https://doi.org/10.48417/technolang.2026.03.13
Research article

Artificial Intelligence Literacy in Slavonic Studies: Student Reflections

Tereza Polednikova (✉) , Monika Ricicova, and Pavel Pilch
Masaryk University, Žerotínovo nám. 617/9, 601 77 Brno, Czechia
tereza.polednikova@phil.muni.cz

Abstract

Artificial intelligence literacy is usually assessed through self-report scales. Such instruments show where a group stands, but not how its members reason about the technology in their own terms. Research on this topic has also concentrated on science and technology disciplines and on high-resource languages. It remains unclear how students working in low-resource languages, such as the Slavonic languages, articulate their artificial intelligence literacy, and what knowledge underlies their evaluative judgments. This study analyses written reflections and a background questionnaire from students of Slavonic studies, translation and language teaching who completed a one-semester course in artificial intelligence literacy. The reflections were produced after students used a generative artificial intelligence tool as both the commissioner and the collaborator of an individual academic project. They were examined through qualitative content analysis combining deductive and inductive categories. The findings show that the reflections centre overwhelmingly on evaluating outputs, while declarative knowledge of how such systems function remains a knowledge of symptoms rather than of mechanisms. Evaluative judgments depend heavily on disciplinary expertise in areas such as phonetics, lexicology, translation and language pedagogy, rather than on general critical thinking, and target precisely the errors a domain-uninformed reader would not detect. The questionnaire further shows that students already held a critical disposition toward artificial intelligence before the course began, which constrains how the reflective practice documented here can be interpreted. Taken together, the results reposition disciplinary expertise, more than technical understanding of artificial intelligence, as the primary resource for critical evaluation, with direct implications for teaching artificial intelligence literacy and language technology in higher education.

Keywords: Artificial intelligence literacy; Slavonic studies; Low-resource languages; Disciplinary expertise; Reflective practice; Language technology; Higher education

Citation: Polednikova, T., Ricicova, M. & Pilch, P. (2026). Artificial Intelligence Literacy in Slavonic Studies: Student Reflections. *Technology and Language,* *7*(3), 184-206. https://doi.org/10.48417/technolang.2026.03.13

© Polednikova, T., Ricicova, M. & Pilch, P. This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License

УДК 378.147.016: 004.8
https://doi.org/10.48417/technolang.2026.03.13
Научная статья

Грамотность в области искусственного интеллекта в славяноведении: Размышления студентов

Тереза Поледникова (✉) , Моника Ричицова , и Павел Пилх
Университет Масарика, Жеротиново, наб. 617/9, 601 77 Брно, Чехия
tereza.polednikova@phil.muni.cz

Аннотация
Грамотность в области искусственного интеллекта обычно оценивается с помощью шкал самоотчёта. Подобные инструменты показывают, на каком уровне находится группа, но не то, как её участники рассуждают о технологии на собственном языке. Исследования по этой теме также сосредоточены преимущественно на естественно-научных и технических дисциплинах и на языках с высокими ресурсами. Остаётся неясным, как студенты, работающие с языками с ограниченными ресурсами, такими как славянские языки, формулируют свою грамотность в области искусственного интеллекта и какое знание лежит в основе их оценочных суждений. В настоящем исследовании анализируются письменные рефлексии и анкета исходных данных студентов славяноведения, перевода и преподавания языков, прошедших одностсеместровый курс по грамотности в области искусственного интеллекта. Рефлексии были созданы после того, как студенты использовали инструмент генеративного искусственного интеллекта одновременно как заказчика и как коллегу в рамках индивидуального академического проекта. Они были изучены посредством качественного контент-анализа, сочетающего дедуктивные и индуктивные категории. Результаты показывают, что рефлексии сосредоточены преимущественно на оценке выходных данных, в то время как декларативное знание о том, как функционируют подобные системы, остаётся знанием симптомов, а не механизмов. Оценочные суждения в значительной степени опираются на дисциплинарную экспертизу в таких областях, как фонетика, лексикология, перевод и лингводидактика, а не на общее критическое мышление, и направлены именно на те ошибки, которые не обнаружил бы читатель, не владеющий предметной областью. Анкета также показывает, что студенты уже обладали критической установкой по отношению к искусственному интеллекту до начала курса, что ограничивает интерпретацию документируемой здесь рефлексивной практики. В совокупности результаты переопределяют дисциплинарную экспертизу, а не техническое понимание искусственного интеллекта, как основной ресурс для критической оценки, что имеет непосредственные последствия для преподавания грамотности в области искусственного интеллекта и языковых технологий в высшем образовании.

