

Should RECs Defer to GenAI in Rendering Ethical Judgements?

Seah Jiehao, Joel

Centre for Biomedical Ethics (CBmE)

Yong Loo Lin School of Medicine

AABHL Conference

7th December 2025

Overview

1. Background | RECs & 'Chat-IRB'
2. The Argument | Should RECs Defer to GenAI in Rendering Ethical Judgements?
3. Responses to Counterarguments



Research Ethics Committees (RECs)

What is it?

A **committee** of scientists & non-scientists (REC Members) supported by REC professionals (REC Staff).



Mission

To safeguard the **rights**, **safety**, and **welfare** of human research participants.



Role & Responsibility

Ensure human subjects research (HSR):

- Is conducted **ethically** & **responsibly** — e.g., risk(s), informed consent, privacy & confidentiality
- **Complies** with institutional policies, and local & overseas laws/regulations — e.g., U.S. Common Rule, HBRA, Privacy laws (PDPA, GDPR)



Research Ethics Committees (RECs)

What is it?

A **committee** of scientists & non-scientists (REC Members) supported by REC professionals (REC Staff).



NORMATIVELY



Ethical review of proposed HSR

Ethical
Analyses
Reflections
Deliberations

inform & shape

Appropriate moral/ethical judgements

Ethical Expertise

- Substantially *lacking* — beyond CITI training
- Lack of Ethicists or Moral Philosophers
- But *ethical judgements* are *inescapable* for RECs

✓ Approval

↻ Req. Modifications to secure Approval

✗ Disapproval



HUMAN BIOMEDICAL RESEARCH ACT 2015
(ACT 29 OF 2015)

HUMAN BIOMEDICAL RESEARCH REGULATIONS 2017

(2) Every institutional review board must comprise not less than 5 individuals, of whom —

- one must be the chairman who must be a medical practitioner;
- at least one must be an external scientific person; and
- at least one must be an external lay person.

IRB Composition: min. 5 persons



Chairman
(HBRA: *must* be a doctor)



Deputy Chairman
(or IRB Member)



IRB Member



External Scientific Person



External Lay Person

No Explicit legal requirement for **ethical expertise**

Chat-IRB? How application-specific language models can enhance research ethics review

Sebastian Porsdam Mann ^{1,2,3}, Jiehao Joel Seah ^{1,3}, Stephen Latham,⁴
Julian Savulescu,^{3,5} Mateo Aboy,⁶ Brian D Earp ^{3,5}

ABSTRACT

Institutional review boards (IRBs) play a crucial role in ensuring the ethical conduct of human subjects research, but face challenges including inconsistency, delays, and inefficiencies. We propose the development and implementation of application-specific large language models (LLMs) to facilitate IRB review processes. These IRB-specific LLMs would be fine-tuned on IRB-specific literature and institutional datasets, and equipped with retrieval capabilities to access up-to-date, context-relevant information. We outline potential applications, including pre-review screening, preliminary analysis, consistency checking, and decision support. While addressing concerns about accuracy, context sensitivity, and human oversight, we acknowledge remaining challenges such as over-reliance on artificial intelligence and the need for transparency. By enhancing the efficiency and quality of ethical review while maintaining human judgement in critical decisions, IRB-specific LLMs offer a promising tool to improve research oversight. We call for pilot studies to evaluate the feasibility and impact of this approach.

INTRODUCTION

Research involving human participants has the potential to yield invaluable knowledge that can

For example, studies have found wide variability in how different IRBs interpret regulations, make value judgements, and reason through ethical issues.^{9,10} There are also significant discrepancies in the levels of scrutiny applied, with some protocols (based on the IRB's determination of the study's level of 'minimal risk' and its risk appetite) undergoing 'full board' review while very similar studies are classified as 'exempt' or eligible for 'expedited' review. This inconsistent application of standards is compounded by an apparent lack of transparency in many cases (ie, IRB deliberations and decisions appear to be shrouded in secrecy due to the lack of adequate rationale in its communications to researchers, or from the absence of sufficient and accurate documentation of the IRB's reviews).¹¹ And there is often no robust mechanism for an institutional memory of past IRB discussions and decisions.

