ON WORKING WITH AI: REASONED RECOMMENDATIONS



Orientation – some recent personal anecdotes and thoughts about AI

In May our reunion of H-Star academic collaboration between Stanford and Denmark prompted me to reflect upon AI in research collaboration [Link].

A new white paper from Stanford's Accelerator for Learning – *Responsible Assessment in the AI Era* – prompted me to comment that we already have tried-and-tested pedagogical models that can embrace AI in learning [Link].

In June I made a trip to Germany to explore with Sabine Remdisch our Trans-nation Co-creation project [Link] in a series of visits to centers of AI development – including the new AI campuses at Heilbronn, the Rittal works at Haiger, European Factory of the year 2025, the Airbus plant in Hamburg, and Sabine's project with IKEA, modeling the AI-enhanced office of the future [Link].

A new report from MIT's Ad Hoc Committee on *AI Use in Teaching, Learning, and Research Training* [Link] is an excellent and progressive statement and represents, I suggest, a maturing of policy about the uses of AI in the academy and beyond.

As Gabriella Giannachi and I work on our new book *Archaeologies of AI* (see below), I have taken a few moments to compose this reasoned reflection on working with AI that summarizes what I am learning.

Recommendation: enthusiastically embrace work with artificial intelligence – openly, in research, writing, creating, designing, learning, teaching. Experiment, test capacities to limits – mindfully, critically, responsibly.

And don't defer to the caricatures of technology as demon / technology as savior.

This statement explains the understanding and reasoning behind this recommendation. Many dealings with AI in the academy are policy-oriented on administrative regulation, behavioral norms, the management of student and instructor experience – under an agenda to maintain or restore a supposedly healthy chain of reading, attendance, attention, unaided work, deadlines, testing, and supervision that existed before the pandemic and generative AI. Instead, one might focus on the character of knowledge and learning itself, treating the roll out of AI as an opportunity to rethink assignments, modes of participation, collaboration, and yes, the nature of learning itself.

AI is not a tool kit

A common way of speaking about AI is as a new set of tools, powerful and unpredictable in their agency and impact, that raise questions of regulation in their use and application, now and in the future. This underestimates what is happening.

I suggest we put to one side the metaphor of AI as a tool or tool kit.

Why? AI, and all that is implicated – machine learning, chatbots, artificial general intelligence, agency, data manipulation, task and project management, is much about knowledge and know-how. What is known, how that is generated, and how applied. A standard model of building knowledge is that it is fundamentally about perception and cognition, human minds acting upon engagements with the world,

learning and discovering how it works, and organizing and sharing the insights, the knowledge. Tools and instruments can help — augmenting perception (microscopes, telescopes), measuring (meters and gauges), communicating (writing, transmitting information), and much more. Here tools are something a mind-in-quest-of-knowledge picks up and puts down.

Building knowledge

But knowledge has never been made this way. Knowledge is not discovered by a detached observer contemplating a world held at arm's length. It is made: built through engaged, situated work with materials, instruments, archives, laboratories, field sites, institutions, colleagues, funding, publication — a whole heterogeneous ecology of participants working with and on one another. This is one of the oldest and best-established findings we have — from the philosophies of pragmatism, from epistemology, from the sociology and history of knowledge, from decades of close empirical study of how science and scholarship actually work. Truth itself is fabricated, in the old and honorable sense of the word: *verum ipsum factum* — the true is precisely what is made. And the maker is never outside the making. We are part of what we seek to understand; the knower is entangled with the known, and this entanglement is not a contamination of knowledge but its very condition. Building knowledge has always been project team work: assembling resources, coordinating different kinds of agency, negotiating with things and people that resist and surprise.

AI enters this ecology not as another tool on the knowledge-worker's bench but as a new kind of participant — one that generates, responds, proposes, errs, and provokes. Working with it is what Andrew Pickering calls a mangle of practice: an open-ended back-and-forth in which human intention and the agency of things accommodate one another, and something neither could produce alone takes shape. This is not a metaphor imported to flatter the technology. It is how experimental science, fieldwork, design, and scholarship have always proceeded. We act with the world, not upon it. AI makes this old truth newly visible — and newly consequential.

Wrong-headed policies

Most AI policies now circulating in the academy rest on a very different and much more fragile picture: the solitary author, the individual mind whose skills,

competencies, genius, and flaws are the sole source of the work, and whose unaided output is what (academic) examination and credential certify.

