STIAEM

Slow Tech Innovation

AI, Ethics and New Teaching Methods

From rule-based programming to Machine Learning: a conceptual guide for educators

Co-funded by the European Union

2025-1-IT02-KA220-SCH-000365213

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them... Project reference code: 2025-1-IT02-KA220-SCH-000365213

Università di Foggia
Studi Umanistici, Lettere, Beni Culturali, Scienze della Formazione

STIAEM

Slow Tech Innovation
AI, Ethics and
New Teaching Methods

Co-funded by
the European Union

2025-1-IT02-KA220-SCH-000365213

eduOpen

CC BY-NC-SA 4.0 International

This license requires that reusers give credit to the creator. It allows reusers to distribute, remix, adapt, and build upon the material in any medium or format, for noncommercial purposes only. If others modify or adapt the material, they must license the modified material under identical terms.

Università di Foggia
Studi Umanistici, Lettere, Beni Culturali,
Scienze della Formazione

OER 2.1.1: From rule-based programming to Machine Learning: a conceptual guide for educators
Reference Module: Module 2 - AI Fundamentals and LLM Mechanics
Related Video-Pill: Video 2.1 - Evolution of AI and Machine Learning
Authoring partner: University of Alicante – Spain
Topic / Thematic Focus: Rule-based programming vs Machine Learning, conceptual guide and educator training

Introduction and Learning Objectives

This resource accompanies, together with OER 2.1.2, video-pill 2.1 'Evolution of AI and Machine Learning' and provides its conceptual in-depth analysis. Its purpose is to make comprehensible, even to those without a computer-science background, the historical shift from artificial intelligence based on hand-written rules to systems that learn from data: knowledge that is a necessary premise for any critical use of AI in the classroom.
The text connects to DigComp 3.0 (Cosgrove & Cachia, 2025), the European framework of digital competences, in which competence in AI runs through, explicitly or implicitly, all twenty-one competences, and to the 'pedagogy of AI' pillar of the UNESCO framework for teachers (Miao & Cukurova, 2024), which invites teachers to build the 'pedagogical why' of technological adoption.
By the end of the reading the teacher will be able to:

- distinguish the three major paradigms of AI - hand-written rules, machine learning, deep learning - and recognise their limits;
- explain why tacit knowledge is the bottleneck of the symbolic approach;
- translate the distinction between the machine's probabilistic diagnosis and the teacher's docimological judgement into defensible teaching choices.

1. Symbolic AI: hand-written rules and the wall of tacit knowledge

Let us imagine we want to explain to a computer what a cat is. In the oldest approach - symbolic artificial intelligence - there was only one path: writing rules, one after another. A cat has triangular ears, whiskers, a tail, fur; and so on, for pages of instructions of the 'if... then...' type. A little instruction manual written entirely by hand.
In closed and well-defined domains this approach worked: the expert systems of the 1970s and 1980s made diagnoses on fixed parameters or configured complex machinery with reliable answers. But as soon as the problem became fluid, the machine stalled. The reason is profound and also concerns teaching: much of our competence is tacit. We recognise a student's face among a hundred in an instant, but we could not list the exact rules we used; a teacher immediately grasps whether an argument holds together, whether the register is appropriate, whether there is coherence, but translating that judgement into 'if... then...' rules is almost impossible. Explicit rules are not enough to capture what we know how to do.
The history of AI is also the history of this wall (for a historical overview see Haenlein & Kaplan, 2019). And it was in order to get around it that a second path opened up.

2. The paradigm shift: learning from data

In machine learning the machine no longer executes rules written by a programmer: the relationship between the input data and the output answers is not decided by hand, but estimated from the data themselves. Michael Jordan and Tom Mitchell, in an article published in Science in 2015, define the field as the study of programmes that improve their effectiveness on a specific task through experience, where experience is, concretely, information taken from examples (Jordan & Mitchell, 2015). Improving with experience: that is the difference that changes everything.

Let us return to the cat. Instead of telling the machine 'a cat has whiskers', we show it ten thousand photographs: here there is a cat, here there is not, here there is. And we let it extract the statistical regularities that distinguish a cat from a dog. No one explained to it what a cat is: it inferred it from the examples. The same kind of learning operates in the spam filter, in the face recognition that unlocks a smartphone, in the recommendations of streaming platforms: systems that have learned from millions of examples.