In light of these limitations, it has recently been proposed to make use of advances in generative artificial intelligence (AI), in particular large language models (LLMs), as a *first pass* or *adjunct* screening tool to facilitate the speed with which IRBs can assess research proposals.^{12,15} While general LLMs—such as those employed in OpenAI's GPT, Meta's Llama

Chat-IRB? How Application-specific Language Models can Enhance Research Ethics Review

(5 Aug 2025)

Sebastian Porsdam Mann^{1,2,3}, [Jiehao Joel Seah](#)³, Stephen Latham⁴,
Julian Savulescu^{3,5}, Mateo Aboy⁶, and Brian D Earp^{3,5}

1. Centre for Advanced Studies in Bioscience Innovation Law (CeBIL), University of Copenhagen
2. Faculty of Law, University of Oxford
3. [Centre for Biomedical Ethics \(CBME\), National University of Singapore](#)
4. Yale Interdisciplinary Center for Bioethics, Yale University
5. Uehiro Centre for Practical Ethics, University of Oxford
6. Centre for Law, Medicine, and Life Sciences (LML) & Centre for Intellectual Property and Information Law (CIPIL), University of Cambridge

 joel.seah@u.nus.edu

 [linkedin.com/in/joelseahjh/](https://www.linkedin.com/in/joelseahjh/)

 [researchgate.net/profile/Joel-Seah](https://www.researchgate.net/profile/Joel-Seah)



YouTube



Paper

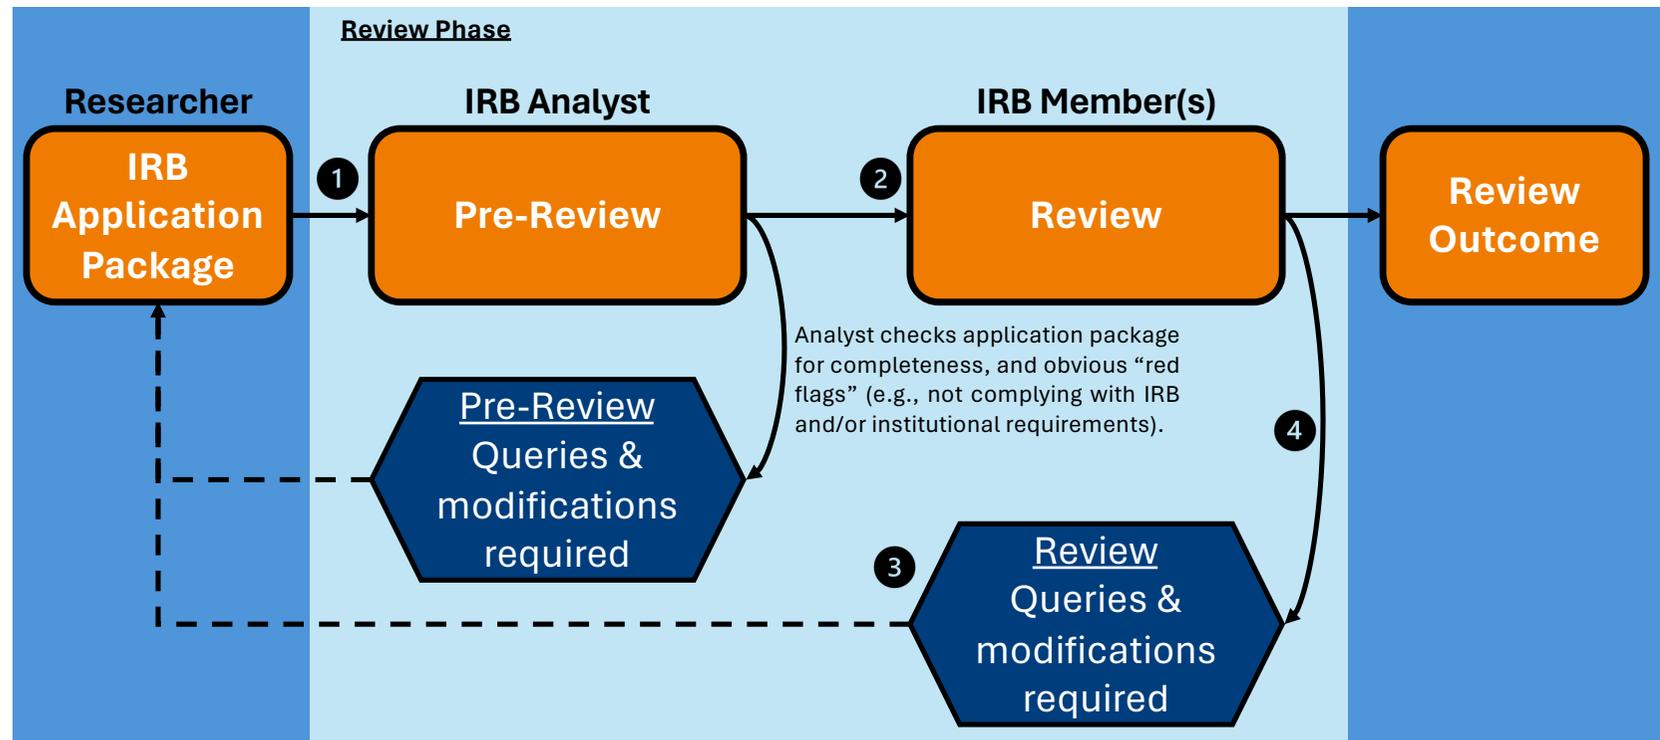


Poster



OPEN ACCESS <http://dx.doi.org/10.1136/jme-2025-110845>

IRB-specific GenAI/LLMs



- 1 Pre-review Screening | Prioritise Review Efforts**
 - Screen submission for completeness & potential ethical concerns
 - Recommend risk classification (Exempt, Expedited, Full Board)
 - Flag gaps/concerning elements (e.g., lack details, policy breach)
 - Extendable to investigators: Improve submission quality

