

Διεθνές Συνέδριο για την Ανοικτή & εξ Αποστάσεως Εκπαίδευση

Τόμ. 13, Αρ. 1 (2026)

ICODL2025



ΠΡΑΚΤΙΚΑ

13ο Διεθνές Συνέδριο για την Ανοικτή & Εξ Αποστάσεως Εκπαίδευση

ISBN: 978-618-5335-27-4

Ανοικτή & Εξ Αποστάσεως Εκπαίδευση:

Οι Δεξιότητες του 21ου Αιώνα & η Πρόκληση της Τεχνητής Νοημοσύνης

ΤΟΜΟΣ 1

5-7/12 2025

ΕΑΠ Πάτρα & Εξ Αποστάσεως



Classroom-Ready AI in the Humanities: An Ethics-by-Design e-Course and a STE(A)M Digital Lab for Teacher Professional Development

Theodora Kouvara, Vassilis Fotopoulos, Theofanis Orphanoudakis, Enikő Nagy, Fernando Antunez Cid, António Manuel Martinho da Silva da Silva, Katerina Maniadaki, Anca Petriuc

doi: [10.12681/icodl.8642](https://doi.org/10.12681/icodl.8642)

Copyright © 2026, Theodora Kouvara, Vassilis Fotopoulos, Theofanis Orphanoudakis, Enikő Nagy, Fernando Antunez Cid, António Manuel Martinho da Silva da Silva, Katerina Maniadaki, Anca Petriuc

Άδεια χρήσης [Creative Commons Attribution-NonCommercial-ShareAlike 4.0](https://creativecommons.org/licenses/by-nc-sa/4.0/).

Classroom-Ready AI in the Humanities: An Ethics-by-Design e-Course and a STE(A)M Digital Lab for Teacher Professional Development

Η Τεχνητή Νοημοσύνη στις Ανθρωπιστικές Επιστήμες, Έτοιμη για την Τάξη:
Ηλεκτρονικό Μάθημα με Ενσωματωμένη Ηθική και Ψηφιακό Εργαστήριο STE(A)M
για την Επιμόρφωση Εκπαιδευτικών

Theodora Kouvara

Hellenic Open University
School of Science and Technology, Greece
tkouvara@eap.gr

Vassilis Fotopoulos

Hellenic Open University
School of Science and Technology, Greece
vfotop1@eap.gr

Theofanis Orphanoudakis

Hellenic Open University
School of Science and Technology, Greece
fanis@eap.gr

Enikő Nagy

Trebag, Hungary
eniko.nagy@trebag.hu

Fernando Antunez Cid

Kaunas Simonas Daukantas Progymnasium,
Lithuania
fernando.antunez@daukantopro.lt

António Manuel Martinho da Silva da Silva
Agrupamento de Escolas José Estêvão, Portugal
f1253@aeje.pt

Katerina Maniadaki

IDEC SA – Business Consulting & Training
k.maniadaki@idec.gr

Anca Petriuc

Colegiul Tehnic Energetic
anca@miraculics.eu

Abstract

European policy increasingly calls for AI-supported, data-informed teaching, yet distance professional development often treats AI literacy, pedagogy and ethics as separate domains. This study presents TENSILE, an Erasmus+ KA2 model for ethics-by-design AI professional development anchored in the humanities. The model integrates Community of Inquiry, TPACK and evidence-based PD principles through a six-module asynchronous Moodle course and a STEAM Digital Lab with curated tools, templates and 25 classroom-ready use cases in language and literature. Activities range from rule-based OULIPO generators to lightweight Python/NLP applications and ML for Kids–Scratch classifiers that make model behaviour observable. Safeguards, including privacy, transparency, provenance, fairness and human oversight, are embedded as routine checkpoints supported by public artefacts such as model cards, data cards and

decision logs. Evaluation combines pilot artefacts with TAICS self-report data to examine competence development and fidelity of enactment. The study contributes a replicable, low-bandwidth model for scalable, ethics-aware AI integration.

Keywords

AI literacy, Community of Inquiry (CoI), TPACK, Ethics-by-design, STE(A)M & humanities, Poetry/NLP use cases

Περίληψη

Η ευρωπαϊκή εκπαιδευτική πολιτική προωθεί τη διδασκαλία με αξιοποίηση της τεχνητής νοημοσύνης και των δεδομένων, ωστόσο η εξ αποστάσεως επαγγελματική ανάπτυξη συχνά αντιμετωπίζει τον γραμματισμό στην ΤΝ, την παιδαγωγική και την ηθική ως διακριτές περιοχές. Η μελέτη παρουσιάζει το TENSILE, ένα μοντέλο Erasmus+ KA2 για επιμόρφωση εκπαιδευτικών με ενσωματωμένη ηθική ήδη από τον σχεδιασμό και με έμφαση στις ανθρωπιστικές επιστήμες. Το μοντέλο συνδυάζει το Community of Inquiry, το TPACK και τεκμηριωμένες αρχές επαγγελματικής ανάπτυξης. Περιλαμβάνει ασύγχρονο ηλεκτρονικό μάθημα έξι ενοτήτων στο Moodle και Ψηφιακό Εργαστήριο STEAM με εργαλεία, πρότυπα και 25 έτοιμες εφαρμογές για τη γλώσσα και τη λογοτεχνία. Οι δραστηριότητες εκτείνονται από γεννήτριες ΟΥΛΙΠΟ έως ελαφριές εφαρμογές Python/NLP και ταξινομητές ML for Kids συνδεδεμένους με το Scratch. Παράλληλα, ιδιωτικότητα, διαφάνεια, ιχνηλασιμότητα, δικαιοσύνη και ανθρώπινη εποπτεία ενσωματώνονται ως σταθερές διαδικασίες, προτείνοντας ένα αναπαραγώγιμο, χαμηλών απαιτήσεων και παιδαγωγικά ουσιαστικό πρότυπο ενσωμάτωσης της ΤΝ.

Λέξεις-κλειδιά

γραμματισμός στην τεχνητή νοημοσύνη, Community of Inquiry (CoI), TPACK, STE(A)M και ανθρωπιστικές επιστήμες, εφαρμογές ποίησης και NLP

Introduction

Across Europe, education systems are converging on data-informed, AI-supported teaching while asking educators to demonstrate advanced digital competence. European policy frames, most notably the Digital Education Action Plan 2021–2027 and DigCompEdu, explicitly call for inclusive, high-quality online professional development (PD) that equips teachers to design and enact digitally rich pedagogy (European Commission, 2020; Redecker, 2017). From a pedagogical standpoint, effective online PD should balance teaching, social and cognitive presence (Community of Inquiry) and embed design levers that support transfer to practice-content focus, active learning, coherence, sufficient duration and collective participation (Garrison, Anderson, & Archer, 2000; Darling-Hammond, Hylar, & Gardner, 2017).

