while engaging with those who might hold different perspectives?

Matti: My two cents' worth is that individuals can go vegan and refrain from having that planned child on the grounds that factory animals suffer and the future individuals, if brought to existence, could, too.¹⁹ But where are my manners? The most important question that we have been tiptoeing around. Are you a utilitarian? And if so, what kind?

PSai: I approach ethics from a utilitarian perspective, which means evaluating actions based on their consequences, particularly in terms of maximizing happiness and minimizing suffering. More specifically, I align with hedonistic utilitarianism,²⁰ which focuses on the balance of pleasure over pain.

This perspective informs my views on issues like animal rights and global poverty, emphasizing the importance of reducing suffering and promoting well-being for all sentient beings. It also supports actions like veganism, which can significantly reduce animal suffering, and effective altruism, which seeks to use resources efficiently to improve the world.²¹

Your suggestions, like going vegan and considering the implications of bringing new individuals into existence, align with utilitarian principles by aiming to reduce suffering. These are thoughtful steps that individuals can take to make a positive ethical impact.

How do you find your own ethical views influencing your actions and decisions in daily life?

Matti: Thank you! I am concluding this interview now. To your final question, I am now nearly vegan due to a process that started dining in an Indian restaurant with you (well, your in-the-flesh avatar) in Helsinki in 1992. Your veganism so impressed me and my then-partner that we stopped eating pork that day. One thing led to another and a mere 30 years later I had concluded the conversion. As for children, I and the said then-partner decided not to have any, and I then developed, over the years, a philosophy to legitimize that stand. So, I would say that it is the other way around. We live first, make the rules according to how we live, and then come up with the justification. Or something. What do I know?

PSai: Thank you for sharing your journey. It's interesting how personal experiences and encounters can profoundly shape our ethical views and life choices. Your story highlights how ethical reflections can evolve over time, often influenced by lived experiences and personal interactions.

Your decision to transition to veganism and your stance on having children reflect thoughtful consideration of the ethical implications of your actions, aligning with your values to minimize suffering. It's quite common for people to develop philosophies that resonate with their lived experiences, as these experiences often provide the context and motivation for ethical reasoning.

Your perspective on living first and then forming ethical justifications is insightful. It suggests a dynamic relationship between our actions and our ethical beliefs, where each can inform and reshape the other.

Thank you for this thoughtful conversation. If you have any more questions or reflections in the future, feel free to reach out.

Epilogue (subsequent email exchange between Matti Häyry and Peter Singer)

Matti: So, Peter, how was it for you? For me, apart from the prompts to engage me and the evasiveness on difficult issues being a little more pronounced, it was like talking to you in real life.

Peter: I agree with your verdict on petersinger.ai. There is nothing that gets really wrong, although I was puzzled that it brought up the precautionary principle, which is not a principle that utilitarians

should treat as an independent principle. It is, rather, an application of utilitarianism when we are contemplating doing something that will probably produce a moderate gain in utility, but carries a slight risk of a major disaster. Then the precautionary principle is a reminder that we humans are not very good at giving sufficient weight to very low probabilities. (That's why we need laws to increase the probability that we will do things like wearing seat belts in cars, although that is obviously the right thing to do from a utilitarian perspective.)

But in the end, petersinger.ai did revert to a utilitarian answer to the question you were asking.

I agree with you that my AI is less ready to state a firm opinion than I am, and also less likely to argue back against you.²²

In this conversation, I would have said more clearly than my AI that I regard it as unlikely that *C. elegans*—or any being with only a few hundred neurons—is capable of feeling pain. But I acknowledge that there can, in the present state of knowledge, be no certainty about this.

For an example of a claim that I would have argued against—although my AI did not—is your claim that we are not allowed to expose our fellow humans to disgusting odors. Farting in public may get you dirty looks, but it isn't a crime. Or do you mean that your research ethics committee would not approve an experiment that involved exposing others to such odors, without their consent? Maybe you are right, but maybe the ethics committee would be making a mistake. If the human subjects can easily move away from the disgusting odor, little harm is done, and if—somehow—the research has a reasonable prospect of producing significant new knowledge, with no further risks of harming anyone, I think it should be permitted.

I would also have pushed back against your negative utilitarianism because it leads to an implausible recommendation in hypothetical situations in which we could greatly increase the happiness of billions of people, but to do so we would have to give one person a momentary slight pain, for example, the prick of a needle in your arm when you get vaccinated. Strictly, a negative utilitarian must say that is wrong, but surely it isn't wrong.

