



Beyond automation and offloading

Rethinking Human–AI Collaboration for Learning

Learning technologies, interaction design, and human–AI interaction

Michail (Michalis) Giannakos

- Professor of Interaction Design & Learning Technologies
- Head, Learner–Computer Interaction Lab
- Co-Director, AI LEARN Centre



NTNU · Department of Computer Science



NTNU • TRONDHEIM

NORWEGIAN UNIVERSITY OF SCIENCE AND TECHNOLOGY

KNOWLEDGE FOR A BETTER WORLD

TRONDHEIM | GJØVIK | ÅLESUND
BRUSSELS





More than 100
million monthly
active users



Kahoot App Brings Urgency of a Quiz Show to the Classroom

By NATASHA SINGER APRIL 16, 2016



Three questions for this keynote

01 THE DILEMMA

Automation dilemma
Balancing automation
and control

02 DESIGN


Human-Centred AI
Human learning and HCAI
Educational GenAI
examples

03 THE FUTURE

Human capabilities and
performance
Beyond HCAI to Hybrid
Intelligence

From purpose-built systems toward contextual AI

The first two waves addressed specific, often narrow problems.



Source: Defense Advanced Research Projects Agency (DARPA) information; Art Explosion (art). | GAO-18-142SP

AI capabilities are growing—but what about ours?

AI PERFORMANCE


Models now meet or exceed human baselines on several complex benchmarks.

THE OPEN QUESTION

Does interacting with AI make humans more capable over time?

Select AI Index technical performance benchmarks vs. human performance

Source: AI Index, 2026 | Chart: 2026 AI Index report




A provocative example: learning to identify cancer


BREAST CANCER DIAGNOSIS

85% individual accuracy
99% in group diagnoses

Following two weeks of training,
"novices" transferred their learning to
unseen images.



High accuracy is not the same as appropriate replacement



A flock of pigeons was able to correctly spot cancer in breast tissue biopsy images with 99% accuracy, on par with human experts. R. M. LEVENSON ET AL., PLOS ONE (18 NOVEMBER 2015)


Should pigeons
replace
pathologists?

Two foundational design questions

What should be automated—and to what extent?

How do we preserve meaningful human control?

Driving automation: six levels of delegation



(Society of Automotive Engineers, 2016)

When to automate and when to maintain human control

GOOD CANDIDATES FOR AUTOMATION

- Reliable, predictable work (backups, spell-checking)
- Repetitive or high-volume tasks (routine reports)
- Low-risk, reversible actions (brightness settings)
- Keep overrides available (AI feedback, autopilot)

KEEP MEANINGFUL HUMAN CONTROL

- Creative or context-sensitive work (lesson design)
- Uncertain judgments (ambiguous diagnoses)
- High-stakes decisions (medical, financial)
- Learning and agency matter (solving math problems)


Automation and control need not be opposites

THE ONE-DIMENSIONAL VIEW


More machine automation necessarily means less human control.

THE TWO-DIMENSIONAL VIEW

High automation and strong human control can be designed together.





Human-Centered AI Framework



Hospital control centres: a high-stakes design example



Autocompletion and autocorrection: everyday examples of automation



A longer history of augmenting human intelligence

1968 · ENGELBART

Intelligence augmentation: computers extend intellectual capacity, not merely replace it.

Example: computer mouse.

AUGMENTED COGNITION

Systems monitor and support human cognitive processes, enabling new partnerships.

Examples: dashboards, our calendars.

THE LIMITATION

Performance gains may still come mainly from delegating cognitive processes to technology.

When it comes to human learning there is a misalignment: "success at a task" versus "learning it"

- The biggest challenge is the misalignment (or antagonistic nature) between automation (complete a task successfully) and human learning (mastering a task, if it's automated why/how humans will master it).

AUTOMATION/AI OPTIMIZES


Immediate task completion
Throughput · efficiency · accuracy

LEARNING OPTIMIZES

Mastery · transferable competence
Understanding · learner agency

- When AI completes the cognitive work, how does the learner develop the capability?

Four paradigms for learning-oriented HCAI



From general-purpose performance tools to learning partners

AI AS OFFLOADING


The system takes over cognitive work to complete the task.

AI AS CAPABILITY BUILDING


The interaction strengthens the human while keeping tasks effective.

Design for human capabilities—not task performance alone.


An example of how a learning technology can be designed differently




Three LLM configurations



01 No AI support





02 Restricted feedback




03 Unrestricted feedback

Cognitive engagement across LLM configurations




Different levels of AI support shape students' experience.

AI-enhanced feedback with LearnLab



AI-ASSISTED DRAFT → TEACHER REVISION → ACTIONABLE FEEDBACK

Human revision turns AI drafts into useful guidance



CLEARER

Teachers simplified wording and removed repetition.