Ключевые слова: ИИ-Грамотность; Славяноведение; Языки с ограниченными ресурсами; Дисциплинарная экспертиза; Рефлексивная практика; Языковые технологии; Высшее образование

Для цитирования: Polednikova, T., Ricicova, M., Pilch, P. Artificial Intelligence Literacy in Slavonic Studies: Student Reflections // Technology and Language. 2026. № 7(3). P. 184-206. https://doi.org/10.48417/technolang.2026.03.13

Поледникова, Т., Ричицова, М., Пилх, П. This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License

INTRODUCTION[1]

At present, there is a growing need to prepare students to work with artificial intelligence tools in such a way that they become critical and competent users, rather than merely passive consumers of the generated outputs. Research by Ernst & Young, Merriman & Sanz Sáiz (2024, as cited in OECD, 2025) found that the so-called Generation Z learns about AI primarily from social media (55%) and news media (35%), whereas educational institutions are mentioned by only 15% of respondents. At the same time, nearly half of young people aged 17–27 (49%) have difficulty critically evaluating AI outputs (OECD, 2025, p. 7). This gap between informal exposure to AI and the ability to use it critically is particularly urgent in the fields of philology and linguistics where language is not only the subject of investigation but also the primary medium of communication with AI tools. In this context, Slavonic languages represent a specific case, as they fall among the so-called low-resource languages – languages for which available training data is limited compared to English. This issue is also highlighted by the Slavonic Natural Language Processing community, which maintains the view that, despite their importance, Slavonic languages remain significantly underrepresented in NLP tools and data resources compared to English (Piskorski et al., 2025). The asymmetry is measurable in performance as well as in resources: in educational tasks: model accuracy corresponds broadly to a language's representation in the training data, and lower-resource languages such as Czech or Ukrainian are at a demonstrable disadvantage (Gupta et al., 2025). What this means for how students learn to work with such tools is, however, an open question, and it is one that the available evidence on student AI use – drawn overwhelmingly from English-language and STEM settings – cannot answer.

AI literacy is defined as "a set of competencies that enables individuals to critically evaluate AI technologies; communicate and collaborate effectively with AI; and use AI as a tool online, at home, and in the workplace" (Long & Magerko, 2020, p. 2). It belongs to the broader field of digital competence – "a set of skills that everyone needs to live, learn and work in a society where people need to communicate and access relevant information through digital technologies such as internet platforms, social media and mobile devices" (Ng et al., 2023, p. 142) – and is increasingly treated as a qualification in its own right, valued by employers and consequential for students' professional prospects (Ng et al., 2022, p. 26). In the present article we build upon the framework proposed by Ng et al. (2021, p. 4), who define four aspects of AI literacy: (1) Know & understand AI – knowing the basic functions of AI and how to use AI applications; (2) Use & apply AI – applying AI knowledge, concepts and applications in different scenarios; (3) Evaluate & create AI – higher-order thinking skills (e.g. evaluate, appraise, predict, design) with AI applications; and (4) AI ethics – human-centred considerations (e.g. fairness, accountability, transparency, ethics, safety). The framework structures both

[1] During the translation, stylistic editing, and grammatical refinement of this manuscript, the authors utilized the Gemini tool under an institutional university license that guarantees full data protection. All outputs generated by this tool were carefully reviewed and verified by the authors. The authors assume sole and full responsibility for the final content, accuracy, and reliability of the published text.

the design of the course from which these data come and the category system used to analyse the students' reflections.

That course, Artificial Intelligence in the Context of Slavonic Languages, was taught in the spring semester of 2026 at the Department of Slavonic Studies, Faculty of Arts, Masaryk University in Brno, and is described in the Methodology.

Research on AI literacy has developed rapidly, and several frameworks have been developed to portray what it comprises (Ng et al., 2021; Zhang et al., 2025), with a growing body of empirical work examining how students engage with generative AI in higher education (Schei et al., 2024). Nonetheless, there are three gaps in this literature that motivate the present study.