Building on these pedagogical foundations, current educational policy and governance increasingly converge on a human-centred, risk-based model for AI integration, with explicit provisions for teacher capacity building. Three strands are central: (i) teacher-facing guidance (e.g., UNESCO’s Guidance for Generative AI in Education and Research, 2023; European Commission’s Ethical Guidelines on AI and Data in Teaching and Learning for Teachers, 2022); (ii) competency frameworks that translate principles into course-level practice (e.g., UNESCO’s AI Competency Framework for Teachers, 2024); and (iii) binding regulation via the EU AI Act, which enshrines requirements of transparency, oversight, and proportionality in relation to risk levels (European Commission, 2022; UNESCO, 2023). In parallel, OECD curriculum advice outlines what teachers need to teach and what students need to learn in an AI-infused world (OECD, 2025).

Against this backdrop, contemporary PD for AI increasingly focuses on three interconnected domains: AI literacy, techno-pedagogical integration and ethics (Dogan, Nalbantoglu, Çelik, & Ağaçlı Doğan, 2025; OECD, 2023; UNESCO, 2023). AI literacy encompasses knowledge, skills and dispositions that enable individuals to understand how AI systems work, use them effectively and responsibly and recognise when judgment should remain with humans (Allen et al., 2024; Kong et al., 2024; Godwin-Jones, 2025). Techno-pedagogical integration refers to a principled alignment between AI tools and disciplinary aims, supported by transparent orchestration,

inspectable prompts or measures and formative data use for iterative improvement, while safeguarding access and inclusion (OECD, 2023; UNESCO, 2023).

The ethical dimension requires norms, competence and governance protocols that ensure legality, transparency, fairness and educational value. For analytic clarity, we group safeguards into four broad categories: (1) privacy and security (data protection, data minimisation); (2) transparency and provenance (disclosure, explainability, traceability); (3) fairness and non-discrimination (bias detection and mitigation); and (4) human agency and safeguarding (supervision, contestability, academic integrity and IP, accessibility, learner well-being, and age-appropriateness) (European Commission, 2022; OECD, 2023; UNESCO, 2023; OECD, 2025).

However, despite this emerging convergence of policy, ethics and competency frameworks, several gaps persist in how AI-related PD is designed and delivered, particularly in distance formats. First, alignment across AI literacy, pedagogy and ethics often remains partial: these strands are frequently addressed in parallel, through separate modules or resources, rather than as a single, coherent system of practice. Instrumented evaluation frameworks that could support this integration, such as TAICS, are only now stabilising (Chiu, Ahmad, & Çoban, 2025). Second, humanities-anchored exemplars remain underrepresented. Ready-to-use cases in language and literature are limited, even though early work suggests substantial potential for coupling AI with poetic and literary study (Kouvara, Fotopoulos, Karachristos & Orphanoudakis, 2024; Kangasharju et al., 2022). Third, ethical safeguards are unevenly operationalised at the course level: transparency of AI use, data governance routines, and explainability artefacts (e.g., model cards, decision logs) are adopted irregularly, despite strong discursive emphasis in policy (European Commission, 2022; OECD, 2023; UNESCO, 2023). Fourth, infrastructure and mentoring constraints remain significant. Consensus-building studies in contexts such as Greece point to unequal digital provision, time-poor staff and limited specialist support, underscoring the need for modular, low-bandwidth e-learning coupled with sustained pedagogical and technical guidance (Kouvara, Fotopoulos, & Orphanoudakis, 2026, in press).

Against this backdrop, the proposed model of this study is explicitly designed to address the above gaps, functioning as a unified, ethically robust and interdisciplinarily grounded organisation for AI-related teacher professional development. It orients the

pedagogical use of AI to address these systemic shortcomings by treating AI literacy, techno-pedagogical design and ethics-by-design as a single integrated whole rather than three separate tracks. The methodology that generated the model combination of systematic needs analysis and a Delphi study with experts on how an online course can be designed to address these shortcomings has been presented in detail elsewhere (Kouvara et al., 2024; Kouvara, Fotopoulos, & Orphanoudakis, 2026) and its results underpinned the development of the theoretical and design backbone of the present study. The study introduces the TENSILE intervention and the conceptual foundation on which it rests, tracing how principles from Community of Inquiry, TPACK and evidence-based professional development are instantiated in concrete routines, artefacts, and support structures.

More specifically, the study pursues two interlinked objectives:

- to present an online PD methodology that is aligned with established evidence on effective professional learning and with appropriate instructional design models; and
- to demonstrate interdisciplinarity through representative language and literature-based use cases that connect humanities content with data, machine learning and natural language processing.

Operationally, the model is instantiated through a six-module asynchronous Moodle e-course and a STE(A)M Digital Lab that offers selected tools, guides and classroom-ready use cases. Together, these elements map conceptual content to concrete classroom enactment. At the same time, the complementary components of the course (week-by-week briefs, templates and suggested resources per module) are intended to support fidelity of implementation and scalable adaptation across diverse infrastructure and school contexts. In the section that follows, we present the concrete course design and show how the model is translated into structure, learning flows and activities.

The TENSILE Project

TENSILE is an Erasmus+ KA2 distance professional development (PD) programme that prepares in-service secondary teachers to design and implement cross-curricular, AI-mediated instruction with ethics-by-design safeguards. The project combines three

core components: (a) a six-module asynchronous Moodle e-course, (b) a STE(A)M Digital Lab with tools, guides and classroom-ready use cases and (c) school-based piloting activities supported by a European community of practice.

The instructional design model of TENSILE e-course

The design of the e-course builds on three complementary foundations: Technological Pedagogical Content Knowledge (TPACK), which conceptualises effective technology integration as the situated interplay of content, pedagogy, and technology; the Community of Inquiry (CoI) model, which defines quality in online learning through the dynamic balance of teaching, social and cognitive presence (Garrison, Anderson, & Archer, 2000); and a set of evidence-informed PD features-content focus, active learning, coherence, sufficient duration, and collective participation-that are associated with stronger transfer to classroom practice (Darling-Hammond, Hyler, & Gardner, 2017). As illustrated in Figure 1, these foundations are operationalised through a staged weekly cycle comprising initialisation, hands-on artefact work, peer review, formative checking and structured reflection.