(Consider the concept of plagiarism in the academy and intellectual property. One must cite the originating mind behind an idea and not appropriate without compensation/acknowledgement what is conceived as the property of that originating mind. In an unresolved tension, at the same time knowledge is produced in institutions and communities of learning, where sharing, copying, collaboration, even theft of ideas and findings, are essential.)

In such a model of learning and knowledge, AI can only appear as a foreign, alien element in what should be pure — the mind working on the world to figure out what is going on. It is commonplace to hold that generative AI is at best a way to augment, speed up, automate, substitute for what are essentially individual competencies (analyzing, assessing, authoring); at worst AI is a way of faking what individuals cannot genuinely do themselves, because they are deferring to the tool kit.

This model of building knowledge was always a fiction. No scholar has ever worked unaided; no author has ever created alone; the individual talent has always been sustained by collaborators, editors, traditions, instruments, and institutions that the title pages of books and journal articles typically omit. And learning has never been well-described as the acquisition of a body of knowledge in the mind of the learner, transmitted by an instructor. Policies built on this fiction are defending an authenticity that never existed, and this is why they misread the present moment so badly: they experience the visibility of collaboration as a crisis of fraud.

Trust and participation

This misreading has consequences, and we should voice objection. A policy that begins from suspicion ends in policing. Across the university sector we now see the apparatuses of surveillance and policing assembling: detection software, proctoring, overseeing behavior and “effort,” rituals of exposure and confession in which students must clear themselves of the presumption of fakery, and the steady conversion of pedagogical questions into questions of compliance with administrative norms. Whatever this is, it is not education. It corrodes the one condition that learning cannot do without — the collaborative trust between teacher

and student in participatory-learning that makes real intellectual risk-taking possible — and it teaches the wrong lesson thoroughly: that what matters is not the work but the appearance of having followed the rules.

Nor should rigor be equated with the absence of tools. Rigor lives in the work — in the seriousness of the problems engaged, the discipline of the craft, the honesty of the judgement — not in the conditions imposed around it. Where education has become hollow, a credential acquired while the experience of learning is abbreviated or skipped, the remedy is not to police the hollowness but to redesign the experience until it is worth having. That is harder work than enforcement. It is the work.

From the understanding set out here follow some specific commitments and recommendations.

- **Engage** AI as a collaborator whose contributions, like those of any collaborator, should be tested, questioned, and judged.

- **Transparency**. Disclose AI involvement in work, acknowledge any substantial intellectual debt. Open up the black-box of research, learning, authoring.

- Remain **answerable** for collaborative project work: within an ecology of many agencies, accountability is precisely the human part, and it cannot be delegated.

In teaching, the same understanding applies.

Invite students to work with AI transparently and critically, and to treat that work as itself an object of inquiry: what did this exchange produce, what did it foreclose, where did the machine’s proposals need resisting, what was learned in the resistance? Some capacities are built only through unmediated engagement with materials and difficulty, and parts of our work will deliberately set the machines aside — always with the reasons given, because a rule without a rationale teaches only compliance. The larger aim is the formation of practitioners who can navigate

a knowledge ecology that now includes these agencies — with judgement, honesty, and skill.

Knowledge ecologies

Such recommendations reach well beyond the academy.

Responsible AI use also requires attention to the wider ecology from which AI systems emerge.

These systems are sustained by data, energy, water, computational infrastructures, institutional investment, human labor, intellectual property arrangements, and corporate priorities. They involve questions of political economy, of access, privacy, ownership, bias, environmental cost, and dependency.

Critical competence therefore means asking not only whether an output is useful or correct, but also: What system produced it? What assumptions and exclusions shape it? What information am I providing to it? Whose work and knowledge sustain it? What dependencies am I creating? What consequences extend beyond the immediate task?

Such questions are not external additions to learning. They are part of understanding the socio-technical conditions under which contemporary knowledge is made.

The capacities at stake here — working openly with many kinds of collaborator, human and otherwise; testing what one is given rather than accepting or forbidding it; matching the mode of work to the situation at hand; remaining answerable for the whole — are the working core of creative and effective practice everywhere: in design studios and research labs, in business and enterprise, in civic life and community governance. Every one of these is an ecology of knowledge and action in its own right, with its own materials, constraints, and stakes. What I describe here is not a compliance regime for coursework. It is a model of working well, sensitive to context, in a world where the company we think and work with with has grown.