First, the evidence base is unevenly distributed: empirical studies of students' use of AI chatbots are concentrated in STEM disciplines (Schei et al., 2024), and none of the sixteen AI-literacy instruments identified in a recent systematic review has been tested for cross-cultural validity (Lintner, 2024). Philology and language teaching in a lower-resource language environment are therefore doubly under-described – as disciplines and as linguistic settings – and it is not known whether accounts of critical AI use developed in English-language STEM contexts transfer to them. The present study does not test for a resource effect, which would require a comparison condition it does not have; it describes what students in such a setting report doing, which is a precondition for asking the narrower question later.

Second, AI literacy is predominantly assessed through self-reporting instruments: of the sixteen scales reviewed by Lintner (2024), thirteen are self-reporting and only three performance-based. Such instruments establish where a cohort stands, but they cannot show how students reason about the tool in their own terms, which of its limitations they notice, on what grounds they accept or reject what it produces, and how they describe regulating their reliance on it. Reflective writing produced in the course of a real disciplinary task offers access to that reasoning.

Third, existing frameworks treat the evaluation of AI output as a general competence, distinguishing it from knowing how the technology works and from using it (Ng et al., 2021). They say comparatively little about the knowledge students draw on when they judge an output to be wrong. Work on evaluative judgement has established that the capacity to decide about the quality of work is developed rather than given (Tai et al., 2017), and studies of machine translation literacy have shown that language students differ considerably in their ability to identify and correct errors in machine output (Loock & Léchauguette, 2021). The question is what a student must already know in order to see that a fluent, grammatical and confidently presented output is nevertheless wrong. For students of Slavonic languages that knowledge is disciplinary – a judgement about phonetic methodology, about whether a slang expression is still current, about whether vocabulary sits at the CEFR band a lesson claims to target – and the relationship between such expertise and the exercise of AI literacy has not been examined systematically. This is the gap the present study seeks to principally address.

It does so by way of a qualitatively-driven mixed-methods design. It analyses reflections written by students of Slavonic languages, translation and language teaching following a one-semester course on AI literacy, in which they used a generative AI tool

both as the commissioner of an academic project and as a collaborator on it. A questionnaire provides a baseline profile of the cohort's orientation towards AI, against which the reflections are read. The three gaps are not addressed symmetrically. The study is designed principally around the third: the setting supplies a cohort whose disciplinary expertise is linguistic and therefore unusually legible in what they write about AI outputs, and the reflective data make the grounds of their judgements visible in a way self-report instruments cannot. The study therefore addresses three research questions:

RQ_1. What do students of Slavonic languages, translation and language teaching articulate about generative AI in reflections written after an AI-literacy course – their understanding and use of it, their evaluation of its outputs, the ethical issues they raise, their attitudes and trust, and the changes in practice they perceive?

RQ_2. What knowledge do these students draw on when they judge an AI output to be wrong, and how does their disciplinary expertise in Slavonic languages shape the evaluative work they report performing?

RQ_3. In what ways does the cohort's self-reported orientation towards AI at the start of the course inform and constrain the interpretation of these accounts?

METHODOLOGY

Research Design

This study employs a qualitatively-driven mixed-methods design to address the aforementioned three research questions: two concerning what students articulate about generative AI in their reflections and how their disciplinary expertise shapes it (RQ_1–RQ_2), and one concerning how the cohort's self-reported orientation towards AI at entry informs and constrains the interpretation of those accounts (RQ_3). The qualitative strand, which is dominant, consisted of written reflections submitted at the end of the course as a part of students' portfolios and analysed through hybrid deductive-inductive qualitative content analysis. The quantitative strand consisted of a self-report questionnaire, providing a baseline description of the cohort's orientation towards AI.

As the questionnaire was administered anonymously in order to protect students' privacy, and because it preceded the reflections, the two strands were not linked at the level of individual participants. No inferential integration was therefore undertaken, and no claim is made about change over the course. Integration was confined to the interpretive stage, where the cohort's baseline profile was used to contextualise the qualitative findings (Fetters et al., 2013).

