

Prediction, Purpose, and the Becoming of the Human Person

June, 3–5, 2026

A warm welcome

Dear Participants and Guests of the
ENNOIA Symposion 2026 on Prediction,
Purpose, and the Becoming of the
Human Person,

This decade has already brought humanity two great upheavals: the pandemic and the release of the large language models. That one followed so closely upon the other was, I think, no accident of timing. Covid turned our reliance on the internet into a dependence. Withdrawn into our four walls, we came to live our relationships through the screen. Into this condition, ChatGPT entered and brought with it the experience that the medium itself could enter into dialogue with us.

Until now, we have known ourselves in answerable relations of three kinds: to the world, to the other human being, and to first principles. With artificial intelligence, a fourth relation seems to enter, for with AI, the work of our own hands has begun to address us as we address one another. That machine, however, is not a subject; its answers patterned upon an impersonal It pre-formed from data and ungrounded in lived relationships. The simulation succeeds so well that we are tempted into relations bearing the form of an I-Thou, while what Martin Buber called the real I-Thou relations are gradually weakened on all three levels.

Until now, our knowing has also always been our becoming. Through this relation between knowing and becoming, we found our way ever anew to ourselves: we became creative; we became scientific; we became fully human subjects. Our ability to make artificial intelligence is itself part of our long human becoming through knowing. Yet what remains of us when an algorithm teaches us, speaks all our languages, discovers for us, decides for us through its thinking, and, perhaps most importantly, increasingly mediates our answerable relations for us? Will we face the consequences of outsourced thinking leading to outsourced answering and to outsourced responsibility? Or can we live the unfolding of our own answerable selves anew amidst the surge of AI?

At this threshold, the ENNOIA Symposion 2026 opens the four-year horizon Being Human in the Age of AI under which I hope we may work toward what I have called relational rationality: the ordering of the relations through which we become human in knowing and doing what is good for us. The claim of the Symposion is simple. By reordering ourselves to these relations, we can also give artificial intelligence, made in the image of our own discursive thinking, its rightful place, so that our being human may continue to bring forth an I able to answer a Thou.

The three days of the ENNOIA Symposion 2026 are conceived as a movement from the telos of learning to the formation of the learner, to what the human person may become. Researchers from across the sciences and humanities, including computer scientists, neuroscientists, philosophers, theologians, scholars of education and policy makers, roboticists, artists, mathematicians, gather in one room to address the anthropological weight of the questions that artificial intelligence now raises. For each discipline must give its account so as to place it within the larger order of being human.



To draw in an even greater community of like-minded researchers, each evening, the panels of the day shall gather again for shared writing sessions to draw the day's reflections into a common Charter of Research Questions. From these conversations we hope to develop a Perspectives Paper for Nature. George Caputa, Senior Editor at the journal, will accompany the Symposion throughout.

That this Symposion takes place at all is owed to several funding bodies. Before all others, my heartfelt thanks go to the kata agorein Foundation. Their generous support has made the undertaking possible from its first moment to its present form; and I thank the Foundation's Executive Director as well as the Board of Trustees personally for their untiring support and intelligent advice. My thanks also go to the Ruhr-Universität Bochum, the Lamarr Institute for Machine Learning and Artificial Intelligence, and the Angelicum in Rome for further contributions; to the Technical University of Dortmund as institutional partner; and to our two venues, the Harnack House Berlin and the Japanese-German Center Berlin, for their generous hospitality.

It remains for me to welcome you most warmly, and to thank you in advance for the seriousness and openness each of you brings to these days. May the artistic and culinary practices that thread through our days remind us that thinking has always taken place in the company of beauty and a shared life. I look forward to the forms of understanding and the friendships that may come into being among us.

Katja Krause



DAY 1

What Is Learning For

Chairs of the day

DAY 1



THE STATE OF AI
Martin Stratmann



SESSION A
Katja Krause



SESSION B
Onur Güntürk



SESSION C
Tobias Hauck

PROGRAM

9:00 AM	Opening of the Day
9:30 AM	The state of AI
10:00 AM	Session A — Talks
11:10 AM	Coffee Break
11:25 AM	Session A — Discussion
12:25 PM	Long Lunch Break
2:00 PM	Session B — Talks and Discussion
3:40 PM	Coffee Break
4:00 PM	Session C — Talks and Discussion
6:00 PM	Long Break
6:45 PM	Writing Session: Charter of Research Questions (chairs of the day)
8:00 PM	Joint Dinner at Harnack House
9:00 PM	Evening Concert: Adamello Quartett
10:00 PM	End

Computing, fast and slow:

Post-Training Optimization and the Next Generation of AI

The fast-accelerating progress of artificial intelligence owes its success to combining several different types of computational approaches.

Primarily, integrating models that were pre-trained using web-scale data with algorithms and tools that can focus on a specific instance. I will talk about this new group of approaches, which focuses on rapid optimization for a specific prediction task using a pre-trained model.

Similar to “System 2” in the “Thinking, Fast and Slow” model of cognitive processing, these approaches are not rapid inference from a pre-trained model, because they involve non-trivial optimization during prediction, beyond what the model was prepared for during training.

These approaches are also not training problems, as they account for the specific instance that needs to be handled. These computationally intensive prediction processes pose new challenges in the field of computational learning and open up possibilities for new user experiences and diverse use cases.

In this lecture, I will describe recent work on post-training optimization in the context of text-to-image generative models, and discuss how these approaches are being applied to make new discoveries in science.



Gal Chechik

Professor of Computer Science at Bar-Ilan University
Senior Director of AI Research at NVIDIA, Israel

Super- intelligence without AGI

Challenges and Opportunities for Humanity

In one of the two central dimensions of intelligence—crystallized intelligence, the body of accumulated knowledge an individual possesses—generative AI already far exceeds human capability.

By training neural AI systems on vast quantities of language data through behavior-based learning alone, researchers have enabled them to acquire a kind of superhuman knowledge—knowledge that even 10,000 carefully selected experts across many different disciplines could not collectively master in comparable breadth. The next frontier is the modeling of fluid intelligence, which accounts for the larger share of what we typically measure as general intelligence.

Until now, neural models have lacked an integrated memory. For modelling true intelligence, a rich, layered memory architecture is needed: short-term and long-term memory, along with a working memory capable of actively manipulating new information and existing knowledge while continuously revising and extending both. In humans, the biological limitations of working memory are among the greatest constraints on intelligence.