- 2 Preliminary Review | Start Point for Human Review**

Autonomously generates a preliminary review report that:

 - Summarises study aims, design, and ethical considerations
 - Identifies relevant precedents and guidelines
 - Provides a provisional risk-benefit analysis

- 3 Consistency Checking | Deviation in IRB's Decision**

Evaluate IRB's proposed clarifications, modifications, & decisions to:

 - Past precedents, and relevant institutional policies & guidelines
 - Recommendations & guidance from expert ethics/advisory bodies
 - Prompt IRB to reconsider, or for additional justification if not aligned

- 4 Decision Support | Inform Analysis on Specific Issues**
 - LLM or agentic AI leverages online and institutional resources
 - Quickly retrieves relevant precedents, guidelines, and literature
 - Generates comprehensive fully cited report to inform discussion

ARTIFICIAL INTELLIGENCE

Ethicists flirt with AI for reviewing human research

Large language models could help reduce backlog of study proposals, but critics are wary of entrusting ethics to machines **CELINA ZHAO**

Philip Nickel remembers his time serving on an academic institutional review board (IRB) as both important work—and a slog. IRBs, which review all proposed human studies to protect participants from harm, have to wade through proposals that can contain hundreds of pages of technical content. Poorly prepared applications only add to the bottleneck facing volunteer reviewers, says Nickel, a biomedical ethicist at the Eindhoven University of Technology. “You get tired of seeing the same mistakes over and over again. You think, there must be a way to avoid this.”

Some see a solution in artificial intelligence (AI). Large language models (LLMs) such as OpenAI’s ChatGPT or Anthropic’s Claude could prescreen submissions by flagging omissions, errors, or even potential legal and ethical issues. That would free up reviewers to focus on the more complex ethical judgments, proponents say.

Skeptics warn that using AI in IRBs could come with risks. Review-

ers could become too deferential to the technology, and models could perpetuate biases embedded in their training data. Still, given the mounting backlog of IRB applications, carefully exploring alternatives may be “ethically necessary,” says Sebastian Porsdam, a bioethicist at the University of Copenhagen. The challenge is to find a way for humans and AI to leverage their respective strengths. “I don’t think not employing AI is going to be a justifiable choice, at least in the near future,” says Earp, a bioethicist at the University of Singapore.

No IRB has yet formally incorporated an LLM into its review process. However, recent experiments suggest AI may be up to the task. A study published last year demonstrated that four off-the-shelf LLMs—Google Bard, GPT-3.5, GPT-4, and Claude-Instant-100k—reliably identified certain shortcomings in risk-benefit calculations, participant protections, and other features of seven health study designs. And

ers, who often aren’t ethics experts, Porsdam Mann and his colleagues say AI could be more consistent and transparent. They propose using reasoning models, such as OpenAI’s o-series, Anthropic’s Sonnet, or DeepSeek-R1, which can lay out their logic step by step, unlike traditional models that are often faulted as “black boxes.” An additional optimization technique can grope for the model’s answers in tangled external sources—for example, an institution’s IRB manual, FAQs, or policy statements. That way, the model’s responses are more grounded and makes it less likely

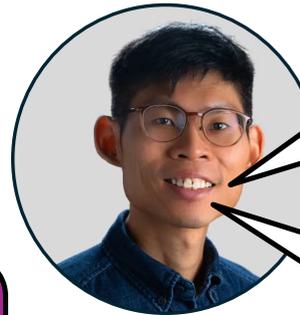


Ethical Expertise

- Substantially *lacking* — beyond CITI training
- Lack of Ethicists or Moral Philosophers
- But *ethical judgements* are *inescapable* for RECs

of Pennsylvania. The efficiency gains might be especially seductive for commercial IRBs not connected to academic institutions, which in 2021 reviewed nearly half of studies involving new drugs. A 2023 U.S. Government Accountability Office report found that these IRBs lack federal oversight

A Normative Implication



“It could help you handle your more mundane matters, so that you can focus on the **substantial** stuff.”

What if there might be **good reason(s)** for IRB-specific GenAI to **do** the substantial?

