

Ethical Aspects of AI Integration in Higher Education

A Roundtable Summary

Prof. Eli Salzberger
Dr. Avigail Ferdman
Dr. Shira Ahissar
Oshrat Katz Shacham



Ethics | 08/26



Ethical Aspects of AI Integration in Higher Education

A Roundtable Summary

Prof. Eli Salzberger

Dr. Avigail Ferdman

Dr. Shira Ahissar

Oshrat Katz Shacham

August, 2026

No part of this publication may be reproduced without prior written permission from the Samuel Neaman Institute except for the purpose of quoting short passages in review articles and similar publications, with an explicit indication of the source.

The opinions and conclusions expressed in this publication are those of the author(s) and do not necessarily reflect the opinion of the Samuel Neaman Institute.

www.neaman.org.il

Table of Contents

Rethinking Norms	3
Big Ideas to Consider	4
Necessary Epistemic Shifts.....	6
Appendix A: Roundtable Agenda and List of Participants.....	7

Background and Rationale

This document is based on a roundtable convened by the Ethics Research Group at the Samuel Neaman Institute to explore the ethical implications of integrating artificial intelligence into higher education. Held on 17 May 2026 at the Samuel Neaman Institute, the roundtable brought together scholars and experts in law, philosophy, technology, education, and higher education policy to discuss the ethical, epistemic, and institutional challenges raised by the growing use of AI in higher education. Following an opening plenary discussion, participants divided into two working groups: one focusing on research and the other on teaching. The full agenda and list of participants are provided in Appendix A.

Rather than providing a chronological summary of the discussion, this document develops the key questions and insights that emerged into a broader analytical framework. It examines the normative, epistemic, and institutional implications of AI in higher education, with particular attention to the conditions under which academic knowledge is produced, evaluated, attributed, recognized, and trusted.

The document is organized around three interconnected themes. The first examines the need to rethink academic norms in response to the growing use of AI in research, teaching, and learning. The second explores broader institutional questions concerning the role and purpose of universities in an AI-mediated knowledge environment. The third considers the epistemic shifts that may be needed to preserve meaningful expertise, intellectual development, and trust in academic work.

Rethinking Norms

The current normative order has collapsed, not just eroded. Several participants converged on the view that what is happening is not incremental disruption but a structural breakdown of the mechanisms by which scientific communities establish, enforce, and reproduce norms. The key argument, made most sharply by Niva Elkin-Koren, is that scientific validity was never just about the content of claims but about the processes and institutions through which communities agreed on what counts as valid knowledge — peer review, journals, grants, academic advancement. These mechanisms are now simultaneously flooded and circumvented.

The normative vacuum is real and not self-resolving. Arnon Keren argued that norms around AI in research will not simply emerge organically, as they did with earlier disruptions (he used the printing press as analogy). Organic norm development requires time and a period of relative stability; neither condition currently holds. The pace of technological change outstrips any plausible rate of normative adaptation. This is not just a practical

problem — it means there is no shared baseline from which to adjudicate what is permitted, appropriate, or praiseworthy.

The diversity of potentially legitimate norm regimes complicates any universal standard. Keren also raised an important pluralist point: multiple incompatible normative frameworks could each be internally coherent and legitimate. An institution that bans all AI use and an institution that treats AI-generated text as authorially unproblematic could both be principled. The current problem is not that we have the wrong norms — it is that we have no settled norms and therefore no shared basis for evaluation.

Norms about attribution and credit need to be rebuilt from scratch. Elkin-Koren suggested that the very concept of authorial credit — the central organizing principle of academic labor — may be obsolete in its current form, not merely in need of adjustment. Eldar Haber's account of writing a full 70-page law review article in nine hours using AI agents made this viscerally concrete: if this is replicable, the system of credit that assigns value to intellectual labor no longer maps onto actual epistemic contribution.

Transparency declarations are a partial but real response. One concrete normative proposal that gained some traction: requiring researchers to attach full interaction logs with AI systems to submitted work, so that evaluators can assess the nature and degree of AI involvement. This was proposed not as a solution but as a minimal condition for honesty — a way of making the human contribution legible rather than allowing it to be obscured.