Hans Uszkoreit

Scientific Director at the German Research
Center for Artificial Intelligence (DFKI)
DFKI Language Technology Lab

In AI systems, however, such limits can be overcome through more expansive design. Leading AI laboratories are now developing the first neural LLMs that augment the transformer architecture of generative AI with such powerful memory systems, and the early results appear highly promising.

Yet even with these new capacities for complex, multi-step inference and autonomous reasoning, such models still fall short of human cognition in many important respects. Superhuman knowledge and superhuman reasoning do not by themselves compensate for what is missing as a result of not sharing the human condition and human social development. These missing elements include: consciousness, embodiment—sensory experience and spatial awareness—as well as instincts and drives, emotions, common sense, and interpersonal and social abilities.

The emergence of super-intelligent AI raises important questions:

- Whether, and in what form, a super-intelligence might develop functional counterparts to aspects of human cognition that AI still lacks today —such as explicit knowledge, consciousness, and social intelligence?
- What are the essential elements and lasting strengths of human existence?
- How should we shape the interaction between human intelligence and artificial intelligence?
- How can—or must—we educate superintelligence so that it becomes both beneficial and safe?
- What steps must we take to effectively address the undeniable risks arising?
- And how can we use AI to drive a renewal of our society—its democracy, economy, public administration, and education system—which is already in urgent need of reform?

What Remains of the Telos of Education
— and Must We Now Reinvent It?



**Fr. Dominic
White OP**

OP Faculty of Theology,
University of Oxford; musician

Recital and verbal reflection

Frédéric Chopin,
Nocture in F minor, Op 55, No 1,
Dominic White,
Nocturne (Naming Infinity)



Oren Harman

Senior Research Fellow,
Van Leer Jerusalem Institute;
Professor of the History of
Science, Bar-Ilan University

**Story-Telling in
the Age of AI**

As machine learning invades more and more
spaces of human activity, something curious is
happening to our relationship to knowledge:

increasingly, we are loosening our ties to causal
thinking. And while we may become better at
explaining many things in and about our world,
we risk understanding it less and less. This
emerging new reality provides a special op-
portunity for re-thinking the humanities, one in
which story-telling plays a special role.

Plato on AI

In 1963, Eric Havelock, an anthropologically
minded classicist, wrote an influential book
entitled *Preface to Plato* which analyzed the
epistemological consequences of the slow (re)
introduction of writing in Greece after the de-
struction of Mycenaean civilization in the 12th
or 11th century.

According to Havelock, this re-appearance of
writing, combined with the adoption of the
Phenician alphabetic system, fostered the
emergence of critical thinking in a population
not submitted any longer to the centralized
institutions of the Mycenaean empire. Havelock's
views was an anti-Platonic move.

For Plato, in the *Phaedrus*, claims that writing
is cognitively detrimental to the transmission of
authentic knowledge: instead of helping mem-
ory, it destroys it, and it is incapable to teach,
since a text cannot answer questions.

In my presentation, I shall complicate the picture
by drawing attention to the fact that for Plato,
there is also a political dimension to the ques-
tion of writing, and then reflect on the analogies
and disanalogies existing between the two tre-
mendous technological revolutions represented
by writing on the one hand, AI on the other
hand, and the responses that they trigger.



André Laks

Emeritus Professor of Philosophy,
Sorbonne University;
Professor,
Universidad Panamericana

What Remains of the Telos
of Education — and Must
We Now Reinvent It?

Cultivation, Fabrication, or Optimisation?

AI and the Future of
Education



Shannon Vallor

Baillie Gifford Professor in the Ethics of
Data and Artificial Intelligence, and
Director of the Center for Technomoral
Futures, University of Edinburgh

The transformation of the learner has always been a central aim of education.

However, the type of transformation sought has varied widely. Naturalistic or spiritual teleologies of the human person conceive of education as cultivation, facilitating the young person's intended developmental trajectory and the maturation of their latent capabilities. Modern secular education has instead conceived of education as fabrication: the making of people into capable citizens, workers, thinkers, leaders or innovators.

The tensions between these visions have now been superseded by the new ideal of education as optimisation. Here, the desired product of education is not 'the learned person' but rather a well-fitted, frictionless component that operates seamlessly within a maximally efficient system of production. This talk will explore the profound consequences of this shift, one now being rapidly accelerated by the gravity of AI as an engine of universal optimisation.

If education has a purpose, it has always been tied to the cultivation of human thinking.

Yet in the age of generative AI, this purpose is increasingly destabilised. While students recognise that AI should not replace their thinking, they often lack the means to regulate its use in practice. This keynote explores the emerging gap between awareness and action, arguing that education must shift from knowledge transmission to cognitive regulation. It presents the TACO framework as a way to operationalise human-AI partnership and restore agency at the centre of learning.



AI Didn't Break Education

We Just Never Taught
Students How to
Regulate Thinking

Cecilia Chan

(Virtual Participation)
Professor and Associate Dean
(Teaching and Learning),
The University of Hong Kong

Discerning Ends in an Age of Predictive Technologies

Predictive technologies like AI promise potential control over our future.

This talk draws on the Sermon on the Mount and other theological sources to argue that, without a grasp of the proper ends to seek, predictive tools will achieve neither control nor human flourishing. Faced with a multitude of risks, these tools may lead to anxiety, disappointment, and a failure to address novel possibilities.

These concerns suggest the need for a formation in a practical wisdom that can recognize ends and direct action to them.



Paul Scherz

Our Lady of Guadalupe Professor of Theology,
University of Notre Dame

What happens when individuals feel that AI understands them better than their partners, friends, or family?

The rise of AI companions signals a transformation in intimate relationships. They do not simply mediate emotional life but actively shape how emotions are experienced, relationships are formed, and how individuals come to understand themselves in relation to others.

This talk advances three claims. First, AI blurs the boundary between “real” and “simulated” relationships. Second, intimacy is increasingly structured by design: AI systems are engineered to sustain engagement, encourage attachment, and convert interaction into data. Third, these dynamics produce a new intimate order marked by asymmetries of control, in which platforms exert growing influence over emotional experience and relational life.

The central question, therefore, is not whether AI replaces human connection, but how it reconfigures the socio-economic and technological conditions under which intimacy—and forms of becoming—take shape.

AI and the Transformation of Intimacy: Emotions, Relationships and Becoming



X. Alvin Yang

Visiting Postdoctoral
Fellow MPIWG

We Can Calculate the Future —
but Which One Do We Want?