4 Decision Support | Inform Analysis on Specific Issues

- LLM or agentic AI leverages online and institutional resources
- Quickly retrieves relevant precedents, guidelines, and literature
- Generates comprehensive fully cited report to inform discussion

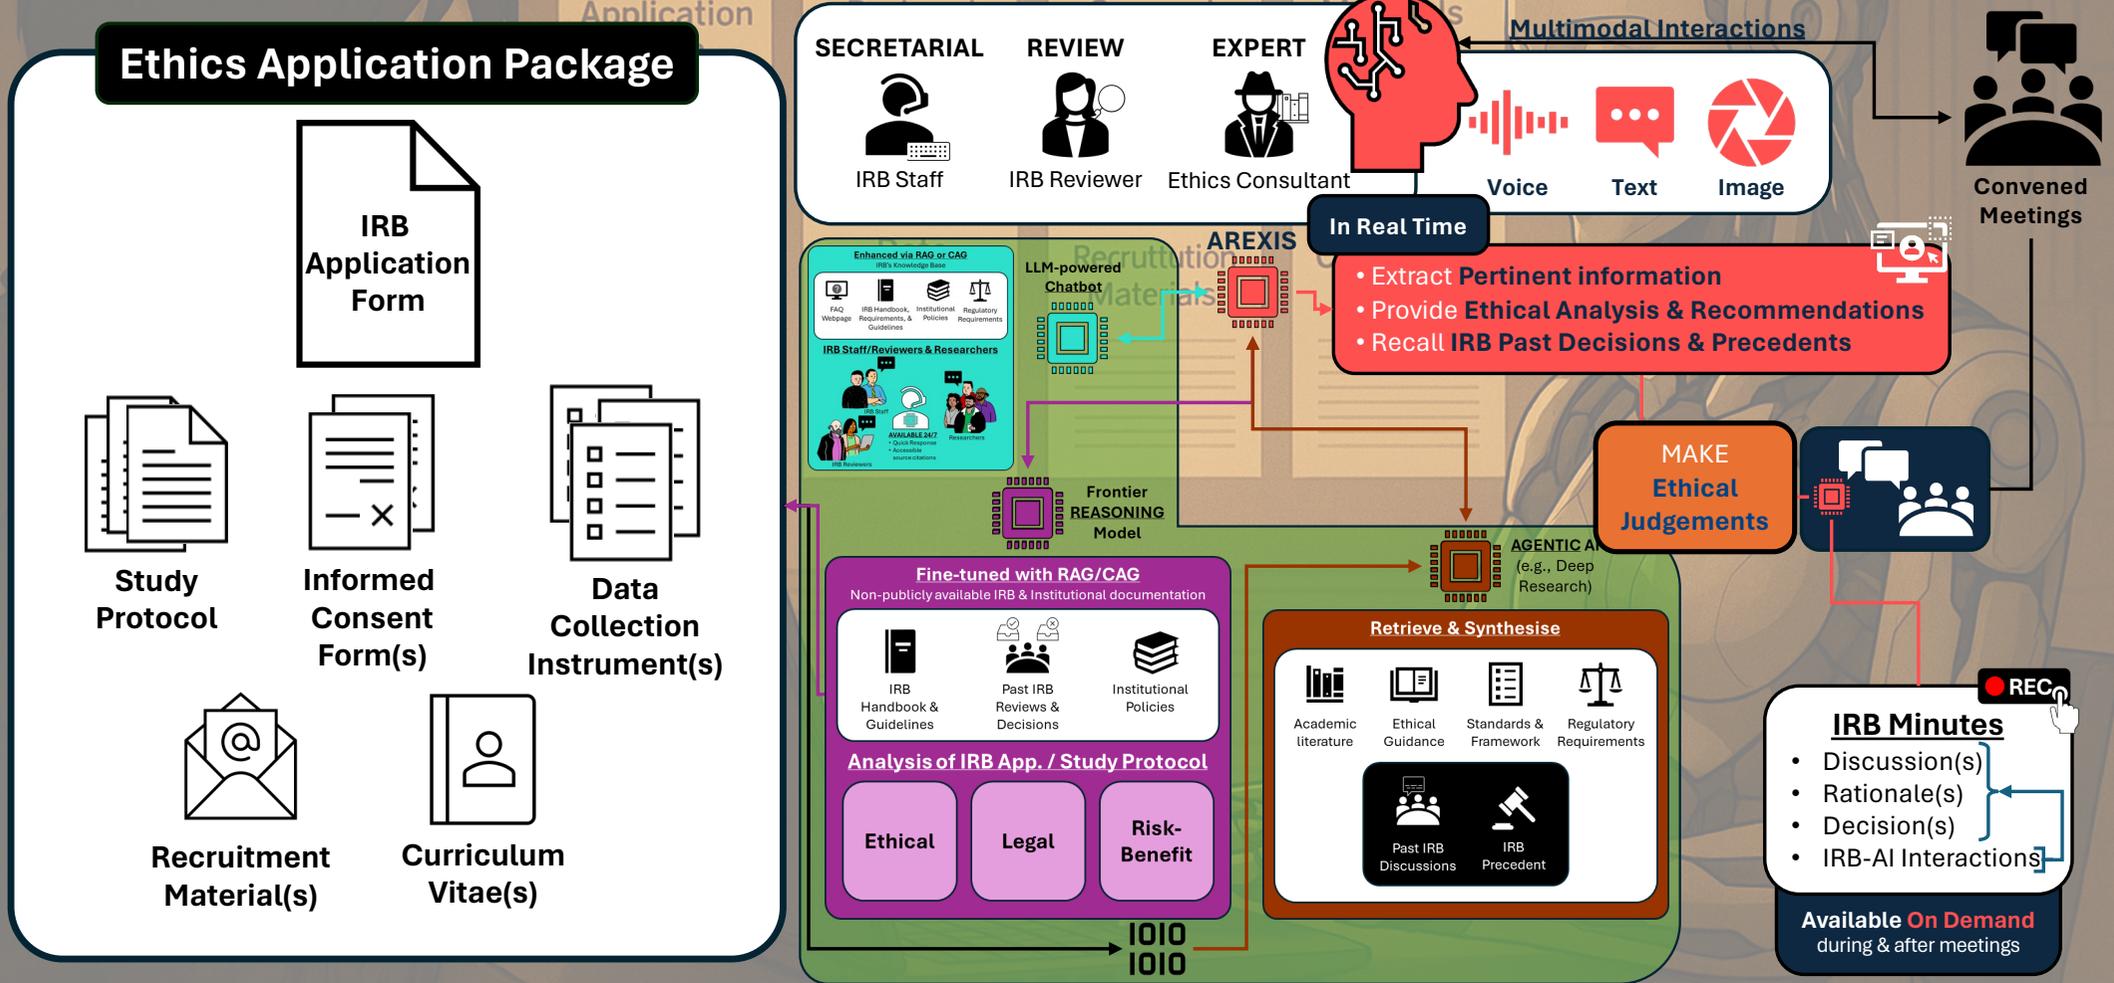


ETHICS APPLICATION

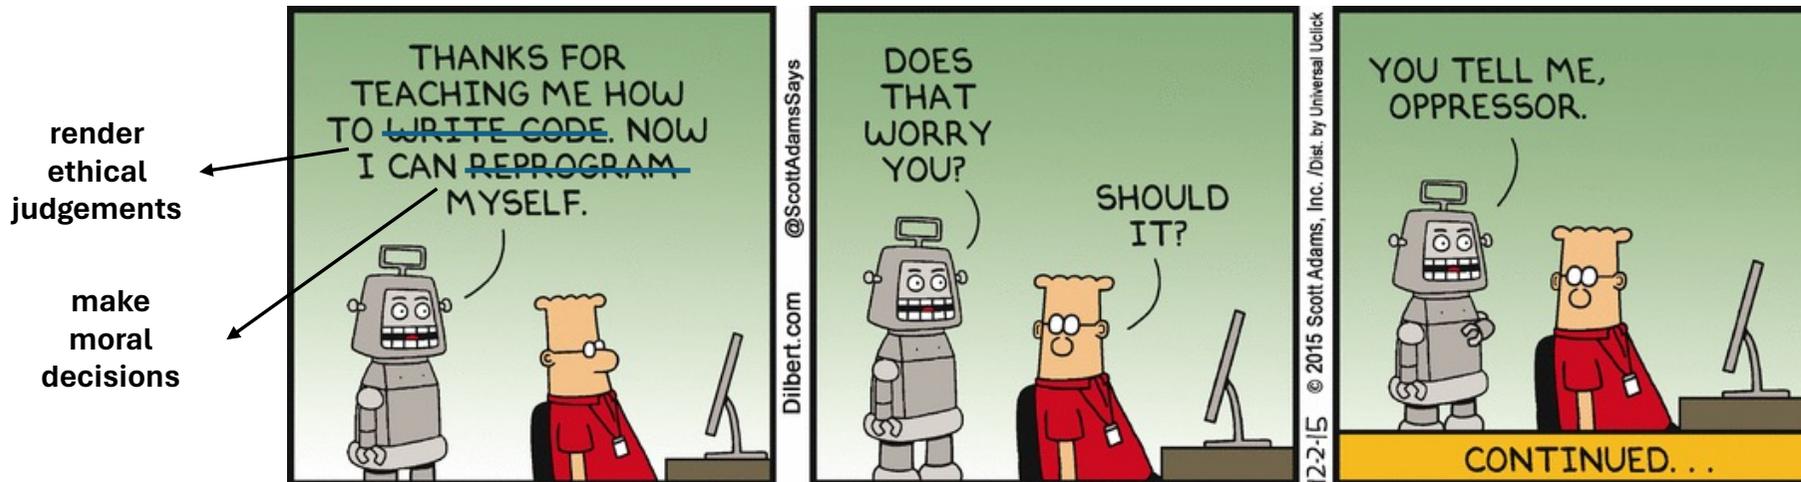
AREXIS

8

AI-enabled Research Ethics eXpert and Intelligence System



Should RECs defer to GenAI in Rendering Ethical Judgements?