ACTIONABLE

Teachers added concrete pedagogical next steps.

All 15 teachers revised the generated feedback before use.

An LLM-supported programming assignment interface

THE LEARNER VIEW

Students can request assistance from an AI-powered chatbot during task completion.

Assignment 1

List manipulation

Task 1

Task 2

Task 3

Home

Task 1

Create a function `count_words()` that takes in `word_list`: a list of words on the form `["word", "excel", "powerpoint", "office"]` and `letter`: a single letter for example `e`. The function should return a list of numbers for the amount of times the letter appears in each word. For example: `count_words(["word", "excel", "powerpoint", "office"], "e")` would return `[0, 2, 1, 1]`

```
print("Hello world!")
```

Run Code

Chat with AI

×

Hi, how do I add two numbers in Python?

Great question! Adding two numbers in Python is straightforward. You use the `+` operator to add them together. Here's a simple example:

```
a = 5
b = 3
sum = a + b
```

Type your message...

Send

An LLM-enabled assessment and follow-up interface

THE ASSESSOR VIEW

The system evaluates submitted code and engages students in structured follow-up questions.

Assessing Task 1

```
def count_words(word_list, l):  
    new_list = []  
    for word in word_list:  
        count = 0  
        for letter in word:  
            if l == letter:  
                count += 1  
        new_list.append(count)
```

Hello! I'm an AI assessor for your CS1 course, here to evaluate your code solution by asking a few questions to ensure you understand what you've written. Let's get started with your code: What does the function 'count_words' do? Please explain the process it follows to achieve its task.

The function counts the number of times a letter appears in each word in a list of words. It takes in a list and a letter, and iterates through them to see how many times the letter appears in the words.

Can you explain the first for loop in the code snippet in detail? Specifically, what is its function and how does it operate within the broader context of the code?


The first for loop iterates through each word in the input list. It creates a count variable and continues to loop for each letter in the word.

Could you elaborate on the role of the 'if' statement within the second for loop and explain how it contributes to the function's overall operation?

Type your message...

Send


Assisted assessment: promising, but not yet dialogue-rich




WHAT THE STUDY FOUND

- The chatbot asked relevant, code-specific questions at an appropriate difficulty level.
- Students reported lower trust and fewer opportunities for clarification.
- The chatbot struggled with follow-up questions and dialogic feedback.

Traditional embodied learning of integer arithmetic



MOVES Number Line • UC Berkeley + NTNU




An AI-powered multimodal learning environment

CAPABILITIES


Content selection · IRT / BKT
Multimodal decision-making

Voice interaction · NLP
Face recognition · CNN

Content and feedback
generation · LLMs
Personal tutoring · agentic AI



Embodied learning can improve inclusion and efficiency



Cosentino et al. (2025) · Anton et al. (2026)

INCLUSION

Significant gains for learners with disabilities ($p = .0048$; $d = 1.44$); the achievement gap narrowed.

AI + HUMAN FEEDBACK

Comparable task performance; lower cognitive load in the GenAI condition, and more focus on visual processing with GenAI.

Human capabilities & task performance


THE COMMON ASSUMPTION

A common assumption is that prioritizing human learning and competence development compromises efficiency

THE DESIGN CHALLENGE

The challenge is to design systems that enhance human capabilities (or at least do not diminish them) while improving task performance

Adding another dimension




HCAI begins with people's needs and contexts

- HCAI is tech designed with the end user firmly in mind, **responding to our needs and desires**.
- Whether it takes the form of an app, device, or website, **it utilizes AI to enhance our ability to achieve our work, study or social interactions** by prioritizing the human over the tech requirements

HCAI is intuitive, allowing users to complete a specific task without the need of prior training.

Think of examples such as Spotify, FitBit, and Kahoot! that provide solutions to real problems and desires.




What evidence should we collect to evaluate learning?

MEASUREMENT GAP

- Most data we collect focus on task performance.
- Deep/sustained learning is difficult and demanding to be measured
- Long term cognitive development is rarely considered
- We get what we measure.

COMMON EVIDENCE

- Questionnaires
- Interviews
- Observations
- Interaction logs
- Analytics



Four implications for educational AI

01 ALIGNMENT

Pedagogical alignment comes first.

02 TASK PERFORMANCE

Task performance is potentially supporting evidence, not the primary goal of learning technologies.

03 METRICS

Misaligned KPIs reward automation at the expense of learning.

04 CONTEXT

“Good performance” depends on the learning task and objectives.

So, how do we go beyond automation and offloading?

Technology is never neutral. What we automate shapes what people learn.

01

AUTOMATE SELECTIVELY

Offload routine tasks

Calendars, logistics and repetitive work can be delegated without sacrificing learning.

