Input for the National AI Committee

Matthew Landers
Associate Professor
Høgskulen på Vestlandet

1. Hva er den største utfordringen med KI i høyere utdanning? (What is the biggest challenge with AI in higher education?)

The biggest challenge with AI in higher education is with trying to understand the big picture. Where is this technology going? Many of the debates in HE revolve around what AI models can do now, when the real issue is what will AI models be able to do in 2, 5, 10 years? In fact, it may be more accurate to say that many debates in HE center on what AI models could do 2 or 3 years ago, as many academics either refuse to stay informed about the state-of-the-art, or they simply cannot keep up. This means that many academics simply do not know where the technology is, or where it is going. Cambridge professor, Harry Law, wrote recently:

> But right now, many academics who speak to policymakers or the press badly underestimate the capabilities of the best models. They dismiss LLMs out of hand and don't engage with the substance of the technology.
>
> […]
>
> Accepting the reality of the situation is the best way to produce timely and relevant work. But that requires getting familiar with the technology so that the public debate is grounded in clarity (Law, 2025).

I discuss the difficulty involved with understanding the technology below.

2. Hvordan ser teknologien og høyere utdanning ut ti år fram i tid? (What will technology and higher education look like ten years into the future?)

Early on, AI labs identified two knowledge domains that their models would likely disrupt first: Medicine and Education. A decades-long love-affair between education and technology sectors has paved a smooth path for rapid deployment of AI in the lower- and higher-education spheres. A global focus on developing digital competencies means that students now master the use of tablets and computers from a very early age. Digital tools and internet connections are ubiquitous in every classroom, making access to AI models likewise universal.

Predicting the near- and far-future of AI is a difficult task. Even within the AI sector, timelines and predictions about AI vary. However, most AI labs (OpenAI, Anthropic, Google DeepMind, xAI) agree that achieving the industry goal of artificial general intelligence (AGI) (Landers, 2025) within the next 10 years now appears likely (Browne, 2025). If this prediction proves to be

true, traditional higher education institutions (HEIs) will be under growing pressure to justify their existence as knowledge gatekeepers. Not only will AGI provide expert-level knowledge across every domain of human inquiry, it will offer a 'gold standard' of education—personalized and adaptive learning experiences for students/users—potentially decentering the traditional classroom as the social locus of instruction.

Frontier AI models have already disrupted traditional education in significant ways. In 2025, Gemini and OpenAI introduced "Deep Research" functionality, which can compose thesis-length research papers in mere minutes, making use of tools to search the internet for academic articles. The next five years will see even more education-focused applications come to market. Importantly, a recent investigation shows that the performance of AI models on long-horizon task benchmarks has improved *exponentially* over the last five years, with performance measurements now doubling every seven months (Kwa et al., 2025, p. 23). This is an incredible rate of progress, which (if extrapolated just a few years into the future) should sound alarms in the education sector.

Conversations about how these technologies will (or should) be used by students and faculty at HEIs, has not kept pace with the current rate of advancement. Furthermore, an exponential rate of improvement means that long-term predictions about the impact on education are more difficult to make, not less. Our imaginations cannot reliably grasp the import of such a rapid pace of advancement. The task is made more difficult by the fact that many of the abilities that emerge in frontier models cannot be predicted before they appear (Landers, under review). Recent attempts to predict the next **two** years(!)—based on data, insider knowledge, expertise, and trends—have produced incredible narratives about the near future (Kokotajlo et al., 2025).

3. Hvordan kan utdanningene forberede studentene på et samfunn og arbeidsliv der KI får økt betydning? Hva blir viktigere, og hva blir mindre viktig å lære? (How can the programmes prepare students for a society and working life where AI is becoming increasingly important? What becomes more important, and what becomes less important to learn?)

The obvious answer is that students need to use the technology in preparation for the future. However, use of AI technology cannot be handled in the same haphazard way that the push toward digitalization (*digitalisering*) in education was. The need to prepare students to use AI has to be balanced with meaningful training in its ethical use. AI is a technology unlike any other technology in the history of our species. Intelligence cannot be confused with simple digital tools. We cannot expect students to 'figure out' how to use AI ethically on their own. This will require significant scaffolding (far beyond a set of regulations and an ethics workshop). The task is made more difficult by the fact that 'ethical AI'—as it relates to education—is currently in a nascent state of discovery and debate. We cannot let the cart get before the horse, as we may have done with other technologies.

4. Hvordan skal man gjennomføre eksamen og vurdering? (How should you carry out the exam and assessment?)

AI changes the plagiarism/cheating paradigm. In the age of AI, we cannot think about assessments from the traditional perspective. I have written elsewhere:

> From an administrative standpoint, the bigger problem is that AI detection tools do not provide educators with forensic evidence, portions of both the original and plagiarized text, to compare in the investigation of cases of cheating. In traditional cases of plagiarism, teachers need only locate the original published text being copied, or claimed as one's own work, to prove the case. With LLMs, there are no original published texts to discover. Evidence in the traditional sense does not exist. Rather, one must trust a detection algorithm tool to predict (accurately) the probability that a text has been generated by AI. Unlike traditional cases, with LLMs, the only 'evidence' available is the prediction of an algorithm. Understandably, university administrators have been squeamish about pursuing investigations of cases of suspected plagiarism using such non-traditional and potentially unreliable evidence. For this reason, AI-based plagiarism cannot be understood within the same paradigm as traditional plagiarism. The nature of the problem has changed. With that change should come the recognition that 'proof' in the traditional sense may not be an achievable goal in the investigation and prosecution of plagiarism for some time to come. Thus, our efforts should be directed at prevention, not detection (Landers, 2024).