Big Ideas to Consider

The signal-to-noise problem at scale is existential for academic infrastructure. Elkin-Koren's point about asymmetric costs — it is now cheap to produce content and very expensive to evaluate it — describes a structural failure that cannot be corrected by working harder within existing institutions. The conference example she cited (over 800 papers submitted by a single lab to a major venue) suggests that peer review as currently practiced cannot survive the volume problem, regardless of quality.

The institutional question: what are universities for? Erez Brown pushed hard on this: the ethics of AI in research cannot be resolved without first settling what the purpose of the research university is. Universities have systematically avoided this question. But without an answer, the AI discussion collapses into technical management of symptoms. The mission question precedes — and should generate — the normative framework.

The cognitive division of labour has a new, powerful entrant. AI is not just a tool but a new participant in the cognitive division of labour. Every normative question about AI in research is partly a question about what role this participant should play, under what conditions, and with what accountability. This framing is useful because it avoids both the naïve "AI is just a tool" view and the alarmist "AI will replace researchers" view.

The automation of scientific judgment, compared to scientific production, is the deeper threat. Haber's account focused not only on writing but on evaluation: if AI can serve as reviewer, grant assessor, and judge of novelty, then the entire epistemic authority structure of science (who decides what is new, what is valid, what matters) is at risk of being absorbed into systems controlled by corporations. Elkin-Koren made this point sharply: when AI is used both to generate and to evaluate, and both systems are trained on the same prior corpus, the capacity to genuinely identify novelty may be systematically compromised.

Automation of research tasks risks prioritizing efficiency over the goal of genuine inquiry. Efficiency does not leave room for making mistakes and learning from errors, and the curiosity that accompanies the process of learning from mistakes.

The distinction between intellectual formation and professional credentialing matters enormously. Shimon Marom drew a consistent and important line between *Bildung* — the formation of a researcher through genuine intellectual engagement — and the credentialing function of the university. His argument was that the AI crisis is actually an opportunity to recover the former from the bureaucratic distortions imposed by the latter, including the metric culture of publication counts, impact factors, and citation rates. In other words, to forgo rote teaching in favour of education. Students ought to be tested according to the way they integrate their skills, rather than spewing knowledge at an exam. Most participants implicitly agreed that the current incentive structure had already corrupted the research enterprise before AI arrived, following Ilana Ram's observation that AI is accelerating and making visible pre-existing deformations.

Effort as a Constitutive Principle in Research and Teaching. Yuval Dror emphasized that artificial intelligence, for the first time, performs integrative processes on behalf of both students and researchers, thereby effectively eliminating the need to make an effort. The problem lies in the fact that we are giving students and early-career researchers tools that allow them to bypass the stage of effort; yet without effort, there is no learning. This has direct implications for how we should think about teaching and the training of graduate researchers: we should adopt a principle according to which AI assistance supports learning when it makes us exert effort and avoid AI use when it replaces cognitive work by providing answers and bypassing the investment of effort.

Necessary Epistemic Shifts

From a scarcity to an abundance epistemology. Elkin-Koren made this explicit: we have always operated in an epistemology of scarcity — data, literature, analytical capacity were all limited resources that justified careful, slow, accretive knowledge production. We are now in an epistemology of abundance, but our evaluation criteria, institutional structures, and conceptions of intellectual merit still presuppose scarcity. The epistemic shift required is not simply about using new tools — it involves reconceiving what counts as an original contribution, what genuine expertise consists in, and what the function of academic knowledge production is in a world where generation is nearly free.

From academic integrity framing to epistemic trust framing. Avigail Ferdman proposed this reframing: the problem of AI in research is not primarily a problem of honesty (academic integrity) but of epistemic trust — the conditions under which a scientific community can rely on the outputs and judgments of its members. This is a different and deeper problem. Academic integrity violations are discrete rule-breaking events; erosion of epistemic trust is a systemic condition that cannot be resolved by detection and punishment.

The validation problem cannot be solved by "human in the loop." Elkin-Koren was explicit that the standard regulatory response — insert a human reviewer at key stages — is not adequate either scientifically or institutionally. When the human reviewer is evaluating outputs of systems trained on historical knowledge, the reviewers themselves are subject to path dependence and cannot easily identify failures that are structurally invisible within the shared paradigm. This is an epistemic problem, not merely a procedural one.