Note, however, that a machine which learns by induction does not possess concepts as we possess them. It builds probabilistic regularities starting from examples and, if the examples are unbalanced or contain errors, it absorbs them. It also learns our prejudices. Talking about data quality is therefore not merely a technical matter: it is a pedagogical and ethical one.

3. Deep learning: representations teach themselves

Until the early 2000s machine learning retained a limitation: we had to indicate which features of the data were relevant. To recognise spam, an expert decided by hand which clues to count - the presence of certain words, the number of links. It was the human being who chose 'what to look at'.

Deep learning overcomes this limitation. Deep artificial neural networks, described by LeCun, Bengio and Hinton in the journal Nature in 2015, learn the representations of the data by themselves, organised on increasingly higher levels of abstraction (LeCun, Bengio & Hinton, 2015). Let us imagine showing a deep network the photograph of a face: the first layers, the shallowest ones, learn to recognise simple elements - edges, contrasts, corners; the intermediate layers combine them into more complex forms - an eye, a nose, a mouth; the deepest layers put everything together, and there is a face. No one programmed this hierarchy: it built itself.

The results came suddenly. In 2012, in a famous image-recognition competition called ImageNet, deep convolutional networks made a sudden leap in accuracy, leaving all previous methods behind; from there the same idea extended to machine translation, voice recognition and text analysis. The point to take home is this: the machine builds an internal representation of the problem, in a space that no one designed by hand. It is not thought: it is a statistical geometry learned from data.

4. Designing Effective Pre-Solved Exemplars

The effectiveness of few-shot prompting depends less on the number of examples than on their quality. Research on demonstrations used in in-context learning indicates that the characteristics and selection of examples can significantly influence model performance

(Min et al., 2022). Teachers should therefore treat exemplars as carefully selected pedagogical resources rather than interchangeable pieces of prompt text.
An effective exemplar should first be clear. It should make the expected relationship between the task and the response understandable. Secondly, it should be relevant to the new task. Examples that are too distant from the intended activity may encourage the AI to identify superficial similarities rather than the desired pattern.
The exemplar should also be accurate and pedagogically legitimate. If an incorrect or misleading example is included in the prompt, the AI may reproduce the same weakness with greater consistency. Finally, the educator should be able to explain why the example was selected and what function it performs. This is particularly relevant to DigCompEdu Area 2: Digital Resources, which requires educators to select, create and manage digital resources in ways that are appropriate to educational objectives (Redecker, 2017).

5. Attention and language models

We now come to the piece that made today's large language models possible. In 2017 a research group led by Ashish Vaswani published 'Attention is all you need', presenting the Transformer architecture. The problem it solved was precise: before the Transformer, models read sentences one word after another, sequentially, and as they went on they tended to forget the initial words, losing the context.
«We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely» (Vaswani et al., 2017, abstract).
The Transformer changes the approach by relying entirely on a mechanism called attention (self-attention): to understand the meaning of a word, the model looks at the same moment at all the other words in the sentence, not only the nearby ones. Let us take an ambiguous word: 'bank'. Is it the bank where money is deposited, or the bank of a river? The model establishes this by relating that word to all the others, in a single stroke. The practical advantage is considerable, because the computation can be parallelised and training becomes faster and more accurate: the numbers confirm it, with 28.4 BLEU points in English-to-German translation and 41.8 in English-to-French. From this architecture descend directly BERT and GPT, the models that you will use in the classroom. And they are still systems that calculate probabilities: they do not 'understand' what they write.

6. What it means for those who teach

At this point the right question is: what changes, for those who teach? What changes is that two levels must be kept clearly distinct. On one side there is the machine's probabilistic diagnosis, that is, what the software suggests; on the other there is the teacher's docimological function, that is, assessment. The two levels never fully coincide and are not interchangeable. Some concrete examples. An automatic proofing tool, trained on standardised texts, tends to flag a colloquial expression, a dialect or an informal register as an error; but if that student is fifteen years old and is recounting a personal experience, that register may be perfectly appropriate. The software applies a statistical average and does not know it. The same applies to adaptive platforms that suggest the 'ideal' learning path: those trajectories arise from statistical correlations over large numbers, they are interpretive hypotheses and not pedagogical truths.