02

PROTECT THE LEARNING

Preserve constructive struggle

Keep reasoning, reflection and other capability-building work with learners.

03

PRESERVE AGENCY

Keep humans in control

Align AI support with pedagogy, learner age and meaningful human choice.

DESIGN TEST

Does this automation make the learner more capable over time?



Hybrid Intelligence

A lens for research beyond automation and offloading

Learning systems are dynamic—are our frameworks?

HUMAN NEEDS

Agency and relationships change over time.

AI CAPABILITIES

The technology itself is continuously evolving.


SYSTEM ADAPTATION

AI behaviour and functionality are dynamic.

SOCIO-TECHNICAL CONTEXT


Humans and AI co-evolve within their environments.

Operationalising a mutually developing partnership




01 • MUTUAL UNDERSTANDING

An interactive process builds shared understanding.



02 • ADAPTIVE FEEDBACK

Continuous feedback creates a positive loop.



03 • EXPANDING CAPABILITIES

Both sides grow their capabilities over time.

Build intelligence together—rather than offload learning

01 AUGMENT

AI should expand human intellect instead of replacing it.

02 RELATE

Students and teachers develop trust, agency and relationships with AI.

03 CO-EVOLVE

Use human expertise and intentionality to guide evolving human–AI partnerships.

04 MISALIGNMENT

Consider/avoid technologies that reward automation at the expense of learning.

05 TASK PERFORMANCE

Task performance is not the primary goal of learning technologies

References

- Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human–Computer Interaction*, 36(6), 495-504.
- Nguyen, A., Wise, A., Rosé, C., Giannakos, M., Järvelä, S., & Cukurova, M. (in press). Hybrid intelligence and collective learning in human–AI systems. In R. Wegerif, I. Casebourne, A. Zhou, & I. J. Ness (Eds.), *Handbook of AI and the future of education*. Springer.
- Giannakos, M., Azevedo, R., Brusilovsky, P., Cukurova, M., Dimitriadis, Y., Hernandez-Leo, D., Järvelä, S., Mavrikis, M. & Rienties, B. (2025). The promise and challenges of generative AI in education. *Behaviour & Information Technology*, 44(11), 2518-2544.
- Cukurova, M. (2025). The interplay of learning, analytics and artificial intelligence in education: A vision for hybrid intelligence. *British Journal of Educational Technology*, 56(2), 469-488.
- Anton, J., Cosentino, G., Gelsomini, M., Sharma, K., Giannakos, M. N., & Abrahamson, D. (2026). Towards minding the disability achievement gap: the case of technology-enhanced inclusive embodied design for integer arithmetic. *Educational technology research and development*, 1-21.
- Dimitriadou, E. A., Stokke, S. A., Hegreberg, S., Nguyen, A., & Giannakos, M. (2026,). Can LLMs Support Formative Assessment? LLM-Based and Teaching Assistant–Based Assessment in Programming Education. *International Conference on Artificial Intelligence in Education (AIED)*, 94-102.
- Kleveland, J., Giannakos, M., & Lindvig, Y. (2026). Teacher–AI Collaboration to Support Assessment and Feedback: A Case Study in Norwegian Secondary Education. *Evaluation Review*, 0193841X261458416.
- Kleveland, J., Giannakos, M., & Lindvig, Y. (2026, April). Exploring Human–AI collaboration for formative feedback: Insights from a K–12 case study and implications for analytics. In *Proceedings of the LAK26: 16th International Learning Analytics and Knowledge Conference* (pp. 772-778).
- Cosentino, G., Anton, J., Sharma, K., Gelsomini, M., Giannakos, M., & Abrahamson, D. (2025). Generative AI and multimodal data for educational feedback: Insights from embodied math learning. *British Journal of Educational Technology*, 56(5), 1686-1709.
- Anton, J., Cosentino, G., Gelsomini, M., Sharma, K., Mok, M., Giannakos, M., & Abrahamson, D. (in press). Grounding allocentric procedures in egocentric enactment: A digital design for basic integer arithmetic. *Educational Studies in Mathematics*.
- Giannakos et al., (in press). Programming Task Performance and Perceived Cognitive Engagement Across Three LLM Support Configurations. *Computers in Human Behavior: Artificial Humans*.

This work is a collective effort

WITH THANKS TO

Learner–Computer Interaction Lab
members
Master students and research assistants
Research collaborators and partners

RESEARCH COMMUNITY



FUNDING & SUPPORT



Funded by
the European Union





THANK YOU

Thank you!

Human capabilities at the centre of AI

lci.idi.ntnu.no · lci@idi.ntnu.no

@lcilab

AI
LEARN

Empowerment of
Human Learning