This understanding is not mine alone, far from it. It draws on long traditions of thought about how we relate to the world and build knowledge within it. and is a

widening community of educators, researchers, and institutions committed to integrating AI thoughtfully into learning and research rather than policing it. This is exemplified most recently by MIT's Institute-wide report on AI in education (2026), which rejects surveillance and detection, insists that every AI policy carry a pedagogical rationale, and calls for the renewal of experiential, project-based, socially embedded learning: augmentation, not automation [Link].

The educational case for project-based studio-based active and participatory learning is made in the book *Creative Pragmatics for Active Learning in STEM Education*, edited with Connie Svabo, Tamara Carleton, Chunfang Zhou (2025) [Link].



The company gathered around the instrument. Joseph Wright of Derby, *A Philosopher Giving a Lecture on the Orrery* (1766) — candlelight, faces of all ages leaning in around a machine that models the cosmos, wonder rather than dread. It is public domain, endlessly reproducible, and its subject is literally an epistemic ecology: knowledge happening *between* people and a device.

Archaeologies of AI: archives, robots and the ecology of intelligence

Such an argument about AI is developed at length in a book I am now completing with Gabriella Giannachi, *Archaeologies of AI: archives, robots and the ecology of intelligence*. We ground an argument for embracing and experimenting with these new companions in the long-term story of human ecologies — the hundred-thousand-year record of a species that has never thought alone, and has always made itself through what it makes.

Archaeologies of AI argues that artificial intelligence is best understood as archival intelligence: the automation and scaling of operations — inscription, classification, summarization, retrieval, delegation — that societies have performed within infrastructures of memory for five thousand years. Against the two framings that dominate current debate — AI as unprecedented rupture on a path to superintelligence, and AI as a narrowly technical matter of engineering and governance — we offer a third: an archaeology of the present, tracking continuities and discontinuities in how knowledge has been stored, staged, delegated, owned, performed.

Our approach is distinctive in combining science and technology studies (knowledge as performance) with archival theory, archaeology, and narratology. We show that whenever societies address what is now called AI, they draw upon a durable repertoire: personae (the custodian, the reader of signs, the oracle, the hidden operator), plots (the runaway, the counterfeit, the animating word, threat and subterfuge), and settings, architectures, from the Mesopotamian temple store to the data center. This mythology is not commentary on the technology; it guides action — pricing companies, drafting regulation, scripting fear and desire. What is genuinely new about AI lies not in intelligence, which was always distributed and performed rather than possessed, but in political economy: scale, speed, capitalization, and planetary substrate. The archive is the past captured; capitalization is the future pledged; contemporary AI stands where the two meet.

The message is practical, and distinctively so. Intelligence, on our account, is an ecology (oiko-logy — the management of habitat): performed across bodies, tools, institutions, archives, and environments, in a coevolution of humans and technics as old as knapped stone, fired clay. The ecology of intelligence names the distributed milieu — bodies, tools, institutions, archives, environments — within

which knowing is performed; ecology is thus an aspect of the performance of knowledge, the frame we take from science and technology studies.

“Artificial” intelligence is not the opposite of the living but a new participant in life’s oldest habit(at) of composing novelty through partnership — sympoiesis. Taking symbiosis seriously — after Margulis and Bateson, for whom the unit of survival is organism-plus-environment — dissolves the binaries of technology and biology, machine and milieu, and with them the fantasy of the autonomous machine. In their place we offer working instruments: responsible embodiment, a diagnostic locating responsibility across the material, labor, perceptual, ecological, archival, and temporal dimensions of distributed systems, including accountability for the futures a system pledges; and a design orientation — acting-with our machines, building commons of archive and memory, labor and energy. The horizon is sur-vival: living-on together, the shared maintenance of human and machinic intelligence within the fragile living systems on which both depend.



Working with clay. An Egyptian statue [Link] confabulated as the clay Golem. The word traditionally written on the forehead is emet (□□□), the Hebrew word for “truth”. According to Jewish folklore, particularly the famous legend of the Golem of Prague, writing this word on the clay creature’s head brought it to life. To deactivate, one would erase the first letter (aleph), changing the word to met (□□), which means “dead”.