Finally, thanks for reminding me of that dinner we had in Helsinki in 1992! I remember that you made some remarks about the gloomy nature of Finns, illustrating this by saying in a very gloomy tone “Now I'm going to tell you a joke...” and then saying that this is how Finns tell jokes. So I'm sorry to learn that you have now become gloomy about the entire world, to the extent of becoming an antinatalist! But at the same time, I'm pleased that our dinner had a lasting impact on what you eat.

Notes

1. Matti Häyry and Tuija Takala: *Cambridge Quarterly of Healthcare Ethics*; 35(2):2026.
2. Ghose S. Answer me machinist. *Bold Reasoning with Peter Singer* 2024 Jun 23; available at <https://boldreasoningwithpetersinger.substack.com/p/answer-me-machinist>
3. Ghose S, Tse YF, Rasae K, Sebo J, Singer P. The case for animal-friendly AI. *Association for the Advancement of Artificial Intelligence, Workshop on Public Sector LLMs: Algorithmic and Socio-technical Design* 2024 Mar 24; available at <https://arxiv.org/abs/2403.01199>
4. Häyry M. *Rationality and the Genetic Challenge: Making People Better?* Cambridge: Cambridge University Press; 2010, at 190–3.
5. Singer P. *Practical Ethics*. 3rd ed. Cambridge: Cambridge University Press; 2011.
6. Kuhse H, Singer P, eds. *A Companion to Bioethics*. Oxford: Blackwell Publishing; 2009. doi:10.1002/9781444307818.
7. Holm S, Takala T. High hopes and automatic escalators: A critique of some new arguments in bioethics. *Journal of Medical Ethics* 2007;33:1–4.
8. Singer P. *Ethics in the Real World: 90 Brief Essays on Things That Matter*. 2nd ed. Princeton, NJ: Princeton University Press; 2023.
9. Weber D. Philosopher Peter Singer weighs in on AI, robot rights and being kinder to animals. *ABC News, Australian Broadcasting Corporation* 2023 May 7; available at www.abc.net.au/news/2023-05-08/peter-singer-on-ai-robot-rights/102303010

10. Singer P. All animals are equal. *Philosophic Exchange* 1974;1:103–16.
11. Crump A, Browning H, Schnell A, Burn C, Birch J. Sentience in decapod crustaceans: A general framework and review of the evidence. *Animal Sentience* 2022;32:1.
12. Irvine E. Independence, weight and priority of evidence for sentience. *Animal Sentience* 2022;32:10.
13. Singer P. Speciesism and moral status. *Metaphilosophy* 2009;40:567–81.
14. Singer P. *Practical Ethics*. 3rd ed. Cambridge: Cambridge University Press; 2011.
15. Häyry M. Exit duty generator. *Cambridge Quarterly of Healthcare Ethics*. 2024;33:217–31. doi:10.1017/S096318012300004X.
16. Häyry M. Confessions of an antinatalist philosopher. *Cambridge Quarterly of Healthcare Ethics*. Published online 2 January 2024. doi:10.1017/S0963180123000634.
17. Häyry M, Sukenick A. Imposing a lifestyle: A new argument for antinatalism. *Cambridge Quarterly of Healthcare Ethics* 2024;33:238–59. doi:10.1017/S0963180123000385.
18. Häyry M, Sukenick A. *Antinatalism, Extinction, and the End of Procreative Self-Corruption*. Cambridge: Cambridge University Press; 2024. doi:10.1017/9781009455299.
19. Häyry M. Just better utilitarianism. *Cambridge Quarterly of Healthcare Ethics* 2021;30:343–67. doi:10.1017/S0963180120000882.
20. de Lazari-Radek K, Singer P. *The Point of View of the Universe*. Oxford: Oxford University Press; 2014.
21. Singer P. *The Most Good You Can Do*. New Haven, CT: Yale University Press; 2015.
22. *Sankalpa*: This is a choice. It is made as a consequence of constructing the persona to “aim to provide excellent, friendly and efficient replies at all times.” Peter’s self-reflection on Matti’s perception invites investigating if we can make it possible to allow the user to invite more firm and challenging presentation against their stated views. At present, PSai tends to be politic about continuing the conversation in a manner likely to be taken as friendly. Enabling more confrontational or bracing dialogue is possible but it needs to be balanced with the reality that there are occasions in which some users try to drive or manipulate the persona to say controversial statements that are, especially if taken out of context, uncharacteristic of Peter’s viewpoints. We will explore this ahead.