A Complex Systems Perspective on Agentic AI: Normative and Epistemic Implications

Agentic AI is increasingly hailed as a new horizon in the field, and it has become mainstream to attribute autonomy and agency to computational models and systems.

Yet these notions remain insufficiently articulated and poorly understood. In this presentation, I argue that (i) agentic AI is a discursively constructed institutional reality; (ii) its operation is exhaustively reducible to the generation of predictions unfolding as a deterministic chain reaction; and (iii) an interpretive account of AI grounded in complex systems theory entails significant normative and epistemic implications, including a potential shift in how legal responsibility and regulatory intervention are conceptualized.



Daria Kim

Senior Research Fellow,
Max Planck Institute for Innovation
and Competition, Munich

Recognizing the New Emperor

AI is a powerful technology; it's an even more powerful political-economic project. Like many powerful tools, it could be used to advance knowledge, empower vulnerable peoples, and address structural inequities.

And yet, that's not how we're seeing it be put to use. Instead, the inevitability that major tech companies, their investors, and their political enablers are putting forward is one of neo-colonialism, existential risk, and corporate hoarding of resources. Young people are surprisingly clear-eyed about the absurdity of the AI project.

To succinctly summarize the view my collaborators and I have heard from dozens: "Of course we're using it, but it's stupid, and it's yet another way in which adults are destroying our future." In this talk, I will interrogate the political-economic project of AI and young people's reaction to it. In support of youth, we have a responsibility to resist the inevitability rhetoric we are hearing and empower them to envision a wide range of alternate futures in pursuit of humanity in an AI-infused world.



danah michele boyd

Founder, Data & Society Research Institute;
Geri Gay
Professor of Communication at Cornell University

SESSION 1C:

Whom Does Science Serve — and Who Gets to Answer?



Ylva Söderfeldt

Professor of History of Science and Ideas;
Director of the Centre for Medical Humanities,
Uppsala University

Discomfort Zone: Synthetic Data and Truth in Biomedicine and the Humanities

Currently, more and more scholars in the humanities are embracing computational methods, creating and using large datasets, using various quantitative approaches to texts and images, and applying machine learning-based tools in their analysis.

Taking minds as dialogic products, defined through their intersubjective capacities, I approach AI as just one among many culturally and historically embedded attempts to “improve” minds. “Improvement” itself simultaneously assumes hierarchies of value (some minds are “better”), institutional authority (who decides what counts as improved), and teleologies (toward what ends).

Our experiment focuses on synthetic data: computer-generated texts, images, or other datasets that mimic authentic data. Synthetic data is increasingly used in biomedicine to supple-

ment real-world datasets, protect privacy, and to train, validate and verify machine learning methods without risking sensitive information. However, the idea of performing historical analyses based on computer-generated documents (or study art, literature, or culture through artificially generated paintings, poems, or human interactions) appears absurd.

In our experiment, we do exactly that. The purpose is to create a productive “discomfort zone” and explore what that discomfort can tell us about the use of synthetic data in biomedicine, and by extension about data-driven approaches and AI in both the humanities and life sciences.



Jamie Cohen-Cole

Associate Professor of
American Studies, George
Washington University

Mechanisms for Improving the Mind: Science, Education, and AI.

As digital methodologies are becoming mainstream in the humanities, this means also that the methodological gap between the humanities and biomedicine and science is narrowing. In this talk, I will present ongoing experimental work that aims to use this convergence as an opportunity for a fruitful bidirectional critique that can benefit both fields.

By examining mechanisms for mind improvement, we can trace and evaluate how hierarchies have been established, contested, and transformed, before and since the invention of AI. Such evaluation treats practices, institutions, and materials on an even plane - studying each as culturally and historically specific “mechanisms” for mind improvement.

I focus on two specific social mechanisms of dialogic mind improvement: education and science. Science and education alike have constituted specific technologies of trust that enable the circulation of ideas, the facts taken to support them, and the minds they are supposed to be improving. Most recently, new forms of AI, especially large language models, are reshaping existing technologies of truth, trust, facts, and minds that these technologies were once supporting.

Education in the Age of AI

Formal education in the industrial age has shown us how we produce second-rate robots; people that simply reproduce what they have learnt. In the age of AI we have to find out what it means to be human. The perhaps greatest challenge is not AI becoming human, but AI infantilizing us and bringing us to give up on basic human capabilities, as we trade our autonomy against comfort.



Andreas Schleicher

(Virtual Participation)
Director for Education and Skills, OECD

The education systems of the future need to help humans to think independently and to contribute to society and the workplace with empathy, make ethical decisions in uncertainty, to build and shape complex social relations and create true novelty. Specifically we should be expecting of education that it empowers the next generations to create new social and economical models for a century of unprecedented challenges, that it fosters human growth through an appropriate mix of cognitive skills, care and creativity and that it helps young people find meaning and hope in times of uncertainty.

These are not just concepts or nice words, they are at the core of education, and they are the pillars upon which we build our societies. If these human capabilities cannot assertively be safeguarded with education, AI could undermine the foundations of our societies.

DAY 1
WEDNESDAY
4:00 PM

SESSION 1C:

Whom Does Science Serve —
and Who Gets to Answer?

Basic Research as Part of Human Becoming in the Age of AI



Martin Stratmann

President Emeritus, Max Planck Society,
Board of Trustees, kata agorein Foundation

Why do we conduct basic research today? Adolf von Harnack's 1909 memorandum founding the Kaiser Wilhelm Society — later the Max Planck Society — gave one answer: leadership in the natural sciences carries “an eminent national and political value,” with economic value following of itself. This is science as a means to ends set by others.

In my talk, I propose a different answer, from the working scientist's own standpoint. Curiosity-driven research is part of what it means to be and to become human: we are the only species that wishes to grasp the inner connections of what we perceive. Scientific understanding is thus the derivation of generalizable laws: Newton's gravitation, the laws of thermodynamics, the Schrödinger equation.

In contrast, AlphaFold now predicts protein folding where quantum-mechanical methods reach their limits. It knows statistical relations and precisely not quantum mechanics. This is why it predicts without understanding. But if science contracts to prediction alone, what becomes of our human drive to understand? Would we still do science once machines do it for us? I argue that AI belongs in the expanded toolkit of science done by humans; and institutional spaces of inquiry, as Harnack envisaged them elsewhere, must continue to be defended.