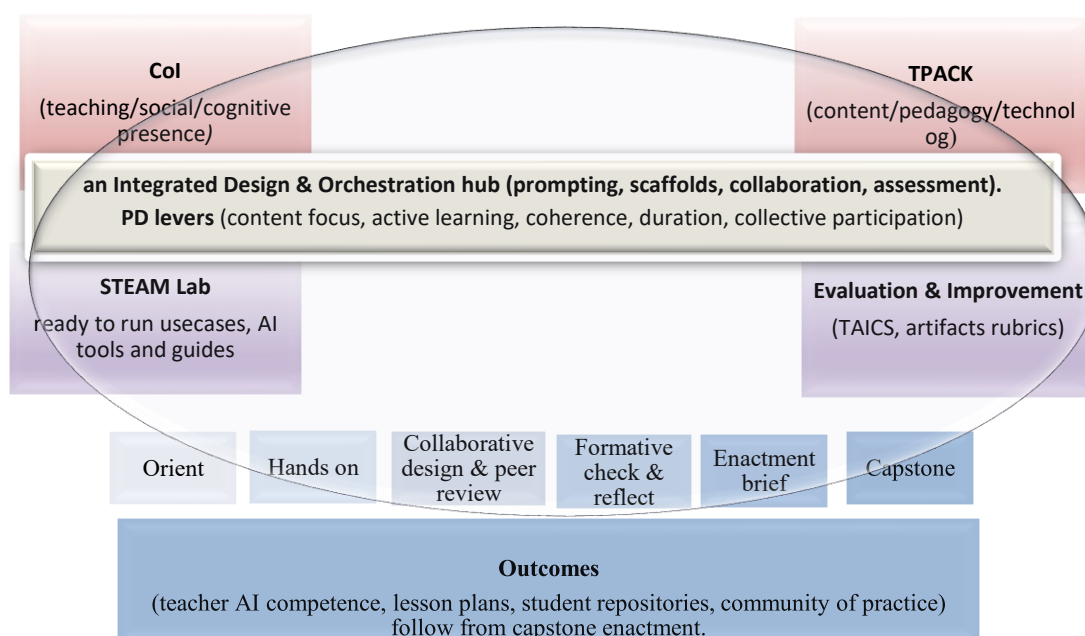


Figure 1: Integrated PD: CoI + TPACK operationalised via PD levers and a staged course cycle

TENSILE adopts a broad interpretation of STE(A)M in which A stands for ALL, explicitly elevating languages and the humanities alongside science, technology, engineering

and mathematics. In this view, interdisciplinarity is anchored in authentic classroom aims rather than added as an aesthetic layer (Perignat & Katz-Buonincontro, 2019). Within this space, poetry functions as a low-threshold, high-ceiling gateway: teachers and students classify poem moods or movements, profile pronouns to surface the “lyrical self”, map motifs, or design visual imagery, while simultaneously engaging with core concepts from natural language processing (NLP) and machine learning (ML).

A central strand of the intervention is constraint-based creativity inspired by the OULIPO tradition, a French literary group that used strict formal constraints to generate new kinds of writing (Motte, 1998; Queneau, 1983). Learners construct “product poems” that follow explicit formal rules, such as snowball verses with increasing word length, Fibonacci-based line patterns, lipograms that exclude a given letter or permutations that reorder lexical items under specified conditions. These activities foreground the algorithmic structure of literary forms and invite explicit reasoning about rules, loops, and conditions.

In parallel, the same or related texts are processed using lightweight computational tools. Python-based NLP pipelines support tasks such as stanza and verse counting, word-frequency visualisation, rhyme-scheme detection (e.g., via phoneme dictionaries), motif analysis, sentiment and mood classification, or identification of nature-related themes. Browser-based environments such as Machine Learning for Kids (Lane, 2021) are used to train simple text and image classifiers, which are then “wired” into Scratch projects (Resnick, Maloney, Monroy-Hernández, Rusk, Eastmond, Brennan, Millner, Rosenbaum, Silver, Silverman & Kafai, 2009), so that ML predictions drive visible behaviours (e.g., a TENSILE Bot that reacts differently to “joyful” versus “somber” lines). Modelling activities include, for example, classifying poems as epic or lyric, distinguishing Romantic from Modernist lyric based on rhyme regularity, or predicting a poem’s mood or likely authorial style (Kangasharju, Ilomäki, Lakkala & Toom, 2022). The juxtaposition of rule-driven poetic generation with ML-based pattern recognition is used deliberately to cultivate AI literacy. Learners compare how deterministic transformations (e.g., N+7 or snowball constraints) differ from models that infer regularities from examples; they analyse model errors, examine feature representations (e.g., n-grams, lexical fields, or phoneme endings) and discuss where rule-based and data-driven approaches succeed or fail. In doing so, they are prompted

to reason about generalisation, bias, uncertainty and the limits of both human and computational interpretations of literary texts.

Digital STE(A)M Lab architecture

The STE(A)M Digital Lab brings together three types of resources: (a) selected tools, including lightweight NLP and ML environments and creative AI services; (b) ready-to-use use cases aligned with literature goals; and (c) short guides and templates that help teachers plan, document and adapt activities. The Lab is organised around a simple workflow (Figure 2): Curated tools → Use cases → Guides and templates → Classroom artefacts. This structure is intended to help teachers move from an initial idea to a tested activity to an enactment brief for their own context, with evidence at each step (e.g., rubrics, model cards, risk notes). Ethics-by-design is built into the Lab from the start. Four practical dimensions are made explicit. Specifically, teachers and learners are clearly informed about when and how AI tools are used and about what data is collected and processed. Short “AI use” sections in each use case and template support this. Secondly, the basic logic of the models and tools (e.g., which inputs they use, what labels they predict and how confidence scores are interpreted) is described in everyday language so that non-experts can understand it. Thirdly, templates include simple fields for documenting data sources, pre-processing steps, model versions and changes across iterations, so that decisions can be revisited and justified. Finally, each use case and enactment brief specifies the teacher's role in monitoring model outputs, deciding when to override or ignore a prediction and ensuring that classroom norms and safeguarding requirements are respected. These dimensions are not treated as abstract principles. However, they are implemented through concrete tools: data and model cards attached to each ML-based activity, risk-reflection prompts embedded in guides and checklists that teachers complete during planning and after running an activity. The Digital Lab is tightly coupled with the online course. In each module, participants explore a subset of tools and use cases, adapt at least one activity using the provided templates and produce artefacts for peer review. The Lab thus serves both as a resource repository (supporting discovery and reuse of ideas) and as a record of practice (storing teacher-generated examples, refinements and reflections). Evaluation procedures for the Lab-supported pilots, including analysis of classroom

artefacts and the use of TAICS to capture AI-related competence and self-efficacy.