If the approach shifts from detection to prevention, assessments should be designed to prevent attempts to cheat. There are many ways to accomplish this: hand-written exams, in-school exams, oral exams. Many universities have already pivoted away from the most obviously vulnerable form of assessment: the traditional take-home essay exam. This form of assessment is vulnerable because the actual work of the student is obscured from the assessor; and because the value of the assessment relies exclusively on the academic integrity of the assessed.

There are ways to design assessments that include the regulated use of AI, but this is a different question.

5. Hvilke tilpasninger av vurderingsformer kan gjøres i lys av KI-utviklingen? (What adaptations of assessment forms can be made in light of AI developments?)

Any recommendations that one makes here will likely be out-of-date by next year. This is why I have recommended an organic and iterative approach, in dialog with expertise at the centralized and local levels. I have written elsewhere:

> Going forward, when educators design assessments for their courses, we cannot think solely about how well the task supports learning outcomes. We must consider first how easily certain elements of a given task may be completed by AI-supported technology. We must be familiar enough with the capabilities of the technology to identify vulnerabilities in the design of assessments. In the process, we may find that many traditional tasks—especially those that require students to summarize information and/or compare

perspectives, but which do not encourage strong engagement with metacognition and/or critical reflection—are now *obsolete* (Landers, 2024).

While these claims are still up-to-date, it should be noted that the newest "reasoning models" have drawn closer to exhibiting cognitive capabilities that resemble metacognition. It is no longer unreasonable to expect frontier models to perform well on assessments that require sustained reasoning and critical thinking. *Viva voce* demonstrations of student knowledge may be one of the best ways to ensure that learning is taking place.

6. Hva slags ferdigheter trenger studentene i bruk av KI, og bør utdanningsinstitusjonene gi spesifikk opplæring og veiledning? (What kind of skills do students need in the use of AI, and should educational institutions provide specific training and guidance?)

Right now, the selection of models that are available for use is deeply fragmented. OpenAI alone has seven models available for use, each with slightly different strengths and weaknesses. Some models "hallucinate" more than others. Some models handle data analysis better than others. While we wait for a global model that does everything well, students should become familiar with the current offerings, how they operate, what their strengths and weaknesses are, and which to use when.

I do not consider ethical knowledge a skill, but it does define boundaries for how students engage with and use the technology. This understanding can develop into an important professional skills. The most important thing that HEIs can provide for students is an understanding of the ethics of AI use, as we are uniquely equipped to provide this support.

7. Hvilke muligheter og utfordringer gir KI i undervisning og læringsaktiviter? (What opportunities and challenges does AI provide in teaching and learning activities?)

With AI, intelligence is quickly becoming commodified. AI labs hope to spark a second wave of automation that rivals (and perhaps surpasses) the Industrial Revolution. This means that AI will be used primarily to reduce the amount of labor that we perform for our jobs. Because the automation currently targets intellectual labor (not physical labor, yet), it means that we are offloading cognitive effort. This offers opportunities and challenges for educators.

Opportunities: Students can use AI to help them understand complex subjects. They can extend their own understanding by asking questions that the model can answer or guide them through. The model understands the student's level of understanding and can adapt to the learning needs of the individual. Education can be sped up, or slowed down to personalize learning.

Challenges: Models can be prompted to do the cognitive work for the student. Increasingly, models can complete complex tasks that would have required human intelligence in the past, which means that students are tempted to offload cognitive effort to the model. Frontier AI models are more accurate than ever, but even the most advanced 'reasoning models' still

“hallucinate” information and present it as factual occasionally. AI labs have not successfully cracked the “alignment problem,” which means that AI models may not be fully aligned with human values. Human values are a generalized construct; further alignment with Norwegian values remains an open problem.

8. Hva har KI å si for masteroppgaven og andre langsvarsoppgaver? (What does KI have to say for the master's thesis and other essays?)

Masters theses are a capstone text that students write to show mastery of both research methods and a subject area. Students are awarded M.A. credentials based on the quality of the thesis, making it the most important artifact of their learning journey. One year ago, frontier models were not able to write a thesis-length research text, but the recent launch of “Deep Research” capabilities has changed this. Significant portions of a masters thesis can now be produced in minutes with very little prompting. If we wish to preserve the integrity of the masters thesis, HEIs will need to craft very clear guidelines about the use of AI. These guidelines, as discussed before, will need to be fine-tuned to reflect the values and requirements of individual subject areas. One set of guidelines will not fit all circumstances. For example, decisions will need to be made about how much (if any) data analysis can be turned over to AI. Anonymous data will have to be treated differently, prompting the used of different kinds of model (local vs. hosted in the cloud). Guidelines for different subject areas will need to be thought through and individualized to avoid misaligned regulations.