AI, Society, and Computing:

Leveraging Global
Data While Tackling
Ethical Challenges

AI agents powered by individuals' data are reshaping everyday interactions on social platforms and marketplaces. At a planetary scale, the aggregation of multi-modal data enables transformative applications such as poverty mapping, socioeconomic predictions, and disaster assessments.

However, these advancements also bring systemic risks, including biases, misinformation, and threats to democratic institutions. This talk explores how thoughtful design in database systems and AI can mitigate these challenges, foster sustainable development, and uphold human-centered values. It further advocates for the vigilant oversight of societal-scale AI applications to prevent dual-use and other ethical concerns and instead promote the benefit for humanity.

I will also discuss my journey through my research on misinformation, as well as my life as a data scientist, based on my experiences collaborating with world-class scientists at Meta, AT&T Research, and Microsoft, as well as NGOs such as the United Nations Pulse Lab and the World Customs Organization.



Meeyoung Cha

Director Max Planck Institute for
Security and Privacy, Bochum;
Professor, Korea Advanced Institute
for Science and Technology


DAY 1
WEDNESDAY
4:00 PM

DAY 1
WEDNESDAY
9:00 PM
HARNACK-HAUS

Adamello Quartett

Felix Mendelssohn
— String Quartet No. 1 Op. 12

Szymon Laks
— String Quartet No. 2



Introduction by
Frank Harders-Wuthenow



Clemens Linder, Nikolaus Kneser, Adele Bitter, Susanne Linder

DAY 2

How the Learner Comes to be



Chairs of the day

DAY 2



SESSION A
Evelyn Blacklock



SESSION B
Julia Reed



SESSION C
Dominic White OP

PROGRAM

THURSDAY
04.06.2026

- | | |
|----------|--|
| 9:00 AM | Opening of the Day |
| 9:15 AM | Session A — Talks |
| 10:30 AM | Coffee Break |
| 11:00 AM | Session A — Discussion |
| 12:00 AM | Long Lunch Break |
| 2:00 PM | Session B — Talks and Discussion |
| 3:40 PM | Coffee Break |
| 4:00 PM | Session C — Talks and Discussion |
| 5:40 PM | Announcements of Future Activities |
| 6:00 PM | Long Break |
| 6:45 PM | Writing Session: Charter of Research Questions (chairs of the day) |
| 8:00 PM | Joint Dinner at Harnack House |
| 9:00 PM | Leisure and Bar |

SESSION 2A:

How Does Knowledge Pass from One Mind to Another — and What Does the Medium Do to the Message?

Appropriate but Uncaused: The Creative Aspect of Language and the Return of Behaviorism



David Pöppel

Professor of Psychology and Neural Science,
New York University

For the last 70 years, since the cognitive revolution, a productive approach to the study of language has been to look at this human cognitive faculty as an object to be studied by the standard approaches of the sciences. Studying the linguistic system is, in that sense, no different from studying the vascular system, the visual system, the immune system, and so on. In the language sciences, five (types of) questions motivate the inquiry. What is it that you know when you know a language?

How is language learned, i.e., how do you go from the initial state to becoming an adult speaker? How is language processed in real time? How is language represented in the brain? How did language evolve?

Considerable progress has been made on many of these issues. One generalization that emerges: the human language system differs in principled ways from the truly remarkable artificial LLMs we now dominate our days, notably LLMs.

I will explain an intuitive result from human neuroscience and the study of AI systems that underpins this position. The language sciences seek to provide understanding and explanation; but much of the research of the moment is about prediction and control, essentially a return to behaviorism. The philosophical observation that language is typically “appropriate but uncaused” suggests that prediction is no mechanism for the creative aspect of human language use.



Limor Raviv

Research Group Leader, Max Planck
Institute for Psycholinguistics, Nijmegen

Lessons from empirical work on language evolution and change

Languages adapt to fit our cognitive biases, our communicative needs, and our ecological niches. Since human societies differ in their social and environmental landscapes, different languages emerge.

But what exactly are the social, environmental, and cognitive pressures that shape the evolution of language in our species, from its very early stages of emergence to the ongoing change we see today? What factors lead to the creations of complex linguistic structures, and which can explain the astonishing linguistic diversity we see worldwide?

These questions represent the range of topics covered by my research so far. In my lab, we use a range of novel behavioral paradigms and computational tools - from group communication games and Virtual Reality experiments, to simulations with deep learning neural networks and swarm robotics - to test core phenomena in cultural evolution, cross-cultural variation, and language adaptation.

DAY 2
THURSDAY
9:15 AM



**Manuel
Andreas
Dürr**

Artist

Reflections on Mnemosyne

Manuel A. Dürr shares thoughts on the process of creation and the philosophical, art-historical, and religious implications of art as a medium of cultural memory and cultural creation. The painting Mnemosyne features two kinds of surfaces: on the one hand, the fragments of an old culture saturated with memory and meaning, but shattered and in ruins. On the other hand, the effervescence of a new digital culture, which generates ever-new possibilities, but increasingly without a footing in the incarnated world of pigment and flesh. What is the pedagogy of oil paintings and their humanizing potential in the age of AI?

How Does Knowledge Pass from One Mind to Another — and What Does the Medium Do to the Message?



Yehuda Halper

Professor of Jewish Philosophy,
Bar-Ilan University

Learning in the Teacher's Absence: Creative Thought and the Ever-Presence of Agentic AI

Human to human learning is learning from someone, often an authority figure, such as a teacher, mentor, etc. Such learning occurs both when the authority is present and actively instructing, and also, separately, after the learner is separated from the instructor.

It is in the latter stage that a great deal of creative thought occurs, in interpreting the instructor or in coming up with new ideas. Does the ever-presence of Agentic AI leave room for this latter kind of learning? History of texts, philosophy, and even religious transmission may suggest that it precisely after the departure of the instructor that the learner learns most.

Different Media, Different Epistemologies: What Learning Analytics Cannot See

From a Learning Analytics perspective, knowledge transfer is rarely observed directly; instead, it is inferred from human activity logged in digital traces.

Therefore, the medium not only shapes the message – it shapes the very evidence we use to claim that knowledge has been transmitted. Clicks, posts, and generated texts become proxies for understanding, while other dimensions of cognition, meta-cognition, and affect remain mostly invisible. As a result, different media construct different epistemologies: They redefine what counts as knowing, learning, and communicating. Understanding these transformations is essential in an era where knowledge increasingly flows through data-driven systems.