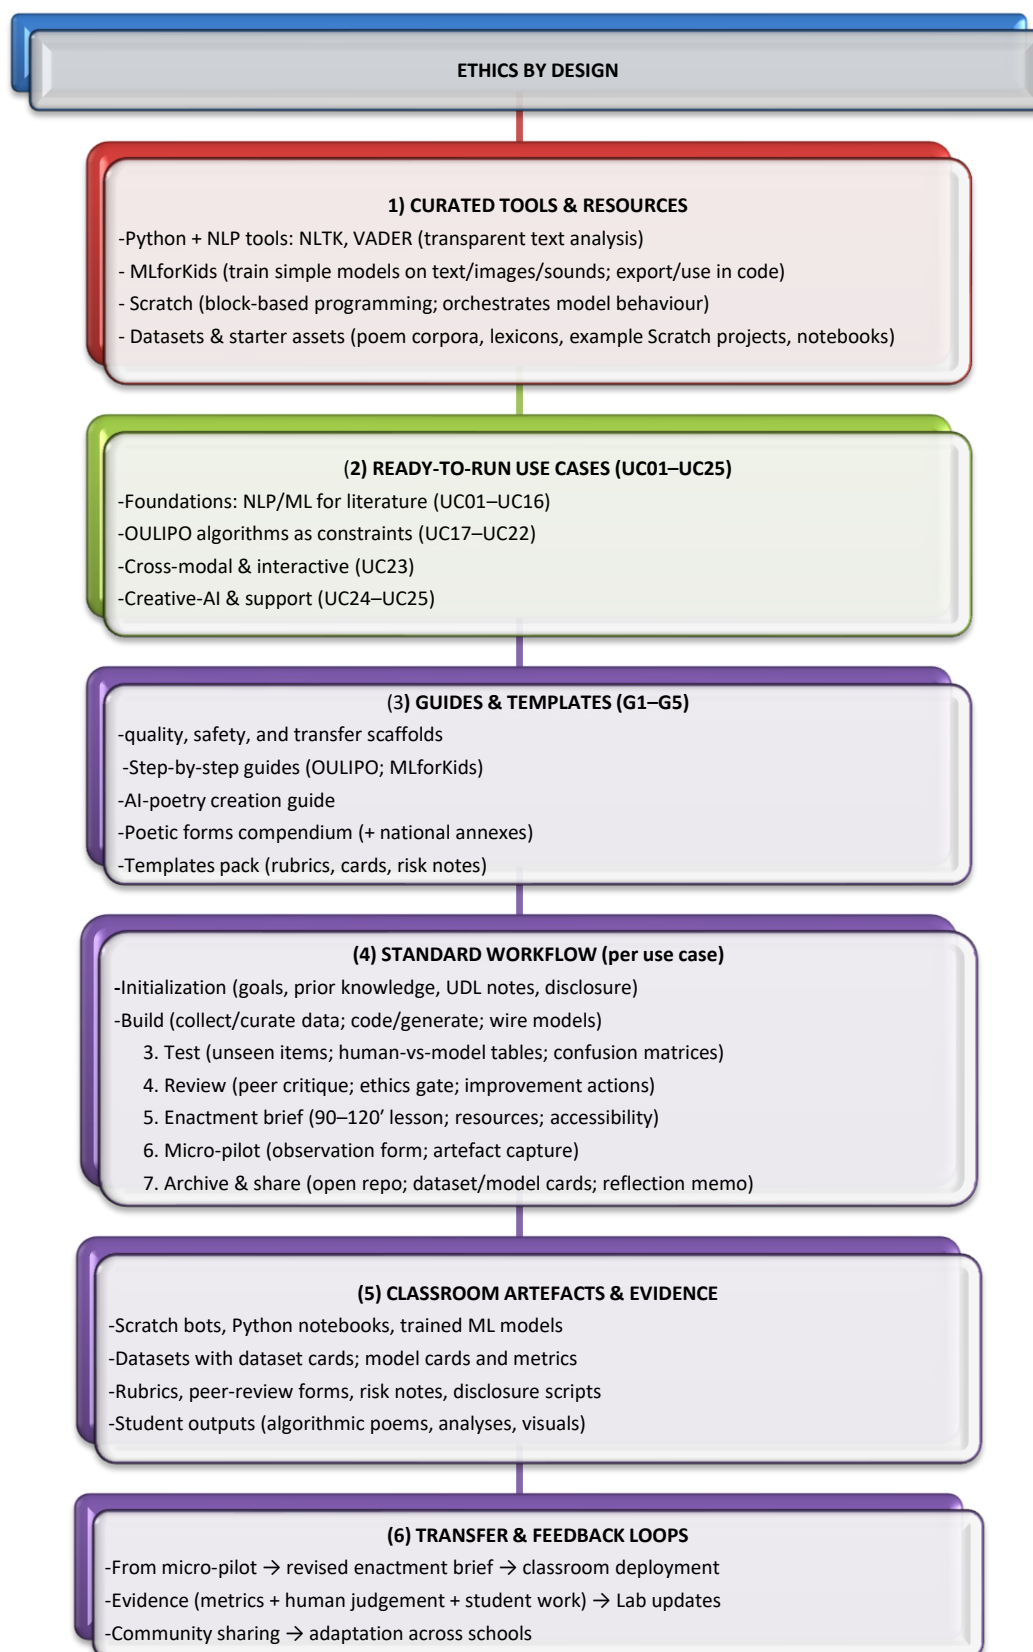


Figure 2: STEAM Digital Lab architecture integrating tools, use cases, workflow under ethics-by-design.

Representative Use-Case Mini-Portfolio

To make the PD architecture concrete (Col + TPACK + the five PD levers under ethics-by-design), the STEAM Digital Lab curates’ classroom-ready, literature-anchored activities that progress from rule-based constraints (OULIPO: French constrained writing group) to lightweight Python/NLP and ML for Kids + Scratch integrations. Each use case below follows the course skeleton (Initialisation → Model/Algorithm construction → Use/Testing → Reflection), embeds guardrails (provenance, disclosure, explainability, human oversight, accessibility) and culminates in visible artefacts that support peer review and classroom transfer. The selection spans: (a) constraint-driven creation (OULIPO); (b) Python/NLP analysis supporting language-arts goals; and (c) MLforKids+Scratch experiences that keep model behaviour observable (e.g., TENSILE Bot). Evidence for the underlying activities comes from the uploaded TENSILE “use cases” compendium and associated files.

Table 1: Mini portfolio at a glance (25 use cases)

UC ID	Title & Core Task	Tools	Typical input → output
UC SA	Text sentiment analysis (EN + non-EN): compute polarity label/scores for a poem.	Python (NLTK VADER; dictionary variant)	poem.txt → compound polarity + label (e.g., positive/negative/neutral).
UC OF	OULIPO_Fibonacci formatting: generate/format lines to follow Fibonacci word or syllable counts.	Python (string ops)	Seed sentence/paragraph → poem lines with 1–1–2–3–5–... structure.
UC RS	Rhyme Schemes (English): detect end-rhyme pattern (A/B/...) using phoneme dictionary.	Python (+ CMUdict-style concept)	poem.txt → scheme string (e.g., ABBA, ABAB) + line-to-label mapping.
UC IR	Image recognition → verses: classify webcam image (cat/dog) and trigger different poem lines/behaviours in Scratch.	MLforKids + Scratch (Video Sensing, optional TTS)	Webcam frame → class label → on-screen sprite says or plays a different verse.
UC RH	Rhymes: suggest rhyming words for a target term to support drafting/analysis.	Python	Target word(s) → list(s) of rhymes; optional stress/phoneme views.
UC SVC	Stanza & Verses Counter: automatically count verses/stanzas for structure checking.	Python	poem.txt → counts per stanza + total lines.
UC WF	Word frequency: compute and visualise term frequencies (with stopword handling).	Python (+ matplotlib)	poem.txt → frequency table + basic plot.
UC PCL	Poems classification_Epic vs Lyric: train a text model and use Scratch UI for predictions.	MLforKids + Scratch	Poem text → class (“epic”/“lyric”) + confidence; Scratch display.