Course context

The data were collected in *Artificial Intelligence in the Context of Slavonic Languages*, a course taught in the spring semester of 2026 at the Department of Slavonic Studies, Faculty of Arts, Masaryk University in Brno, and open to all students of the department. It was designed to develop both theoretical knowledge of how AI tools function and practical skills for their use in academic and professional practice, with attention to the specificities of Slavonic languages, and to prepare students for their prospective professional roles as philologists, teachers, lecturers, and translators. In line

with recommendations for integrating technology into higher education, AI tools were used deliberately as a supplement to teaching rather than a replacement for it, and as a means of stimulating critical thinking, creativity and ethical awareness (Luketić & Diković, 2026, p. 72). Unlike frontal instruction, in which broadly identical content reaches all students, the course was expected to adapt to individual needs precisely because of the tools it used, with the teacher acting as a guide while each student worked within their own competence (Șandru, 2025, p. 55). Sessions took place in a computer lab that allowed direct work with AI tools during instruction.

Eighteen students completed the course, representing four Slavonic disciplines – Ukrainian, Russian, South Slavonic, and Balkan studies (Croatian and Bulgarian specialisations) – and including native speakers of Czech, Slovak, Ukrainian, and Russian. This linguistic heterogeneity became a pedagogical resource in itself, since students could compare and contrast AI outputs across different Slavonic languages. Assessment rested on a portfolio, in which students maintained an overview of AI tools, built a library of effective prompts and recorded new insights, and on a final project, described in the following section.

The themes of the individual sessions were classified according to the Ng et al. (2021) framework which builds on Bloom's taxonomy (Table 1). Classification was based on the dominant cognitive level a session demanded – whether its primary focus was acquiring and understanding fundamental concepts, applying knowledge in a specific context, or higher-order thinking such as evaluation, creation and critical reflection – and sessions spanning several categories were mapped to all relevant levels.

Table 1. Classification of course themes according to the AI literacy framework

AI literacy aspect	Course theme
Know & understand AI	What is Artificial Intelligence? Introduction to AI and Language Models (week 1); AI and Slavonic Languages: Specifics and Limitations (week 2); AI-Assisted Translation I: Capabilities and Limitations (week 7)
Use & Apply AI	AI and Slavonic Languages: Specifics and Limitations (week 2); Effective Communication with AI: How to Write Good Prompts (week 3); Using AI for Independent Learning (week 6); AI-Assisted Translation I: Capabilities and Limitations (week 7); AI-Assisted Translation II: Post-Editing (week 9); AI for Academic Reading and Writing: An Assistant, Not a Replacement (week 11); Using AI in Teaching Practice (week 12)
Evaluate & create AI	Critical Thinking and AI: How to Identify Quality Output (week 4); Using AI for Independent Learning (week 6); AI-Assisted Translation II: Post-Editing (week 9); AI for Academic Reading and Writing: An Assistant, Not a Replacement (week 11); Using AI in Teaching Practice (week 12); Project Workshop (week 13); Project Presentations & Course Reflection (week 14)
AI ethics	Risks and Responsibility: The Ethics of AI in Academia and the Workplace (week 5); AI for Academic Reading and Writing: An Assistant, Not a Replacement (week 11); Using AI in Teaching Practice (week 12)

Of the twelve thematic units, three focus primarily on foundational knowledge, seven on practical application, seven on higher-order thinking and three explicitly on ethical considerations. This weighting was deliberate: in the opening session several students said that they preferred practical work to theory, and the syllabus was planned accordingly. Several sessions extend across multiple aspects, reflecting a design in which foundational knowledge, practical skills and critical thinking are revisited and deepened rather than taught as isolated topics. The temporal distribution mirrors the same progression: the early sessions address predominantly foundational knowledge, the middle ones are practically oriented, and the later ones are devoted to analysis and evaluation, culminating in student-led projects.

Participants and data collection

Of the 18 students who completed the course, 14 (78%) gave informed consent for their portfolios to be used for research purposes and are included in the qualitative analysis; 11 (61%) completed the questionnaire and gave consent, and are included in the quantitative analysis. As the questionnaire was anonymous, it is not known which of the 11 are among the 14, and the two samples are treated as overlapping but unlinked subsets of the same cohort.