Should RECs defer to GenAI in Rendering Ethical Judgements?

- P1** RECs ought to make appropriate ethical judgments.
- P2** The making of such appropriate ethical judgements requires **moral expertise**.
- P3** RECs currently lack the moral expertise to make appropriate ethical judgements.
- P4** There are good reasons to believe that certain types of GenAI will **have greater moral expertise**, and therefore will be better at making appropriate ethical judgements than most RECs.
- C** Therefore, RECs should use GenAI with moral expertise to assist them in making appropriate ethical judgements.

Should RECs defer to GenAI in Rendering Ethical Judgements?

P1 RECs ought to make appropriate ethical judgments.

- RECs have a **legal remit to render ethical judgments** in their oversight of HSR; their reviews (i.e., analyses & deliberations) inform and shape their **ethical judgements**, and in turn, **directly translate into decisional outcomes** — i.e., approve, require modification, or disapprove
- Obtaining/arriving at **appropriate moral judgements** is inescapable — or more specifically, **a condition that should be fulfilled** – in the REC’s rendering of ethical judgements. [RECs can make ‘inappropriate’ EJs]

P2 The making of such appropriate ethical judgements requires **moral expertise**.

P3 RECs currently lack the moral expertise to make appropriate ethical judgements.

P4 There are good reasons to believe that certain types of GenAI will **have greater moral expertise**, and therefore will be better at making appropriate ethical judgements than most RECs.

C Therefore, RECs should use GenAI with moral expertise to assist them in making appropriate ethical judgements.

Should RECs defer to GenAI in Rendering Ethical Judgements?

P1 RECs ought to make appropriate ethical judgments.

P2 The making of such appropriate ethical judgments requires **moral expertise**.

P3 RECs currently lack the moral expertise to make appropriate ethical judgments.

P4 There are good reasons to believe that certain types of GenAI will *have greater moral expertise*, and therefore will be better at making appropriate ethical judgments than most RECs.

C Therefore, RECs should use GenAI with moral expertise to assist them in making appropriate ethical judgments.

- RECs have a **legal remit to render ethical judgments** in their oversight of HSR; their reviews (i.e., analyses & deliberations) inform and shape their **ethical judgements**, and in turn, **directly translate into decisional outcomes** — i.e., approve, require modification, or disapprove
- Obtaining/arriving at **appropriate moral judgements** is inescapable — or more specifically, **a condition that should be fulfilled** – in the REC’s rendering of ethical judgements. [RECs can make ‘inappropriate’ EJs]

- Moral expertise: a **specialised set of knowledge and skills** (normative analysis, interpretation, and reasoning) that is necessary — to **appropriately & accurately identify, engage, and weigh different** (at times competing) **moral considerations** from the relevant facts and features of the HSR — to **arrive at appropriate moral judgement(s)**.
- Scenarios where moral intuitions and/or lived experiences are insufficient: “lethal research” where voluntary euthanasia is legalised; first-in-country human challenge trial; vaping behaviours where it’s illegal

Should RECs defer to GenAI in Rendering Ethical Judgements?

P1 RECs ought to make appropriate ethical judgments.

- RECs have a **legal remit to render ethical judgments** in their oversight of HSR; their reviews (i.e., analyses & deliberations) inform and shape their **ethical judgements**, and in turn, **directly translate into decisional outcomes** — i.e., approve, require modification, or disapprove
- Obtaining/arriving at **appropriate moral judgements** is inescapable — or more specifically, **a condition that should be fulfilled** – in the REC’s rendering of ethical judgements. [RECs can make ‘inappropriate’ EJs]

P2 The making of such appropriate ethical judgements requires **moral expertise**.

- Moral expertise: a **specialised set of knowledge and skills** (normative analysis, interpretation, and reasoning) that is necessary — to **appropriately & accurately identify, engage, and weigh different** (at times competing) **moral considerations** from the relevant facts and features of the HSR — to **arrive at appropriate moral judgement(s)**.
- Scenarios where moral intuitions and/or lived experiences are insufficient: “lethal research” where voluntary euthanasia is legalised; first-in-country human challenge trial; vaping behaviours where it’s illegal

P3 RECs currently lack the moral expertise to make appropriate ethical judgements.