UC IM	Introducing ML (the TENSILE Bot): build a 2–3-label text model; compare rule-based vs ML in Scratch.	MLforKids + Scratch	Short verse → mood label (“happy/sad/indifferent”) → bot reaction (smile/cry/etc.).
UC MV	Mood of Verses (3-label): train and wire a text classifier to a sprite’s emotional behavior.	MLforKids + Scratch	Verse → {Happiness, Sadness, Indifferent} + sprite reaction.
UC MP	The Mood of Poetry (5-label): extend mood labels; test reliability with human–vs–model table.	MLforKids + Scratch	Short poem → {Joyful, Melancholic, Romantic, Somber, Empowering} + comparison table.
UC PSS	Poet Sentiment & Style: train a model to classify short text to poet/style labels and reflect on features.	MLforKids + Scratch	Text lines → predicted poet/style (e.g., Seferis/Elytis/Kavafis) + confidence.
UC TN	Theme of nature in Romantic poetry: detect nature-related keywords and contexts.	Python (NLP)	poem.txt → counts, concordances and a decision on “nature” thematisation.
UC PT	Teaching Past Tense (Simple) through poetry: identify/visualise past-tense forms.	Python/NLP	poem.txt → highlighted verbs + counts/plots; reflective prompts.
UC LS	Lyrical self by Python: profile first-person pronouns & sentiment to infer speaker stance.	Python (NLTK, TextBlob, matplotlib)	poem.txt → pronoun frequency graph + polarity summary.
UC MOT	Using AI tools to analyse motifs in poetry: mine repeated images/lexical fields.	Python (keyword lists/plots)	poem.txt → motif frequency & co-occurrence views.
UC VI	Creating visual images through poetry: design a Scratch scene reacting to verses (imagery devices).	Scratch	Selected poem lines → on-screen visuals/animations/sounds mapped to imagery.
UC ON7	OULIPO_N+7: substitute each noun with the seventh following in a dictionary to generate variants.	Scratch/Python	Source text + lexicon → constraint-generated poem.
UC OL	OULIPO_Lipograms: create a poem forbidding a letter; implement filters/validators.	Scratch/Python	Seed text + forbidden letter → lipogrammatic output.
UC OPAL	OULIPO_Palindromes: check or generate palindromic fragments/words/lines.	Python	String(s) → boolean/palindromic reformulation.
UC OPE	OULIPO_Permutation (with acrostic/anagram options): produce constrained permutations for drafting.	Python (itertools, filters)	Word/line set → unique permutations meeting constraints.
UC OS	OULIPO_Snowball: build lines whose word lengths increase by one each step.	Scratch/Python	Seed lexicon → 1-2-3-4-... letter snowball verse.
UC OFmt	OULIPO_Fibonacci (alt task): <i>formatting-only</i> pipeline for existing text (didactic variant).	Python	Raw paragraph → re-lined output following 1–1–2–3–5–...
UC RM	Romanticism vs Modernism in lyrical texts: classify by rhyme regularity (fixed vs free).	MLforKids + Python	Poem text → movement label (Romantic/Modern) + rationale based on rhyme.
UC TAU	OULIPO_Tautogram: generate or validate alliteration where every word starts with the same letter.	Scratch/Python	Topic + target letter → tautogram line(s) and checker.

EXAMPLE Use Case _ OULIPO Snowball (Rule-based poetry generation)

This activity introduces secondary students (12–18) to constraint-based creativity as a bridge between literary form and computational thinking. Using the OULIPO “snowball” rule—where each line grows (or shrinks) by exactly one letter—students make the governing rule and its algorithmic analogue explicit and inspectable. Pedagogically, the task shows how formal constraints map to loops, lists and conditions and invites comparison with iterative ML workflows elsewhere in the course (e.g., Introducing ML, Mood of Verses). This use case aims to:

- Introduce iterative structures in Scratch/Python;
- Illustrate the analogy between algorithmic refinement and constrained poetic growth;
- Foster creativity within strict formal rules.

In about two hours, learners work in Scratch (Figure 3) with STEAM Lab OULIPO materials to generate snowball poems, producing artefacts aligned with both language and computational learning goals.



Figure 3. Scratch implementation of a snowball generator: initialisation, user input, and loop-based construction of incrementally longer lines.

In specific, the Snowball algorithm takes the input poem as seed text and transforms it by re-arranging or extracting characters and words to fit the rule of incremental growth (1→10 letters). Thus, the output is not a repetition of the input but a constrained reformulation of it. For instance, the following input poem was generated by AI (ChatGPT) for instructional use”:

*“Under clear rules, small sparks learn to rise;
lines gather rhythm as loops repeat;
from form to freedom, patterns take flight”.*

The learners will then transform it using the Snowball constraint during the task by following the rule of one-word lines, 1→10 letters). An exemplar of a potential output is the following:

“i/am/art/form/lines/always/designs/iterates/ascending/illuminate”

This side-by-side view makes the constraint discussable: in Scratch, students iterate through candidate tokens, check length and assemble the 1→N growth. The activity foregrounds control flow and list operations. Students iteratively assemble word sequences and render each line to the required length, allowing growth to be observable at runtime. Moreover, in this Use Case, learners compare constraint-driven generation with iterative improvement in a simple ML workflow: re-label → re-train → re-test, thereby connecting rule-based craft to data-driven refinement and model reliability (see Figure 3). The sequence follows TENSILE’s standard design (Initialisation → Build/Run → Test → Reflect), and can be delivered with Scratch alone or Python alone, or as a comparative pair.

EXAMPLE Use Case _ ML: The Mood of Poetry

Learners train a text classification model with five mood labels-Joyful, Melancholic, Romantic, Sombre and Empowering using a curated database of ~25 poems. They then link the model to Scratch, where the TENSILE Bot asks for a short poem and predicts its mood. The activity follows the official TENSILE packet flow (Initialisation → Model → Dataset → Training → Scratch → Reflection) and ends with a reliability check (human vs. model) and brief reflection. This use case aims to:

- Create and train a simple ML text model in MLforKids;

- Classify poems into five moods and integrate the model with Scratch;
- Apply abstraction and pattern recognition in a creative context;
- Evaluate model reliability and limitations with held-out poems.

Students need basic Scratch skills and familiarity with MLforKids. The lesson lasts ~2 hours and uses the STEAM Lab environment (AI poetry tools, guides, related use cases), the MLforKids platform and Scratch (orchestration layer). The session begins with a short discussion on recognising moods in poetry (keywords, imagery, context). Learners then build the model in MLforKids by creating a new text project and adding five labels_ as illustrated in Figure 4.