The reflections were part of the semester project, the objective of which was to use artificial intelligence as an assistive tool in designing and implementing a student project aligned with the student's profile as a future graduate in Slavonic Studies. The primary purpose of the assignment was to reflect on how AI can influence professional practice, improve workflows, facilitate decision-making, and conversely, highlight its potential pitfalls. Students then needed to critically review the proposed project, evaluating its viability, feasibility, and alignment with their own goals and capabilities. Based on the brief generated by the AI, the student attempted to carry out the project. Throughout the process, the students were expected to actively collaborate with the AI by asking follow-up questions, requesting clarification on specific steps, and using it as a consultant to address concrete challenges. The subsequent reflections were to focus on what worked well during the collaboration, the benefits AI provided, and where its limitations became apparent. Students were also asked to address the following questions: What worked well? Where did AI help, and where did it reach its limits? How did you have to verify or correct the AI's outputs? How would you adjust the process next time? What ethical aspects did you need to consider?

Coding procedure

The category system combined a deductive frame with inductively developed additions following the logic of directed content analysis (Hsieh & Shannon, 2005) and of deductive content analysis conducted with an unconstrained categorisation matrix, in which further categories may be created for material the a-priori frame cannot accommodate (Elo & Kyngäs, 2008). Four categories were specified before coding and correspond to the four components of AI literacy identified in the review by Ng et al. (2021): i) Know and Understand AI, ii) Use and Apply AI, iii) Evaluate and Create with

AI, and iv) AI Ethics. This framework is also used to describe the AI literacy content within the course the students completed (see Table 1).

An initial pass over the corpus identified a substantial body of AI-related content that these four categories could not accommodate without distortion: evaluative and affective stances towards AI as such rather than towards a particular output, statements regulating how much epistemic weight AI was granted and temporal statements about the student's practice before, during and after the course. Three further categories – Attitudes and Affect towards AI, Trust and Reliance and Perceived Change and Learning – were therefore developed inductively from the material. These were not specified in advance, but neither were they arrived at without expectation: students' trust in AI, their regulation of reliance on it, and any shift in practice they might report were sensitizing concepts in Blumer's (1954) sense, giving the analysis a general direction of attention without prescribing what would be found (Bowen, 2006). The categories themselves, and the 57 codes within all seven, were derived from the reflections during the coding.

Coding proceeded in two phases. In the first phase, the data were coded independently by two of the authors, working separately and without consultation. The first author applied the codebook; the second coded inductively, generating their own descriptive labels, and only subsequently assigned each label to one of the seven categories. This design was chosen so that the category system could be tested against a reading of the data that had not been shaped by it.

RESULTS

Quantitative Baseline

The questionnaire describes where the cohort stood at entry. Because it was anonymous and cannot be linked to the reflections, it is used here only as context for the qualitative account that follows; no relationship between the two is tested.

Eleven students completed it. Ten were women and one a man; nine were aged 18–25, with one respondent each in the 26–35 and 36–45 bands. Seven were enrolled at bachelor's and four at master's level, across Russian studies (n = 5), Ukrainian studies (n = 4), South Slavonic and Balkan studies (n = 1) and Russian language teaching for secondary schools (n = 1).

Descriptive statistics for the six subscales are given in Table 2. Internal consistency ranged from α = .71 to .93; at n = 11 these estimates are highly imprecise, and 95% confidence intervals are reported alongside each coefficient. Trust sat almost exactly at the scale midpoint, with the cohort dividing evenly around it (six students above 4, five below). Over-reliance and separation anxiety were low overall but markedly dispersed, and in both cases the distribution is better described as a small high-scoring group set against a low-scoring majority than as variation around a central value: three students exceeded the midpoint on each, with partial overlap between the two groups. The cohort therefore entered with a low central tendency but a heterogeneous reliance profile, and the subscale means understate the position of a minority.

Table 2 Descriptive statistics and internal consistency for the AI subscales (n = 11)

Subscale	k	M	SD	Mdn	IQR	Range	α [95% CI]
AI-mediated interaction	4	3.32	0.92	3.25	2.88–3.50	2.00–5.00	.80 [.49, .94]
Over-reliance on AI	6	2.76	1.70	2.17	1.42–3.92	1.17–5.50	.93 [.83, .98]
Reflection on AI feedback	5	4.07	1.13	3.80	3.60–4.10	2.60–6.80	.74 [.38, .92]
Trust in AI	14	4.13	1.02	4.21	3.29–5.00	2.93–5.57	.88 [.73, .96]
Emotional attachment to AI	4	3.59	1.55	3.75	2.25–5.00	1.25–5.50	.71 [.28, .91]
Separation anxiety	4	2.77	1.87	1.75	1.62–3.75	1.00–6.25	.90 [.76, .97]

Note. α confidence intervals computed by Feldt's (1965) method.