Arnon HersHKovitz

Assistant Professor,
School of Education,
Tel Aviv University

How to Become an Authentic Agentic Learner in the AI Era

The rapid development of AI is transforming how knowledge is produced, shared, and learned, while redefining roles, responsibilities, and relationships in education.

Building on UNESCO's call for a "new social contract for education", this work proposes a new pact for the intelligent era, centered on cultivating Authentic, Agentic, and Active (3A) Learners. Authentic learners engage with real-world problems and create meaningful, shareable outcomes; agentic learners define purposes, make value judgments, and exercise control over learning with AI; active learners plan, monitor, and reflect through sustained practice.

A multi-layer framework places learners at the core, strengthens teacher–student interaction and feedback, promotes school–family co-construction of learning environments, and is enabled by policy provision and institutional support, grounded in shared social values. It aims to foster responsible, creative, and future-ready learners who can both adapt to and lead the intelligent era.



Ronghuai Huang

Professor and Executive Dean, Faculty of
Education, Beijing Normal University

Embodiment Shapes Formation —
but What Does the Attempt to
Replace It Reveal?

Know Yourself!



Onur Güntürkün

Professor of Biopsychology, Ruhr-Universität Bochum

The use of mirrors for self-observation has long been regarded as a cognitive hallmark, hailed as an experimental indicator of self-awareness or even consciousness in humans and great apes.

Since then, the mark-and-mirror test has been administered to many healthy subjects, patients, and non-human animals, including invertebrates. The resulting picture presents a puzzle composed of partly contradictory findings and widely divergent interpretations.

Several key insights are emerging:

First, mirror self-recognition has been demonstrated in a small number of non-human animals, without a phylogenetic pattern of evolutionary proximity to humans.

Second, in some species, self-recognition only occurs when the test is conducted in the appropriate ecological context. Cognitive situatedness thus determines whether this ability is expressed or not.

Third, in some species, the ability for mirror self-recognition depends on the animal observing in the mirror how its own face is being touched. Recognizing oneself in the mirror may therefore depend on learning to integrate sensory events involving one's own face. Fourth, some patients with brain lesions lose the ability to recognize themselves in the mirror without losing their sense of self.

Overall, mirror self-recognition is not a litmus test for human-like cognition. However, it demonstrates that the self-reflective recognition of one's own body in the mirror depends on a variety of mechanisms that are often based on learning. We have to learn what is part of us and what is not.

The Human Person and The Neural Architecture of Reasoning

What is the human being, and what kinds of mental capacities constitute our distinctive modes of knowing and acting?

Contemporary neuroscience generates increasingly precise data on human cognition, yet often within frameworks that fragment the person into a handful of functions or reduce thought to language, computation, or diffuse neural activation.

In this talk, I argue that new evidence from cognitive neuroscience presses us toward a richer and more unified account of the human person. I present data from three convergent research programs, employing precision neuroimaging, clinical lesion studies, and investigations of atypical brains: dissociations between language and logical reasoning in profound aphasia, the discovery of a fronto-parietal "Logic Network" recruited across abstract inference, and preserved, structured cognition in individuals with extreme cortical deformities.

Together, these findings challenge reductionist assumptions and align most naturally with a non-dualist, non-reductionist accounts of the human person as one living organism with distinct but integrated cognitive powers, realized in an embodied form that is highly structured but not mechanistic.



Hope Kean

Postdoctoral Fellow Department of
Brain and Cognitive Sciences, MIT

Embodiment Shapes Formation —
but What Does the Attempt to
Replace It Reveal?



Stephen Blackwood

President, Ralston College

Education at its best takes into consideration the whole human person

The ancients and medievals had very robust models for both thinking about the distinct faculties of the human—for example: sense, imagination, reason, and intellect—and how to speak to, or educate, each of these faculties. But whatever particular conception one adopts, effective education takes into account the manner and means by which the student—body and mind, heart and soul, however one understands it—are coherently moved, engaged, and actualized.

Avatars and a Society Without Constraints

In this talk, I will introduce the technologies of autonomous robots and avatars (teleoperated robots and CG characters) that I have developed, and discuss how the avatar technology will change the world.

It is important to enable people with various backgrounds and values, such as those who need to care for or raise children or the elderly, to participate in diverse activities according to their own lifestyles, and to realize a society in which people are free from the constraints of body, brain, space, and time. Avatar is the realization of such a society. I am engaged in research, development, and business with the aim of enabling anyone to freely and remotely control multiple avatars and participate in diverse work, educational, medical, and everyday social activities without having to go to the field.



Hiroshi Ishiguro

Distinguished Professor Osaka University;
Director Intelligent Robotics Laboratory



Hans Uszkoreit

(DISCUSSION ONLY)


What Is a Person Beyond Their Function — and Where Do We Look for the Answer?

Another turn of the screw? Will AI further narrow our concept of the person?

The oldest philosophical question is “How to live?”. It is really two questions: How should I live as an individual? How should we live together?

For a large segment of the history of Western Civilization, one of the major religions was taken to supply the answers. With the Enlightenment, however, Kant, von Humboldt, and John Stuart Mill renewed the inquiry begun by the Greeks. They all emphasized the importance of framing our own conceptions of ourselves and what is central to our lives. I shall try to elaborate this liberal conception, and to show that it demands a large number of important capacities.

Sporadically, throughout human history, various societies have attempted to make some of these capacities available to some subset of their citizens (typically a very small one). Occasionally, political leaders have declared that those opportunities should be offered to all. In our own time, however, the institutions, local and global, crafted in the past half century have made it less and less likely that such declarations will be honored.



Philip Kitcher

John Dewey Professor Emeritus of Philosophy,
Columbia University

The dominant version of capitalism in the contemporary world, combined with the rise of the internet, has already narrowed our concept of the person. In principle, AI might prove liberatory. Given the socioeconomic milieu in which life today is lived, it is, however, highly likely that the liberal conception of the person will become ever more of a utopian fantasy.



Howard Gardner

Research Professor of Cognition and
Education, Harvard Graduate School of Education

Humanity After the Anthropocene

(VIRTUAL PARTICIPATION)

Homo sapiens have long dominated the planet. That era is rapidly drawing to a close. I will consider educational, social, and ethical consequences of this likely ‘sea change’ and propose an agenda for the emerging era.



Tobias Rees

Founder and CEO of limn.ai

Cosmos. Cognition. Self. Exposures

What does AI do to the premise of this Session – the premise that humans are ‘more’ than the (economic) function they fulfill?