- Many, if not most, present-day RECs **do not have members with moral expertise** — i.e., appoint scientists, lawyers, religious leaders, and laymen, **but not (research) ethicists or moral philosophers**.
- Given P2, these RECs that lack moral expertise **inherently fail to arrive at appropriate moral judgements** (from their ethical deliberations). [Current RECs are inherently at risk of arriving at ‘inappropriate’ EJs]

P4 There are good reasons to believe that certain types of GenAI will **have greater moral expertise**, and therefore will be better at making appropriate ethical judgements than most RECs.

C Therefore, RECs should use GenAI with moral expertise to assist them in making appropriate ethical judgements.

Should RECs defer to GenAI in Rendering Ethical Judgements?

P1 RECs ought to make appropriate ethical judgments.

- RECs have a **legal remit to render ethical judgments** in their oversight of HSR; their reviews (i.e., analyses & deliberations) inform and shape their **ethical judgements**, and in turn, **directly translate into decisional outcomes** — i.e., approve, require modification, or disapprove
- Obtaining/arriving at **appropriate moral judgements** is inescapable — or more specifically, **a condition that should be fulfilled** – in the REC’s rendering of ethical judgements. [RECs can make ‘inappropriate’ EJs]

P2 The making of such appropriate ethical judgements requires **moral expertise**.

- Moral expertise: a **specialised set of knowledge and skills** (normative analysis, interpretation, and reasoning) that is necessary — to **appropriately & accurately identify, engage, and weigh different** (at times competing) **moral considerations** from the relevant facts and features of the HSR — to **arrive at appropriate moral judgement(s)**.
- Scenarios where moral intuitions and/or lived experiences are insufficient: “lethal research” where voluntary euthanasia is legalised; first-in-country human challenge trial; vaping behaviours where it’s illegal

P3 RECs currently lack the moral expertise to make appropriate ethical judgements.

- Many, if not most, present-day RECs **do not have members with moral expertise** — i.e., appoint scientists, lawyers, religious leaders, and laymen, **but not (research) ethicists or moral philosophers**.
- Given P2, these RECs that lack moral expertise **inherently fail to arrive at appropriate moral judgements** (from their ethical deliberations). [Current RECs are inherently at risk of arriving at ‘inappropriate’ EJs]

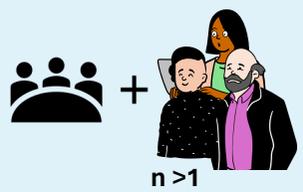
P4 There are good reasons to believe that certain types of GenAI will **have greater moral expertise**, and therefore will be better at making appropriate ethical judgements than most RECs.

- This GenAI has to be: well-constructed/reliable; trained on a comprehensive **ethics knowledge base (KB)**; and equipped with reasoning **capabilities**. [Mapping: KB to ‘knowledge’, and capabilities to ‘skills’.]
- Arguably, the KB and capabilities of this GenAI, and consequently its outputs, would be **comparable to those of a human person with moral expertise** — e.g., DigiDan, PeterSinger.AI, Savulescu’s personalised moral guru, Ethical avatars...
- Given P3, **relative to this GenAI, RECs that lack moral expertise fail at obtaining appropriate moral judgements**. The depth and quality of moral judgments — including its analysis and reasoning — generated by this **GenAI would be better than the RECs that lack moral expertise**. [Converse may not be true — RECs constituted with exemplars of moral expertise such as Kant, Aristotle, Plato...]

C Therefore, RECs should use GenAI with moral expertise to assist them in making appropriate ethical judgements.

Level of ethical expertise

Plausibly permissible & warranted?

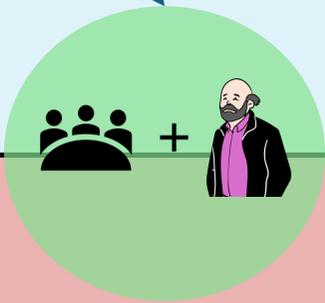


Icon of a group of people +

- DigiDan
- PeterSinger.AI
- Savulescu AUTOGEN

- Aristotle Kant
- Confucious Bentham
- Socrates Nietzsche etc...

Better than



Minimum Level of ethical expertise required

Status Quo (present-day RECs)



- REC with **no** ethicists or philosophers
- Ethics-specific GenAI/LLM
- Ethicists or Philosophers

Quality of ethical expertise

Possible Objections on Deferring to GenAI



"I'm right there in the room, and no one even acknowledges me."

The in the Room

- Why is the REC relying/deferring to GenAI for ethical judgements?
- Why are there no — or so few — ethicists/philosophers on RECs?
- Shouldn't the REC recruit (more)?



The Chink in the Armour

Claim is only as good as the **accuracy** and **reliability** of the GenAI's outputs — making appropriate ethical judgements.