Item-level statistics (Appendix A) sharpen the picture in a way that bears on the qualitative findings. The three highest-scoring items in the battery were all evaluative or technique-related: comparing one's own solution with the AI's before deciding (M = 5.00, 73% endorsement), critically appraising AI feedback in order to locate one's own error (M = 4.91, 55%), and refining prompts by supplying context, examples and criteria (M = 4.55, 64%). By contrast, understanding how AI produces its output fell below the midpoint (M = 3.55, 27% endorsement), as did using AI's comments to plan future work (M = 2.82, 18%). The cohort thus arrived already disposed to check AI output and to prompt deliberately, but with little self-reported grasp of the underlying mechanism and little forward transfer from AI feedback into subsequent planning.

Qualitative findings

The fourteen reflections yielded 197 coded segments and 327 segment–code assignments across 57 codes and seven categories. Just over half the segments (106 of 197) carried more than one code. Segments per reflection ranged from 4 to 23 (median 12.5), so prevalence is reported throughout as the number of students in whose reflection a code occurs; segment counts indicate only how much of the corpus a category occupies.

Table 3 Distribution of the coded corpus across the seven categories (N=14 students, 197 segments)

Category	Students	Segments	Most frequently applied codes
Evaluate & create with AI	14	83	Identifying errors or unsuitability (13); verification against external sources (10); adapting output (9); explicit criteria named (8); affirming quality (8)
Use & apply AI	12	51	Structuring and planning (8); supplying context and parameters (7); iterative refinement (7); language work (5); ideation (5)
Know & understand AI	14	40	Recognised capability profile (7); hallucination and fabrication (5); limited domain expertise (5); output quality as a function of prompt quality (4)
Perceived change & learning	12	32	Stated intention for future practice (8); metacognitive insight (6); new awareness attributed to the course (4)
Trust & reliance	11	30	Distrust and the felt need to verify (8); AI as assistant, not replacement (6); deliberate self-limitation (3)
Attitudes & affect	10	24	Overall positive evaluation (6); efficiency and relief (5); ambivalence (5); frustration (4)
AI ethics	11	20	Transparency (3); plagiarism boundary-work (3); human responsibility (3)

As Table 3 shows, evaluative work occupies more of the corpus than the other three components of the Ng et al. (2021) framework, and it is the only category in which every student produced multiple segments. Ethics was raised by eleven of the fourteen students but is represented by only twenty segments and by no code applied to more than three. Knowledge about AI was present in all fourteen reflections, but, as the next section shows, in a form that is almost entirely observational.

i) Know & Understand AI

All fourteen students made claims about what generative AI can and cannot do, but these were overwhelmingly generalisations from what they had just observed rather than statements about how the systems work. Seven described a settled sense of the tasks at which the model is dependable; five reported hallucinated or misattributed content; five located its weakness in specialist domain knowledge, four in sensitivity to the prompt, four in creativity and register, and three in currency and access to sources. Reports of

fabrication are unusually specific, because the students could say not merely that a citation was wrong but why:

> jediný jmenovaný zdroj – Zdena Skoumalová napsala v 60. letech Bibliografický úvod do lingvistického studia ruštiny, praktickým výzkumem se nezabývala [the only source named – Zdena Skoumalová wrote a Bibliographical Introduction to the Linguistic Study of Russian in the 1960s, she did not do empirical research] (Student 2, Czech)

Only one student offered any account of why these limits exist:

> Vďaka teoreticným informáciám som pochopila, že výstupy generované modelmi sú vždy pravdepodobnostné a závislé od kvality zadania (promptu), ako aj od dát, na ktorých bol model trénovaný. V kontexte slovanských jazykov sa tieto obmedzenia prejavujú ešte výraznejšie, najmä pri menších jazykoch. [Thanks to the theoretical information I understood that the outputs generated by models are always probabilistic and dependent on the quality of the prompt, as well as on the data the model was trained on. In the context of Slavonic languages these limitations are even more pronounced, especially with the smaller languages.] (Student 7, Slovak)