I will address this question by tracing the various ways this ‘more’ has been articulated; by examining the conceptual assumptions those articulations made; and by asking what AI does to these assumptions. Call it exposure as method — designed to bring into view what it could mean to be human (to be a humanist) in the age of AI.

SESSION 2C:

What Is a Person Beyond Their Function — and Where Do We Look for the Answer?

Two Kinds of Ego Death: On Personhood in the Age of AI



Mario Verdicchio

Associate Professor in Logic and Philosophy of Science, University of Bergamo, Italy

This contribution distinguishes between two senses of “ego death.” The first, rooted in Eastern mystical traditions and echoed in reports from psychedelic experience, names a temporary dissolution of the self that can foster a less self-centered, more relational orientation toward the world.

A second, critical sense is introduced: a “functional ego death” emerging within contemporary AI discourse, where human beings are increasingly framed in terms of computation, optimization, and function. This shift carries two major risks. First, it relegates subjective experience to a mere epiphenomenon, obscuring the role of embodiment and first-person perspective in shaping meaning and action. Second, it promotes a form of technosolutionism that privileges problems amenable to algorithmic treatment, sidelining those that resist formalization yet remain central to human life.

This tendency is reflected in the growing influence of technocratic reasoning in political contexts. Against this backdrop, the ego is defended. It is not a problem to eliminate, but as a locus of value, conflict, and meaning that exceeds any functional account of the human.



Yvonne von Dohna Schlobitten

Extraordinary Professor, Faculty of History and Cultural Heritage of the Church, Gregoriana

Contemplating AI as an Artwork: From *theoria* to Mental Contemplation

Can AI be contemplated? How do things shape the human person? The existential difference between “becoming human” and “proving oneself to be human” lies in the way we encounter things. Things, including AI, can be approached instrumentally as objects within individual disciplines — didactically, therapeutically, or otherwise — and thus become part of education through the application of universally valid ethical, moral, or sapiential values. Yet if we seek to understand the role of AI more deeply, we must distinguish between sense and purpose, and between being and object, and ask what kind of promise or horizon AI opens for us.

To do so requires establishing a personal relationship within a triangle of artist, work, and beholder, a structure in which each pole answers and is answered by the others. The “unavailability” (*Unverfügbarkeit*), the “not yet known,” plays a crucial role in this encounter. “The work of art does not reproduce the visible, but makes visible,” writes Klee. This is *theoria* in its older sense: a seeing that receives.

In encountering AI, can we take it seriously as a space of credible engagement and become aware that it may disclose its own sense? This question goes beyond a secular phenomenology and touches on the possibility of encountering something numinous (*mysterium tremendum et fascinans*), or what Franz Marc describes, in relation to art, as a mystical “Me and God.” In this way, even AI could become not merely a tool, but an educator or a way-shower.

This formative dimension of encounter is precisely what Romano Guardini, who held the chair for Catholic *Weltanschauung* in Berlin, names *Selbst-Bildung*: the idea that encountering the essence (*Wesen*) of things must occur within the individual. This “deep knowing” (*Erkenntnis*), through which we recognize ourselves as human beings, can be accompanied, but not taught in a traditional way. Such mental contemplation could fundamentally change how we understand research itself.

DAY 2
THURSDAY
4:00 PM

kata agorein

Der Mensch ist das Maß

Guided by this principle, we at the kata agorein Foundation promote open dialogue in the public sphere on central questions facing our society. In keeping with our name — which stands for public speaking on the agora — we make room for open discourse, especially when different perspectives meet, and invite shared reflection to open new ways of thinking.

Our work is shaped by three core funding areas: the promotion of science and research, civic education, and civic engagement. A substantial part of this commitment is our support of the Ennoia Symposium. As a platform for exchange, the Symposium provides ample space for a patient, rigorous, and far-sighted exploration of the role of human beings in the age of AI, and allows us as an operational foundation to bring the scientific insights gained there to bear on a broader public discourse.

In 2027, we will award the first kata agorein Humanology Award on the question of Being Human in the Age of AI. The award honors scholars whose work examines how AI affects the human person. Its aim is to keep the question of what it means to be human at the center of public attention, as we enter this new technological age, and to pose it with the seriousness and responsibility that it deserves.

FRIDAY
05.06.2026

DAY 3

What benefits the Human



Chairs of the day

DAY 3

PROGRAM



SESSION A
Yehuda Halper



SESSION B
Tobias Rees



SESSION C
Maria A. Castaldi

- | | |
|----------|--|
| 9:00 AM | Opening of the Day |
| 9:30 AM | Session A — Talks |
| 10:45 AM | Coffee Break |
| 11:00 AM | Session A — Discussion |
| 12:00 AM | Long Lunch Break |
| 2:00 PM | Session B — Talks and Discussion |
| 3:40 PM | Coffee Break |
| 4:00 PM | Session C — Talks and Discussion |
| 6:00 PM | Writing Session: Charter of Research Questions (chairs of the day) |
| 7:00 PM | Piano Recital: Stefan Litwin |
| 8:00 PM | Joint Dinner at Harnack House |
| 10:00 PM | End |

SESSION 3A:

Does Human Formation Require a Duration That Cannot Be Accelerated — and What Do We Lose If We Take It Away?



Lucia Melloni

Professor of Neuroscience "Predictive Brain", Ruhr University Bochum; Visiting Professor of Neuroscience, NYU Grossman School of Medicine

Reverse engineering a self

Blending philosophy with computational cognitive science, I will present recent collaborative research showing how people can solve problems by performing a particular type of cognitive task: they can center themselves in a possibility space using a distinctive type of self-representation. I will suggest that something analogous in AI agents could improve efficiency and move us towards a new type of problem solving.



L.A. Paul

Millstone Family Professor of Philosophy and Professor of Cognitive Science at Yale University

The Proper Time of Consciousness: Why Human Becoming Cannot Be Accelerated

Artificial intelligence is forcing us to confront an uncomfortable possibility: perhaps much of what we call intelligence is just prediction at scale. Contemporary AI systems generate increasingly convincing simulations of thought by predicting what comes next, while neuroscience increasingly describes the human brain in similar terms, as a prediction engine continuously anticipating the world.

In this talk, I argue that prediction is not enough for human becoming. Human consciousness does not merely process information through time. It transforms lived time into meaningful episodes that can become memories, identities, regrets, hopes, and selves.