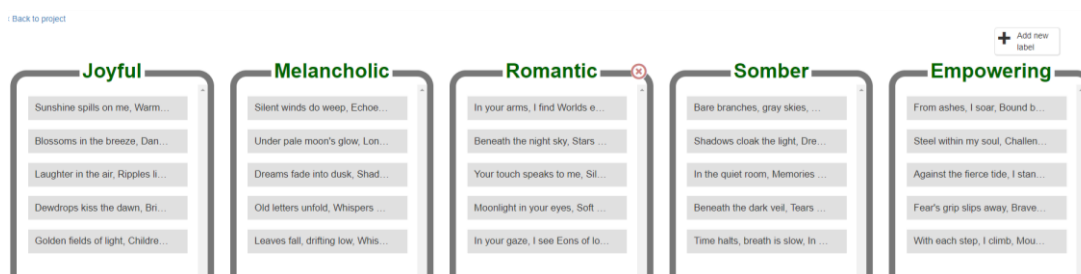


Figure 4. Label setup for Joyful/Melancholic/Romantic/Somber/Empowering before training.

Next, they draw on the curated poetry database of 25 samples, assigning examples to appropriate labels to ensure balance. Once the dataset is ready, the class trains and probes the model by clicking 'Train' and testing sample inputs to verify that the predictions are accurate. At this stage, learners are encouraged to keep notes on dataset balance, edge cases, and any systematic errors. After training, the workflow shifts to connecting the model with Scratch. From the *Make* menu, students open the generated project with ML blocks and program the Bot of Scratch program to ask the user for a short poem, send the text to the ML model, and display the predicted mood, optionally showing confidence values with a threshold (e.g., 50%). The Scratch script and Bot interface appear in Figure 6.

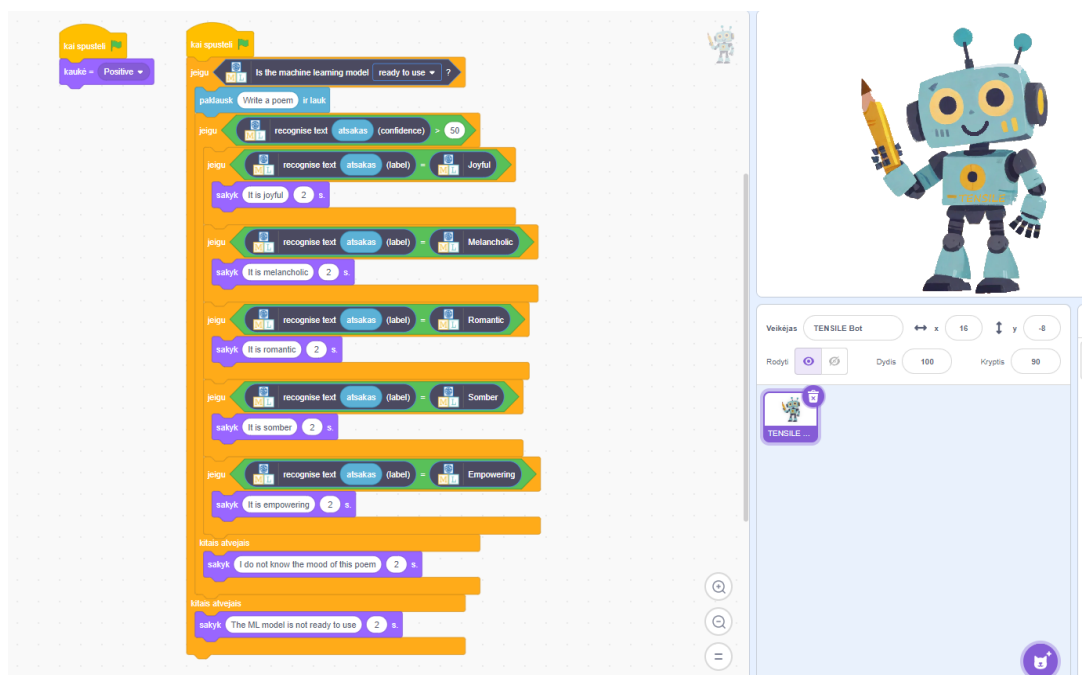


Figure 5: Flow: input poem → ML prediction → Bot says mood; optional confidence threshold.

To evaluate reliability, students complete a model evaluation table by testing five unseen poems and recording their own interpretations alongside the model’s predictions. The comparison template is shown in Table 2. The activity then concludes with a structured reflection, where students discuss when the model becomes confused (e.g., lexical overlap, mixed tone, short inputs), how data quality and label balance affect predictions, what new examples should be added, and what changes might occur through a re-label → re-train → re-test cycle.

Table 2: Students’ judgment vs model prediction on five unseen poems

Number	Students’ interpretation	ML model’s interpretation
1		
2		
3		
4		
5		

The e-learning Course Structure

Building on the integrated Col–TPACK model and the five PD levers outlined above, the Professional Development (PD) programme is implemented as a six-week competence ramp from orientation to classroom enactment and evaluation. Week 1 establishes shared foundations in STE(A)M and AI through engaging activities (including AI-mediated poetry transformations). Week 2 develops core ML understanding by collecting balanced text datasets in MLforKids and wiring models to Scratch, while introducing an initial ethics gate on data minimisation and licensing, and recording baseline TAICS measures. Week 3 foregrounds rule-based creativity (OULIPO) and algorithmic poetry with Python/Scratch, inviting systematic comparison with ML approaches. Week 4 pivots to pedagogy: participants author a TPACK-aligned poetry lesson using simple NLP, incorporating use cases from the STEAM lab and UDL checks, as documented in a concise enactment brief. Week 5 advances to team-based design of STEAM lesson plans and a literature-support chatbot with explicit AI disclosure, validated via cross-team peer review and mentor feedback. Week 6 focuses on effectiveness through co-created quality criteria, analysis of outputs and errors, a final graded project with reflective memo, post-programme TAICS comparison and certification. Ethics-by-design safeguards are embedded throughout (early: provenance, minimisation; later: disclosure, traceability), and each step yields inspectable artefacts that consolidate into classroom-ready materials.

Table 3. Six-Week TENSILE PD sequence: week-by-week focus, representative STEAM-Lab activities, and assessment/evidence

Week	Focus	Core activities	Deliverables	Ethics	Design alignment ¹ (PD · Col · TPACK)
1	Intro to STE(A)M & AI	STEAM Lab tour; AI applied to poetry; brief “rules vs ML” demo; short reflection	Quiz; mini activity; peer review; brief risk note for Poetry generation via AI	Clearly state where/how AI was used	PD: CF, AL, COH, DUR, CP · Col: T (orientation), C (reflection), S (peer) · TPACK: C×P×T (poetry + AI tools; guided inquiry)

¹ **Legend.**

PD levers: CF = Content focus; AL = Active learning; COH = Coherence; DUR = Sufficient duration; CP = Collective participation.

Col: T = Teaching presence; S = Social presence; C = Cognitive presence.

TPACK: C = Content; P = Pedagogy; T = Technology; **C×P×T** denotes integrated alignment.