From this follow a practical consequence and a responsibility. The consequence: the tool has a descriptive and orientative function, whereas giving meaning, verifying relevance with respect to the curriculum and adapting to the context of that specific class remain tasks of human judgement. The responsibility: an algorithm trained on distorted data will reproduce, with statistical consistency, the same distortions. Understanding how the machine works is therefore not a theoretical luxury, but the basis for guiding students towards a critical use of the tools and for preventing what we might call the uncritical delegation of thought.

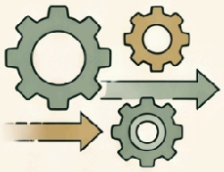
Paradigm	How it works	Everyday example	Consequence for the school
Symbolic AI	Explicit hand-written rules ('if... then...')	Old expert diagnostic systems	Works in closed domains; stalls on tacit knowledge
Machine learning	The machine estimates the relationship between data and answers	Spam filter, recommendations	Biases in the data become biases of the system: an ethical matter
Deep learning	Hierarchical representations learned by deep networks	Image and voice recognition	The machine builds representations, not concepts
Transformer	*Attention* over the whole sequence, parallelisable computation	Chatbots and assistants based on BERT/GPT	It does not 'understand': it calculates probabilities. Judgement remains human

Visual Scaffolding

AI Evolution: From Rules to Probabilities in the Classroom

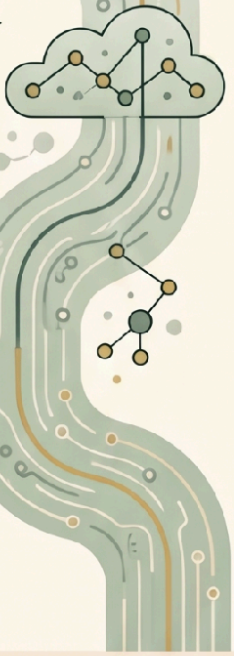
Tracks the evolution of AI from hand-written 'if...then' rules to complex neural networks that learn from data. It emphasizes that while AI has become more powerful, it remains a probabilistic engine that cannot replace a teacher's professional judgment.

The Technological Shift: From Rules to Representation

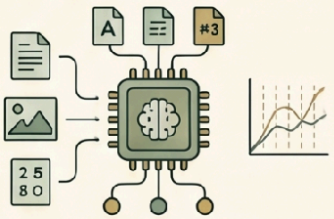


The Symbolic AI Bottleneck
Early AI relied on 'if...then' rules, failing because human 'tacit knowledge' is too complex to code.

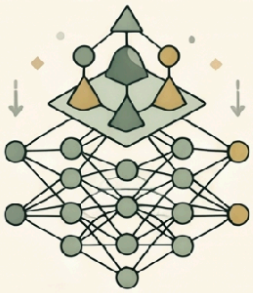




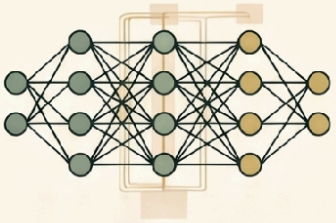
The Machine Learning Revolution



Systems now estimate relationships from data examples rather than following manual human instructions.

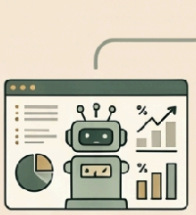


Deep Learning & Neural Networks
Hierarchical layers learn features independently, building statistical geometries instead of human-like concepts.



Transformers & Self-Attention
Modern LLMs use 'self-attention' to process entire text sequences simultaneously as probability engines.

The Pedagogical Reality: Diagnosis vs. Judgment

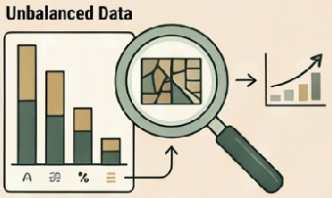


Probabilistic Diagnosis



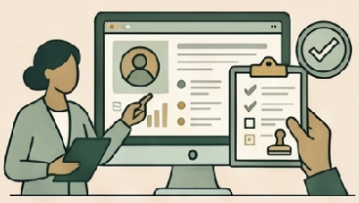
Docimological Judgment

Machines offer statistical averages; teachers provide contextual meaning and irreplaceable professional assessment.



KEY_FINDING: The Ethics of Data Bias
Since machines learn by induction, they absorb and reproduce prejudices found in unbalanced datasets.