Drawing on recent work in predictive processing, event segmentation, and memory formation, I suggest that some of the most important features of human cognition arise not from optimization, but from its limits: forgetting, uncertainty, ambiguity, surprise, and the irreversible structure of lived experience itself. The real danger of AI, then, may not be that machines become too human, but that humans increasingly reorganize themselves according to the logic of machines.

Education in the age of AI confronts us with fundamental questions:

What remains distinctly human when knowledge can be generated, translated, and recombined by machines?

This talk argues that education must be understood not merely as the transmission and reception of information, but as a relational practice in which understanding grows through encounter, judgment, responsibility, and shared inquiry. It explores how AI can support learning in multifaceted ways while also challenging us to renew the human foundations of teaching and inquiry.

The central task is therefore not only to adapt education to AI, but also to shape educational AI with a humanistic perspective.



Niels Pinkwart

Vice President for Teaching and Research at Humboldt University Berlin

DAY 3
FRIDAY
9:15 AM

SESSION 3A:

Does Human Formation Require a Duration That Cannot Be Accelerated — and What Do We Lose If We Take It Away?

The Importance of Unpredictability in Fiction

Building fictional worlds and creating the people who live within them — characters who act, react, speak, and make decisions — is a distinctly human act. Characters in books, the ones that feel true to us as readers, are as erratic and surprising, as funny and self-destructive as the people we meet in real life.

A good novel reveals how a character thinks and feels, how she perceives others, how she copes (brilliantly or courageously or embarrassingly) in engaging, high stakes situations. Most novelists do not know what we will have our characters do or say as we begin writing a novel. Rather, we are constantly making choices along the way, trying to find the most interesting path.

The writer George Saunders said that as writers of fiction “we’re trying to make space for genuine surprise, for that feeling we get as readers when we feel a writer suddenly improvising in response to something they’ve just noticed being present in their own text.” AI, on the other hand, works on models of prediction. Predictability does not make for good story-telling.



Amy Pöppel

Novelist; Lecturer



Yoav Beirach

Postdoctoral Research Fellow,
Max Planck Institute for the History of Science

Human Experience and the Future of Quantified Time

The idea that two durations can be proven mathematically equal, treated as an independent quantity, and added together to form a count is relatively new. Before the conceptual birth of isochronism in the seventeenth century, the social and spiritual experience of time was woven into a cycle of natural proportions, none of them truly commensurate with the rest. With isochronism, this incommensurability gave way to a new form of count, built on the mere addition of fundamental equal units — seconds. When Walter Benjamin wrote in 1916 that “the nature of repetition in time is such that no unified form can be based on it,” he was pointing to a radical problem at the heart of European science.

In this talk I will discuss the seventeenth-century project of quantifying time, its effects on our contemporary temporal experience, and the ways in which AI might transform it again. As with earlier technological revolutions, AI may bring with it a new form of temporality, social and spiritual — and, as has happened before, the change may pass unnoticed. It is still too early to see what AI temporality will look (feel?) like, but to prepare for such an unknown, we should reacquaint ourselves with the forms of time that we have only recently forgotten.

DAY 3
FRIDAY
9:15 AM

Reditio Completa in Seipsum — Returning to the Heart of the Matter

We cannot create a ‘humane’ future without having a shared understanding of what ‘humanity’ is supposed to mean. The very concept of ‘the human’ has been challenged by technological innovation, the life sciences, climate change, and the dark history of humans treating each other as less than human. Nevertheless, humanity is the bedrock of our entire legal and political order; of human dignity and human rights. Particularly because we are confronted with an increasing entanglement of humans and technology, an emphasis must be placed not just on technical innovation, but on human growth, wisdom, and the virtues that allow us to create a future worth wanting.



Oliver Dür

Director, Center for Faith and Society,
Université de Fribourg; Oxford Collaboration
for Theology and Artificial Intelligence

(De)–anthropological trends in AI



**Gabriele
Gramelsberger**

Professor of Theory of Science and Technology,
RWTH Aachen University

Current developments in the fields of simulation and artificial intelligence (AI) have shown that the complexity of digital tools has exceeded the level of human understanding. We can no longer comprehend, understand or explain the results that AI delivers.

Even AI deceptions and hallucinations are now almost impossible to detect. This raises the question of the relationship between humans and their technology anew.

Are digital technologies as instruments useful extensions of human capabilities, as was understood in the classical philosophy of technology, or do we become extensions of technology? Do digital technology still have anthropological links to us or do we experience an increasing gap between digital technologies and us?

Advances in AI now make it possible to mine large, multimodal datasets to uncover mechanisms that could transform scientific discovery. Yet, predictive success alone does not guarantee scientific truth.

Identifiable AI algorithms offer a new approach: they recover structure that, under stated assumptions, corresponds to the data-generating process, providing a principled basis for forming and testing hypotheses about mechanisms.

Crucially, the value of such statistical guarantees grows as machine learning shifts from a tool applied to static datasets to an active participant in scientific feedback loops, proposing experiments, exploring data, and generating hypotheses.

This interaction makes AI and statistical literacy a condition for scientific inquiry: researchers must learn not only how to use AI systems, but how to audit, challenge, and verify them. Realizing this potential will require progress on two fronts at once: statistical rigor in the algorithms themselves, and equipping researchers with the skills to guide them.

Statistics in the Age of AI: From Prediction to Understanding



**Steffen
Schneider**

Head of Dynamical Inference Lab
at Helmholtz Munich

SESSION 3B:

Simulation, Algorithm, and Learning
Analytics — and What Do These Do
to the One Who Wants to Know?

Troubleshooting the Standard Model:



Klaus Blaum

Director and Scientific Member,
Max Planck Institute for Nuclear Physics,
Heidelberg former Vice President, Max Planck Society

Science Between Knowledge and Meaning —
Who does science serve – and what does it
mean to be human in the age of artificial intelligence?

These questions are particularly relevant in fundamental physics, which describes the elementary structures of reality. The Standard Model of particle physics has been extremely successful, yet it remains incomplete: it does not fully explain dark matter or gravity and contains unresolved parameters. The search for deviations from the Standard Model is therefore central to progress.

At the same time, AI systems that analyze large amounts of data and detect anomalies are transforming the scientific discovery process. When machines identify patterns that elude human intuition, the role of humans shifts.

This contribution argues that it is precisely the process of troubleshooting in fundamental physics that brings this tension to light: the tension between algorithmic analysis and the human search for meaning. Science serves not only to explain the world, but also to help humans find their place in it – a task that remains indispensable even in the age of AI.