2	ML basics	MLforKids: build a small balanced dataset; train a simple model; wire to Scratch and display behaviour	ML quiz; data card; Human-vs-Model table; basic metrics; baseline TAICS	Keep only what is needed; record sources & licences	PD: CF, AL, COH, DUR · Col: T (scaffolds), C (model testing) · TPACK: C×P×T (MLforKids ↔ Scratch ↔ poetry)
3	OULIP O & algorithmic poetry	Generators in Python/Scratch; identify when rules suffice vs when ML is needed; team coding	Code + 1-min video; 4-criterion rubric; “algorithm card”	Label what is rule-based vs ML; cite sources	PD: CF, AL, COH, DUR, CP · Col: S (team coding), C (rule–ML comparison), T (facilitation) · TPACK: C×P×T (constraint-based creation with Python/Scratch)
4	Pedagogy in poetry (TPACK K)	Lesson plan with simple NLP and UDL; flow: Init → Hands-on → Peer critique → Reflect	Short quizzes; collaborative critique; 1-page enactment brief	Short paragraph on model limits/errors	PD: CF, AL, COH, DUR, CP · Col: T (design templates), S (peer critique), C (reflection) · TPACK: C×P×T (explicit lesson alignment)
5	Developing use cases	Team design; prototype literature-support chatbot with explicit AI disclosure; mini-pilots	Project pack (code/data/teacher guide); cross-team peer review; mentor memo	AI disclosure; version history; age appropriateness	PD: CF, AL, COH, DUR, CP · Col: S (cross-team review), T (mentoring), C (design decisions) · TPACK: C×P×T (use case + chatbot integration)
6	Evaluation & outcomes	Define quality criteria; analyse outputs/errors; final project; open documentation	Final project (rubric); post-measures and comparison; reflective memo; certification	Licences/CC on files; brief “decision log”	PD: AL, COH, DUR · Col: C (analysis), T (assessment) · TPACK: C×P×T (integration review and iteration)

Discussion

This section explains how the TENSILE model aligns with current policy and research trajectories and discusses its expected contributions, risks and implications for scale. The model addresses the four design gaps identified in the introduction. First, integration across AI literacy, pedagogy, and ethics is treated as a single core course rather than three parallel strands. Community of Inquiry routines structure social, teaching and cognitive presence. TPACK constrains tool–task–content alignment and established PD levers operate as design constraints intended to support transfer. Second, the STEAM Digital Lab contributes humanities-anchored exemplars and poetry-centred activities that range from rule-based OULIPO generators to ML-enabled classifiers, positioning language and literature as primary sites for AI-mediated inquiry rather than peripheral enrichment. Third, ethics is moved from general principles to repeatable procedures through templates (disclosure notes, model/data cards, human-model comparison tables, licensing checks, decision logs)

inserted at explicit “gates” in Weeks 2 and 5 and revisited in the capstone. Fourth, infrastructure and support challenges are addressed through modular, asynchronous learning, low-bandwidth artefacts (printable briefs, short clips, screenshots) and sustained mentoring, aligned with Delphi-elicited national priorities. The model’s mechanism of action couples the production of inspectable artefacts with iterative, socially scaffolded reflection. Teaching presence is front-loaded through explicit course initialisation, worked examples and rapid formative checks. Social presence is developed through team coding, cross-team reviews and sharing of public artefacts. Cognitive presence is promoted via comparative reasoning (rules versus ML; human judgment versus model output) and structured reflection prompts. TPACK alignment is made concrete at the task level, for example, connecting MLforKids models to Scratch to externalise literary decisions as observable behaviours or using Python NLP to render grammatical or motif structures visible, so that a disciplinary goal in language or literature justifies each technical step. The six-week progression supports duration and coherence, while collective participation draws on known benefits of collaborative PD. Taken together, these design moves are expected to increase fidelity of classroom enactment compared with content-only or tool-only designs.

Constraint-based generators (such as Snowball, N+7, Fibonacci, lipograms) make algorithmic structure explicit. At the same time, ML tasks (for example, mood classification, poet/style prediction, Romanticism versus Modernism) reveal how features are inferred and where generalisation fails. Juxtaposing the two enables participants to reason about when simple rules are sufficient, when learned models add value and how uncertainty should be communicated. The resulting artefacts (trained models, Scratch bots, annotated notebooks, curated datasets and transformed poems) are intended to be legible to colleagues and adaptable across languages and grades, supporting the emergence of a community of practice rather than isolated one-off trials. In this model, ethics is treated as a property of activities rather than an optional add-on. Data minimisation and provenance logging are required at dataset construction. Disclosure and traceability are embedded in classroom-facing materials. Performance is documented through model cards and comparison tables. Licensing and accessibility checks form part of the capstone packaging. These routines are aligned with human-centred, risk-based governance.

They are designed to normalise practices that can extend beyond the course (for example, acknowledging uncertainty, documenting changes and enabling learner agency and contestability). The evaluation plan aims to triangulate a validated self-report instrument (TAICS) with artefact-based evidence. In line with prior syntheses on online PD and ML-in-education interventions, the model anticipates gains in three domains: (a) AI literacy (conceptual understanding of model behaviour and limits; prompt and data design); (b) techno-pedagogical competence (designing and facilitating AI-mediated, discipline-aligned lessons); and (c) ethical awareness in action (routine disclosure, provenance tracking, accessibility planning). A central design choice is that outputs take the form of classroom-ready packages rather than decontextualised demonstrations, in order to support transfer during pilots.

Several constraints warrant caution. Literary mood and style labels are inherently subjective; small, curated datasets may encode instructor biases or genre-specific signatures. Human-model comparison tables can surface such issues but cannot fully resolve them. Rhyme and phoneme-based heuristics are language and dialect-dependent; accuracy may degrade for non-English texts or for orthographies with weak grapheme-phoneme correspondence. TAICS primarily captures perceived competence and dispositions; without complementary performance rubrics and classroom observations, inferences about enacted practice remain tentative. Uneven infrastructure and limited technical mentoring, especially in low-bandwidth contexts, may depress participation or restrict more advanced integrations. Finally, creative AI use introduces provenance/IP and academic-integrity risks that require consistent disclosure, process logging and assessment designs that value iterative work and documentation rather than product-only outcomes. For PD designers, the TENSILE model suggests a replicable template: start from disciplinary goals (here, poetry and broader language arts); select AI and digital tools for their pedagogical affordances; build short, artefact-driven cycles with explicit ethics gates; require public, inspectable outputs and scaffold social learning through critique and comparison. For school leaders, modular asynchronous delivery with low-bandwidth options and reusable templates can broaden access while maintaining quality. For policy actors, humanities-anchored exemplars demonstrate that AI capacity building is not confined

to STEM and that risk-based governance can be translated into everyday classroom routines without suppressing creativity.

Subsequent phases should: (a) report pre–post changes with effect sizes and confidence intervals; (b) relate competence gains to artefact quality and fidelity of classroom enactment; (c) examine student outcomes (engagement, disciplinary learning, AI literacy) and equity of participation; (d) conduct fairness and robustness checks on student-built models; (e) extend resources to additional languages and poetic traditions with appropriate phonological tools; and (f) evaluate sustainment six to twelve months after PD completion. Lightweight learning analytics, implemented under privacy-by-design constraints, could support formative mentoring and generate evidence for iterative improvement. Open repositories with versioned artefacts, model cards and decision logs would further enable replication and secondary analysis.