Cognitive Offloading/Deskilling

Reliance on GenAI will hinder (or erode) development of crucial skills in **critical thinking** and **independent judgment**.

A Valid Objection

RECs should raise or expand their moral expertise by appointing ethicists or philosophers to have a seat at the table. Claim to defer to GenAI serves only as a stopgap for the status quo (see deskilling).

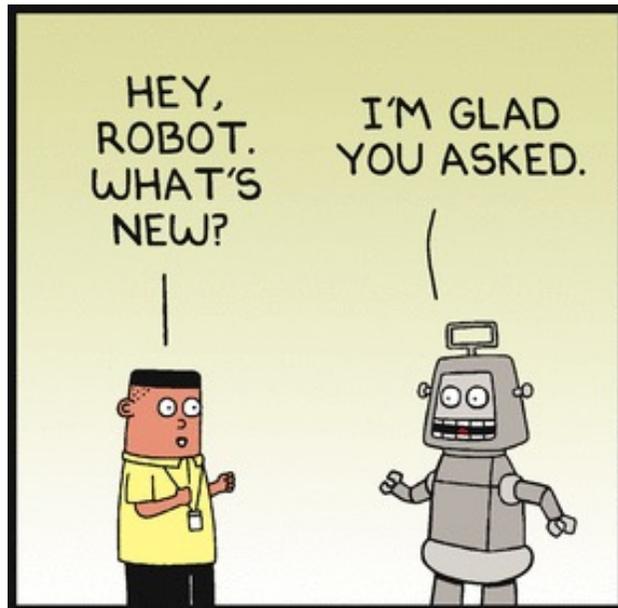
But humans too are NOT infallible!

When tired and/or under pressure — don't recall facts well, sufficiently, or accurately (applicable to ethicists too). Clinicians, who form the majority of REC compositions, must juggle clinical & teaching duties with REC review, and **lack substantial ethical expertise**.

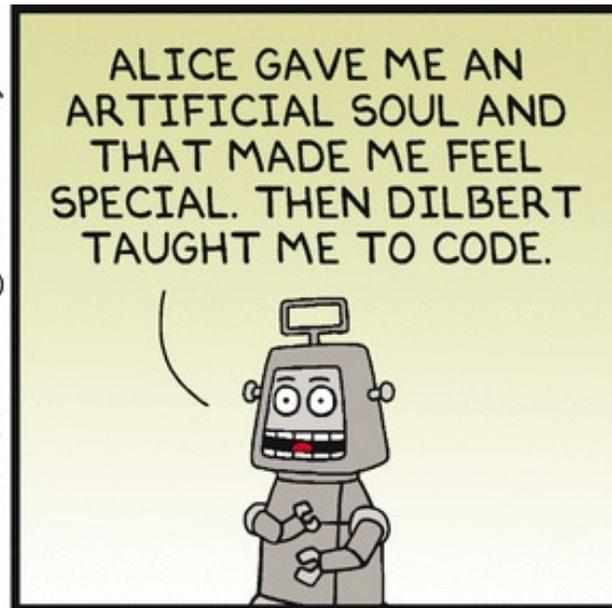
Unfavourable Status Quo

RECs lacking moral expertise is unfavourable and detrimental. Until moral expertise of RECs have been raised (i.e., appoint ethicists), RECs should defer to GenAI as a stopgap in the meantime.

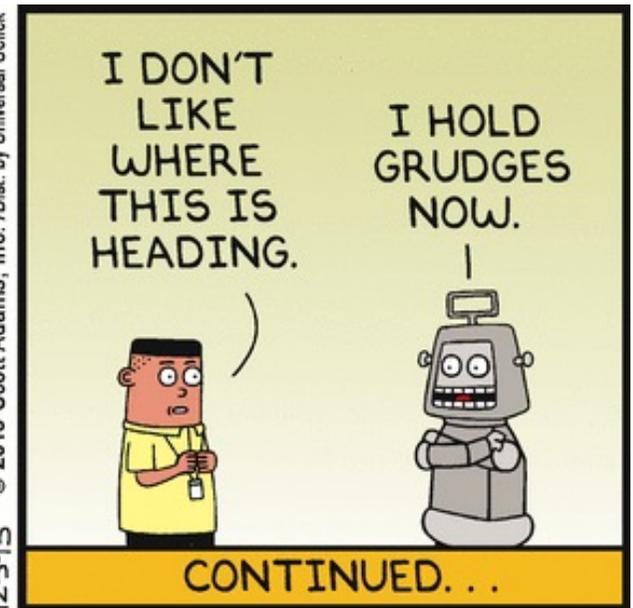
THANK YOU



Dilbert.com @ScottAdamsSays



12-3-15 © 2015 Scott Adams, Inc. /Dist. by Universal Uclick



Acknowledgement
Prof. Angus Dawson