Expertise itself needs reconceptualization. Several participants touched on this from different angles. What constitutes expertise when much of what we called expert performance — synthesis, pattern recognition across large literatures, formal argument construction — can be performed by AI at high quality? The emerging answer, implicit in several contributions, is that expertise in the AI era will center on: (a) knowing what questions are worth asking; (b) having the judgment to evaluate machine outputs critically; (c) taking epistemic responsibility for claims in the full sense — not just knowing that an output is formally correct but understanding why it is true, what its limits are, and where it could fail. This is a shift from expertise as performance to expertise as accountability.

Scientific creativity may need to move toward the logic of art. Irad Yavne suggested that if AI can eventually handle proof and derivation in mathematics better than any human, what remains for the mathematician is something closer to artistic creativity — identifying what is worth investigating, what constitutes an interesting or beautiful problem, what conceptual moves are genuinely novel. This is not a diminished role, but it requires a different self-understanding on the part of researchers, and different criteria for evaluating and rewarding research work.

Appendix A: Roundtable Agenda and List of Participants



Invitation to the Roundtable: Ethical Aspects of AI Integration in Higher Education

How is artificial intelligence reshaping the epistemic environment of the university, the conditions under which knowledge is produced and disseminated, and the responsibilities of academic institutions, researchers, faculty members, and students in the use of AI for research, teaching, and learning?

Can AI tools be integrated in ways that enhance learning and research without compromising academic integrity, human judgment, literacy, or the cognitive capacities that higher education is intended to cultivate?

And is there a need for a higher education policy to address these issues?

**The Ethics Research Group at the Samuel Neaman
Institute is pleased to invite you to join us in
.discussing these important questions**

Sunday, 17 May 2026

09:00–13:50

Samuel Neaman Institute for National Policy Research

Panel Members (in alphabetical order)

Prof. Niva Elkin-Koren, Tel Aviv University

Prof. Eldad Haber, University of Haifa

Prof. Arnon Keren, University of Haifa

Prof. Shimon Marom, Technion and the Samuel Neaman Institute

Dr. Ilana Ram, Technion and the Samuel Neaman Institute

Session	Activity	Time
	Registration	09:00–09:30
Opening	Welcome and Introduction to the Topic	09:30–09:50
Session 1	Opening Plenary: Five Points of View	09:50–11:05
	Coffee Break	11:05–11:25
Session 2	Parallel Working Groups: Teaching and Research	11:25–12:40
	Reports from the Working Groups	12:40–12:50
Session 3	Plenary: Integrating Key Insights	12:50–13:40
Closing	Concluding Remarks and Next Steps	13:40–13:50

Full List of Participants (in alphabetical order by surname)

Dr. Shira Akhisar, University of Haifa; Samuel Neaman Institute

Prof. Rosa Azhari, Samuel Neaman Institute

Prof. Erez Braun, Faculty of Physics, Technion; Samuel Neaman Institute

Mr. Yoav Ben-Arie, Samuel Neaman Institute

Dr. Tzipi Buchnik, Samuel Neaman Institute

Dr. Yuval Dror, Sociologist

Prof. Niva Elkin-Koren, Faculty of Law, Tel Aviv University

Dr. Reuven Gal, Samuel Neaman Institute

Prof. Orit Hazzan, Faculty of Education in Science and Technology, Technion; Samuel Neaman Institute

Prof. Eldar Haber, Faculty of Law, University of Haifa

Prof. Arnon Keren, Department of Philosophy, Faculty of Humanities, University of Haifa

Prof. Shimon Marom, Samuel Neaman Institute

Dr. Nohad Ali, Samuel Neaman Institute

Dr. Anat Fligelman, Samuel Neaman Institute

Ms. Ronit Piso, Samuel Neaman Institute

Dr. Abigail Freadman, Department of Humanities and Arts, Technion; Samuel Neaman Institute

Dr. Ilana Ram, Faculty of Education in Science and Technology, Technion; Samuel Neaman Institute

Prof. Eyal Shimoni, Tel-Hai University; Samuel Neaman Institute

Prof. Eli Salzberger, Faculty of Law, University of Haifa; Samuel Neaman Institute

Prof. Irad Yavneh, Faculty of Computer Science, Technion; Samuel Neaman Institute



neaman.org.il