9. Hva skal til for at faglige ansatte har kompetanse til å håndtere KI på en utdanningsfaglig god måte? (What does it take for academic staff to have the competence to handle AI in a good educational way?)

I have argued elsewhere that “AI systems are progressing at such unprecedented rates that it will be impossible for most educators to achieve and maintain an acceptable level of knowledge about the current capacities of AI-supported technology” (Landers, 2024). As an alternative, I proposed the following approach:

> […] if institutions of higher learning wish to address these vulnerabilities successfully, they will need to create dedicated departments/committees to research and stay up-to-date with the state-of-the-art. These departments would need to act as consultants for faculty members and administrators, relieving both parties of the impossible task of maintaining high levels of competence in the domain.
>
> The proposed department would develop adequate expertise about the current capacities—and imminent advances—of AI systems. In addition, it would develop adequate expertise about AI-supported tools and applications that are being brought to market. With this expertise, the department could generate and maintain a list of sanctioned and unsanctioned AI-supported tools, while providing guidance for faculty about state-of-the-art AI-detection. Lastly, the proposed department would act as a

> technological and educational hub for faculty and administrators, ensuring that an adequate level of competence is maintained at the institutional level (Landers, 2024).

In short, HEIs cannot expect academic faculty to develop and maintain the competencies needed to understand artificial intelligence at a high enough level to be able track rapid advancements in the technology, while also maintaining high-level competencies in their own areas of research. I assert that "innovative institutions will need to develop internal specialists to bridge the gap between expert and non-expert competency levels. Universities that fail to supply the necessary expertise will create a massive knowledge gap in the AI readiness of their faculty and administrators" (Landers, 2024).

10. Hvilke behov er det for nasjonale tiltak og løsninger kontra lokale når det gjelder KI i høyere utdanning? (What are the needs for national measures and solutions versus local ones when it comes to AI in higher education?)

Expertise cannot be centralized. Expertise should be distributed in a way that facilitates an effortless sharing of knowledge at the local level. Otherwise, decision-making will be slow and strenuous, hindered by distance from knowledge-centers and by bureaucratic processes. Bureaucracies should be created to make local decision-making easier, not more laborious. We want individual HEIs to be part of the innovation, not just to follow rules without understanding.

11. Hvilke særtrekk ved utdanningene – fagområder og enkeltutdanninger – fordrer at KI blir håndtert annerledes enn i høyere utdanning generelt? (What special characteristics of the education programmes – subject areas and individual programmes – require that AI be handled differently than in higher education in general?)

The use of AI should **always** be contextualized for the subject area. Each subject area should be given the power to fine-tune regulations for AI use in that subject, dependent on the skills that students are expected to acquire to master that subject. For example, in subjects that have a language requirement (Norwegian & English), where the student's language-level is being assessed, it would be more important to consider restricting the use of AI for writing or revising texts. Otherwise, teachers lose the ability to assess student knowledge and progress. The student's actual knowledge is obscured by the AI's abilities. In turn, we run the risk of graduating students who lack the very skills that a degree certifies they have at the end of their education. This would represent an act of fraud, since the whole point of accredited HEIs is to certify that graduates leave with the knowledge and skills required for their degrees. HEIs that cannot vouch for the knowledge and skills of their graduates damage the reputation of education in Norway.

Works cited:

Browne, R. (2025, March 17). *AI that can match humans at any task will be here in five to 10 years, Google DeepMind CEO says*. CNBC. https://www.cnbc.com/2025/03/17/human-level-ai-will-be-here-in-5-to-10-years-deepmind-ceo-says.html

Kokotajlo, D., Alexander, S., Larsen, T., Lifland, E., & Dean, R. (2025, April 3). *AI 2027*. https://ai-2027.com/

Kwa, T., West, B., Becker, J., Deng, A., Garcia, K., Hasin, M., Jawhar, S., Kinniment, M., Rush, N., Arx, S. V., Bloom, R., Broadley, T., Du, H., Goodrich, B., Jurkovic, N., Miles, L. H., Nix, S., Lin, T., Parikh, N., … Chan, L. (2025). *Measuring AI Ability to Complete Long Tasks* (No. arXiv:2503.14499). arXiv. https://doi.org/10.48550/arXiv.2503.14499

Landers, M. (2024). Adapting to the Unsanctioned Use of AI-Supported Technologies in Student Assessments. *Higher Education for the Future*. https://doi.org/10.1177/23476311241300608

Landers, M. (2025). Modern Prometheus: Tracing the ill-defined path to AGI. *AI & Society*. https://doi.org/10.1007/s00146-025-02363-1

Landers, M. (under review). Reclassifying Frontier AI Models as a Complex Emergent Technology. Submitted to *Philosophy & Technology*.

Law, H. (2025, June 17). Academics are kidding themselves about AI. https://www.learningfromexamples.com/p/what-academics-get-wrong