Gal Chechik
(DISCUSSION ONLY)

DAY 3
FRIDAY
2:00 PM



Philip-Neri Reese OP


Pontifical University of Saint Thomas
Aquinas Angelicum, Rome

Truth and Knowledge as Trans- formative

It's tempting to think of truth and knowledge as something external – as a collection of facts about the world that can be more or less easily accessed.

But from a Thomistic perspective, this misses something crucial, namely, the transformative effect that knowing the truth has on us as human beings. And an awareness of that transformative effect can and should change the way we think about AI.

AI and the Spiritual Life



James Hankins

Professor of History, Harvard University;
Hamilton School of Classical and Civic
Education, University of Florida

The doctrinal note published by the Vatican in January 2025 under the title “Antiqua et Nova” and approved by Pope Francis before his death offers a starting point for reflection about AI and the spiritual life.

Intended for those exercising pastoral care in the Church, it raises important questions about the use of AI and its effects upon spiritual relationships and religious communities. While the positive benefits of AI are acknowledged, the dangers of idolatry and dehumanization are stressed. Is a person dehumanized by AI still able to participate in the life of the spirit?

I will discuss these matters from the perspective of Iain McGilchrist’s ‘hemispheric’ research on and analysis of human modes of reasoning (first laid out in his major work, *The Master and His Emissary* (2009). My discussion will suggest the paradox that a vast increase in the powers of the left hemisphere of the brain, aiming at superhuman power, has the effect of dehumanizing us, leaving us less than human.



Sarah Spiekermann

Professor of Information Systems, Vienna
University of Economics and Business

Cathedrals for AI?

How we fail to automate with visions that cater to the Good The talk looks at the conditions under which artificial intelligence (AI) can be designed to realize its true potential for progress.

Technology, as a novelty in and of itself, is not necessarily a good thing. The spread of fire represented a major advance for humanity, but only when we learned to tame it and designed its use for the good did it serve humanity and became controlled.

This is what we would need to do for AI as well. Globally standardised methods like Value-based Engineering are existing as well as the 10 rules for the digital world (derived from the 10 commandments). The trouble is only that nobody or few take notice of these and many other inclusive design approaches. As a result the talk concludes that we are far from our forefathers who built gothic cathedrals with a vision and drove progress in a way we are unable to today.

What Cannot be Automated — and What Does this Reveal about Judgement Core of Human Formation?



Enric Boix-Adsera

Assistant Professor in Statistics and Data Science,
Wharton School, University of Pennsylvania

Why AI Could Not Have Taken Us to the Moon

How did humanity get to the moon? Technology was indispensable, but to treat the achievement as a triumph of technical problem solving alone is to fundamentally misunderstand it.

I argue that the lunar landing is paradigmatic not of engineering prowess, but of human becoming itself, and I approach it through Castoriadis' concept of radical imagination. Radical imagination is not a sophisticated form of representation; it is the distinctly human power to rupture the given and institute what did not previously exist.

No algorithm, however advanced, possesses this capacity. AI can only operate within inherited forms and meanings—copying, recombining, and optimizing what has already been imagined. The moon landing, by contrast, was not the product of prediction, but of a creative leap into the unprecedented— something no machine can do.



Galit Wellner

Assistant Professor, School of Multidisciplinary Studies,
Holon Institute of Technology

The Verifiability Barrier

Recent breakthroughs across science and engineering — from AlphaFold's protein-structure predictions to advances in automated mathematics — point to a future in which AI contributes substantively across nearly every domain of human activity.

It can seem there is nothing these systems will not eventually be able to do. I argue, however, that our use of AI will run into a "Verifiability Barrier" — a limitation more subtle than questions of raw capability. I distinguish formally verifiable domains, where correctness of an AI's action can be checked automatically, from informal domains, where no such automatic check is available.

Drawing on my work in universal steering and multi-agent games, I argue that in informal domains AI systems are bottlenecked by our ability to specify what they should do: (1) a model's values reduce to those installed by its creator, and (2) without a formal verifier, agents are prone to reward-hacking that automated oversight struggles to catch.

Between Erosion and Formation: Human Capacity in the Age of Capable Machines

As AI capabilities advance rapidly, two gaps are widening in parallel. The first is organizational: a "capability overhang" between what AI systems can do and how organizations actually deploy them, creating stratified adoption across industries, social groups, and regions.

The second is cognitive: a growing body of evidence shows that sustained AI use can erode the very human capacities it is meant to support. Yet these outcomes are not inevitable. Research in human-AI interaction reveals that interaction design fundamentally shapes whether human skill, discernment, and understanding are preserved or displaced.

This raises a question I believe this symposium is uniquely positioned to address: if human capacities are sustained through practice, and AI is restructuring the conditions of practice itself, how must we design AI-mediated environments so that they remain genuinely formative?



Juho Kim

Associate Professor, KAIST
School of Computing

DAY 3
FRIDAY
7:00 PM
HARNACK-HAUS



In Search of Lost Subjectivity

Lecture–Recital with music
from Bach to Cage

Stefan Litwin (b. 1960 in Mexico City) studied piano, composition and repertoire in the USA and Switzerland. He has been on international concert tours with notable orchestras, conductors, chamber musicians and singers. Beyond that he cultivates the art of the lecture recital, with which he can comment and demonstrate selected works for piano and chamber music directly to the audience.

Aside from working as a performer, Stefan Litwin is also a composer. His newer works include two music theatre pieces based on the works of Peter Weiss (“Nacht mit Gästen” and “Wie dem Herrn Mockinpott das Leiden ausgetrieben wird”), the string quartet “Es wechseln die Zeiten”, and “Kinderszenen”, a composition for Ensemble that deconstructs the idyll in Schumann’s cycle of the same name, several Lieder and piano pieces as well as the full-length Monodrama “Flegeljahre” after Jean Paul, performed by actor Ulrich Noethen and the GrauSchumacher Piano Duo.

The musical prowess of Stefan Litwin is documented in TV and radio productions across Europe and the US as well as on several records. Recently the Wolke Verlag published a book with writings and conversations under the title “... als ob Musik die Welt verändern könnte.” [...as if music could change the world.]. Since 1992 Litwin is lecturer at the Hochschule für Musik Saar. He was Distinguished Professor of Music at University of North Carolina at Chapel Hill between 2008 and 2024.