Paradigm	How it Works	Consequence for School
Hand-written 'if...then'	Learning from data examples	Self-attention probabilities
Stalls on fluid, tacit human knowledge.	Biased data creates biased systems.	Models do not 'understand'; judgment stays human.



EXAMPLE: Human-in-the-Loop Assessment
Tools serve descriptive functions, but verifying relevance to specific students remains a human responsibility.

Conclusions and Metacognitive Reflection Prompts

This resource has reconstructed the shift from AI based on hand-written rules to systems that learn from data, showing why it is a paradigm shift and not an incremental improvement: the machine no longer executes instructions, but estimates statistical regularities; it does not possess concepts, but builds representations; it does not understand, but calculates probabilities. For the teacher this translates into a professional distinction to be safeguarded, that between the tool's probabilistic diagnosis and human docimological judgement.

To transfer these core ideas into one's own practice, the following reflection prompts are proposed:

- In which tasks of your subject is the competence you assess largely tacit, and how does your way of assessing it change knowing that a machine cannot learn it from explicit rules?
- Which example of bias (in language, in marks, in recommendations) could you bring into the classroom to make the dependence on data visible?
- How would you explain, in three sentences, the difference between what an automatic proofing tool suggests and what a teacher decides?

.

Technical Metadata Sheet

Metadata Field	Specification / Parameter Value
Identifier	OER 2.1.1
Title	From Rule-Based Programming to Machine Learning: A Conceptual Guide for Educators
Creator	University of Alicante - Spain
Description	A conceptual guide for teachers without a computer-science background on the historical shift from hand-written rules to systems that learn from data. It distinguishes symbolic AI, machine learning and deep learning, explaining why tacit knowledge is the bottleneck of the rule-based approach. The didactic focus is on translating the difference between the machine's probabilistic diagnosis and the teacher's docimological judgement into defensible classroom choices.
Keywords	symbolic AI; machine learning; deep learning; tacit knowledge; probabilistic diagnosis; educational assessment; AI literacy; teacher training
Language	EN (English)
Affiliation	University of Alicante - Spain
Contributor	STIAEM Transnational Scientific Board (peer review, Rhodes validation protocol)
Media Format	Text; Infographics
Genre	Training Manual; Educational Guide
URL	
Citation	Cosgrove, J., & Cachia, R. (2025). DigComp 3.0: European digital competence framework (Fifth edition). Publications Office of the European Union. https://data.europa.eu/doi/10.2760/0001149 Haenlein, M., & Kaplan, A. (2019). A brief history of artificial intelligence: On the past, present, and future of artificial intelligence. Business Horizons. https://doi.org/10.1016/j.bushor.2018.08.004

	Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. Science, 349(6245), 255-260. https://doi.org/10.1126/science.aaa8415 LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436-444. https://www.nature.com/articles/nature14539 Miao, F., & Cukurova, M. (2024). AI competency framework for teachers. UNESCO. https://doi.org/10.54675/ZJTE2084 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30. https://arxiv.org/abs/1706.03762
Educational Level	Secondary School
Associated Module / MOOC Video Lesson	Module 2 - AI Fundamentals and LLM Mechanics / Video 2.1 - Evolution of AI and Machine Learning
Learning Objectives	distinguish the three major paradigms of AI - hand-written rules, machine learning, deep learning - and recognise their limits; explain why tacit knowledge is the bottleneck of the symbolic approach; translate the distinction between the machine's probabilistic diagnosis and the teacher's docimological judgement into defensible teaching choices.
Rights	Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0)
Date Available	14/09/2026
Date Modified	

References and Webography

Cosgrove, J., & Cachia, R. (2025). DigComp 3.0: European digital competence framework (Fifth edition). Publications Office of the European Union. https://data.europa.eu/doi/10.2760/0001149

Haenlein, M., & Kaplan, A. (2019). A brief history of artificial intelligence: On the past, present, and future of artificial intelligence. Business Horizons. https://doi.org/10.1016/j.bushor.2018.08.004

Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. Science, 349(6245), 255-260. https://doi.org/10.1126/science.aaa8415

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436-444. https://www.nature.com/articles/nature14539

Miao, F., & Cukurova, M. (2024). AI competency framework for teachers. UNESCO. https://doi.org/10.54675/ZJTE2084

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30. https://arxiv.org/abs/1706.03762.