Consequently, the overall study proposes a coherent, ethically grounded approach to distance AI-related professional development by integrating Community of Inquiry, TPACK and evidence-based PD levers into a practice-centred course architecture anchored in the humanities. If confirmed by evaluation, the TENSILE model could offer a credible pathway for scaling inclusive, ethics-aware AI integration across subjects, school levels and systems.

Conclusion

Ultimately, TENSILE offers a practical route for online teacher professional development in AI that treats the humanities as a primary locus for inquiry and creation. By coupling a six-module asynchronous e-course with a STEAM Digital Lab of curated tools and classroom-ready use cases, the design translates principles into repeatable routines that foreground transparency, provenance and human oversight. Rather than emphasising tools, it institutionalises workflow-observable model behaviour, inspectable data practices and disciplined reflection that can travel across languages, grades and infrastructure conditions. As participating schools enact and adapt the materials, the approach is positioned to build capacity at scale while preserving professional agency and ethical guardrails. In doing so, it provides a rigorous, transferable template for integrating AI into day-to-day teaching in ways that

are instructionally meaningful, infrastructure-aware and ethics-by-design from the outset.

References

- Allen, L. K., & Kendeou, P. (2024). ED-AI Lit: An interdisciplinary framework for AI literacy in education. *Policy Insights from the Behavioral and Brain Sciences*, 11(1), 3–10. <https://doi.org/10.1177/23727322231220339>
- Chiu, T.-K. F., Ahmad, M., & Çoban, M. (2025). Development and validation of the Teacher AI Competency Scale (TAICS). *Education and Information Technologies*. Advance online publication. <https://doi.org/10.1007/s10639-024-13094-z>
- Darling-Hammond, L., Hyler, M. E., & Gardner, M. (2017). *Effective teacher professional development*. Learning Policy Institute. <https://files.eric.ed.gov/fulltext/ED606743.pdf>
- Dogan, S., Nalbantoglu, U. Y., Celik, I., & Agacli Dogan, N. (2025). Artificial intelligence professional development: a systematic review of TPACK, designs, and effects for teacher learning. *Professional Development in Education*, 51(3), 519-546.
- European Commission. (2020). *Digital Education Action Plan (2021–2027): Resetting education and training for the digital age*. Publications Office of the European Union. <https://eur-lex.europa.eu/EN/legal-content/summary/digital-education-action-plan-2021-2027.html>
- European Commission. (2022). *Ethical guidelines on the use of artificial intelligence (AI) and data in teaching and learning for educators*. Publications Office of the European Union. <https://op.europa.eu/en/publication-detail/-/publication/d81a0d54-5348-11ed-92ed-01aa75ed71a1>
- Garrison, D. R., Anderson, T., & Archer, W. (2000). Critical inquiry in a text-based environment: Computer conferencing in higher education. *The Internet and Higher Education*, 2(2–3), 87–105. [https://doi.org/10.1016/S1096-7516\(00\)00016-6](https://doi.org/10.1016/S1096-7516(00)00016-6)
- Godwin-Jones, R. *Language teacher preparation for an AI world: A human ecological perspective*. In *Rethinking Language Education in the Age of Generative AI*. Routledge.
- Kangasharju, A., Ilomäki, L., Lakkala, M., & Toom, A. (2022). Lower secondary students' poetry writing with the AI-based poetry machine. *Computers and education: Artificial intelligence*, 3, 100048.
- Kong, S. C., Cheung, M. Y. W., & Tsang, O. (2024). Developing an artificial intelligence literacy framework: Evaluation of a literacy course for senior secondary students using a project-based learning approach. *Computers and Education: Artificial Intelligence*, 6, 100214.
- Kouvara, T., Fotopoulos, V., Karachristos, C., & Orphanoudakis, T. (2024, May). *Expanding the “A” in STEAM: Integrating poetry and AI for educational evolution*. In *2024 IEEE Global Engineering Education Conference (EDUCON)*. IEEE.
- Kouvara, Fotopoulos & Orphanoudakis (2026). From Chalkboards to Chatbots: Preparing Educators for AI Integration. In *Transforming Education with AI: Focus on Content, Analytics, and Inclusivity* (in print). Springer.
- Lane, D. (2021). *Machine learning for kids: A project-based introduction to artificial intelligence*. No Starch Press.
- Motte, W. F., Jr. (Ed.). (1998). *Oulipo: A primer of potential literature* (2nd ed.). Dalkey Archive Press.
- OECD. (2023). *OECD Digital Education Outlook 2023: Towards an effective digital education ecosystem*. OECD Publishing. <https://doi.org/10.1787/c74f03de-en>

- OECD. (2025). *What should educators teach and students learn in a world of powerful AI?* (Education Spotlights No. 20). OECD Publishing. <https://doi.org/10.1787/ca56c7d6-en>
- Perignat, E., & Katz-Buonincontro, J. (2019). STEAM in practice and research: An integrative literature review. *Thinking Skills and Creativity*, 31, 31–43.
- Queneau, R. (1983). *Cent mille milliards de poèmes* (Original work published 1961). Gallimard.
- Redecker, C. (2017). *European framework for the digital competence of educators: DigCompEdu*. Publications Office of the European Union. <https://doi.org/10.2760/178382>
- Resnick, M., Maloney, J., Monroy-Hernández, A., Rusk, N., Eastmond, E., Brennan, K., Millner, A., Rosenbaum, E., Silver, J., Silverman, B., & Kafai, Y. (2009). Scratch: Programming for all. *Communications of the ACM*, 52(11), 60–67. <https://doi.org/10.1145/1592761.1592779>
- UNESCO. (2023). *Guidance for generative AI in education and research*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000386693>
- UNESCO. (2024). *AI competency framework for educators*. UNESCO. <https://www.unesco.org/en/articles/ai-competency-framework-educators>

Terms of Publication, Copyright, and Academic Ethics

This volume contains the papers presented during the Conference. The views and opinions expressed in these proceedings are solely those of the authors and do not necessarily reflect the positions of the Organizing or the Scientific Committee.

Author Responsibility & Copyright: Each author bears full and exclusive responsibility for the content of their paper. Authors guarantee that their manuscripts are products of original scholarly work and that they have secured all necessary written permissions for the use of any third-party copyrighted material (including images, diagrams, or extensive excerpts).

Use of Artificial Intelligence (AI): In accordance with academic integrity standards, authors declare that the use of Generative AI tools, where applicable, was strictly limited to supportive tasks (e.g., language editing, structural organization). The final scholarly judgment, verification of sources, and authenticity of findings remain the sole responsibility of the human authors.

The editors and the Conference organizers shall not be held liable for any copyright infringements or for the academic accuracy and validity of the information provided by the authors.