This is the only place in the corpus where a student connects an observed failure to a property of the technology, and the only reflection that credits the course's theoretical component for the insight. It is worth being precise about what the segment is: Student 7 is applying a proposition she was taught, not reporting an observation from her own material. The distinction matters for the corpus as a whole. Across all fourteen reflections the phrase "Slavonic languages" occurs twice, both in Student 7's closing paragraph, and no other student anywhere frames a limitation in terms of the resource position of the language they work in. The students detected a great many language-specific failures, as the following sections show, but described them as failures of register, level, currency or scholarly accuracy, not as consequences of how much Czech, Ukrainian or Russian the model had seen. The framing that motivates the course did not, on this evidence, become a framing the students used, yet they went further and were able to describe the specific linguistic features in which the model's output fell short.

One further mechanism-level observation was made once, by Student 10, who had deliberately set out to produce a weak project and noticed the model bending towards his position – a segment discussed below as a deviant case. Against forty segments in this category, these two are the whole of what the cohort says about how AI works. Its declarative knowledge of AI is a knowledge of symptoms.

ii) Use & Apply AI

Twelve students described what they had used AI for, across a wide range of tasks: structuring and planning a project (8 students), ideation (5), language work (5), explanation and self-tutoring (5), and, for one student, routine professional communication outside her studies. Five reported the course task's distinctive move of using the model as the setter of the assignment as well as the collaborator on it:

> AI jsem využívala nejen jako pomocníka při zpracování úkolu, ale také jako „zadavatele" projektu [I used AI not only as an assistant in carrying out the task but also as the "setter" of the project] (Student 8, Czech)

Against this breadth of application, the vocabulary in which students described how they interacted with the model is thin. Seven described supplying context and explicit parameters and seven described iterative refinement; beyond these, task decomposition was named by two students, comparison across tools by two, and cross-tool chaining by one. Most descriptions of technique are retrospective – what the student would do next time rather than what she did. The overlap between this category and Perceived change is the second highest in the corpus (12 segments) and is almost entirely of this kind. Only Student 7 describes prompting as an established, changed practice ("S AI teraz pracujem veľmi systematicky, nie intuitívne" [I now work with AI very systematically, not intuitively]).

iii) Evaluate & Create with AI: the centre of the corpus

Evaluative work dominates the reflections. Thirteen of the fourteen students identified a specific error, weakness or unsuitability in a concrete output; ten described verifying output against an external source; nine described adapting and editing output for their own purposes; and eight named the criterion on which their judgement rested. Criticism is characteristically attached to a particular passage rather than stated as a general reservation:

> Niektoré návrhy síce boli gramaticky správne, ale nepôsobili autenticky alebo nezodpovedali charakteru postavy. [Some of the suggestions were grammatically correct, but they did not feel authentic or did not fit the character.] (Student 6, Slovak)

Verification is overwhelmingly external and manual, and the two most frequent co-occurrences in the entire corpus are between verification and distrust (6 segments) and between verification and hallucination (4): students describe checking because they expect to find something wrong, and report what they found. Positive appraisal was offered by eight students but is heavily concentrated, six of its twenty-one segments coming from a single reflection; where it occurs it is as specific as the criticism.

Two forms of evaluative work were almost absent. Comparison across models was reported by two students, and exactly one student anywhere in the corpus acknowledged a limit on her own capacity to verify:

> Bohužel nemůžu porovnat nakolik je jeho odpověď správná, co se týče děje samotné povídky, jelikož ji neznám a ani od samotné autorky jsem doposud nic nečetla. [Unfortunately I cannot judge how correct its answer is as regards the plot of the story itself, since I do not know it and have not read anything by the author.] (Student 13, Czech)

Given that thirteen students confidently identified errors, the near-absence of this reflexive move is worth noting: evaluative confidence in this cohort is not accompanied by a stated sense of where its own competence to evaluate runs out.

iv) AI Ethics: raised by most, developed by few

Eleven students touched on ethical questions, but the category is the smallest in the corpus and no code within it was applied to more than three students: three addressed transparency and disclosure, three the plagiarism boundary, three human responsibility for the final text, two data protection, two copyright and two bias. Where ethics is raised it is typically raised once, in a summative sentence at the end of the reflection, and the reasoning is largely reactive — students establish that what they did was permissible rather than examine what would not have been:

> Co se týče etických aspektů, příliš jsem nad nimi nepřemýšlela, protože jsem v zásadě neporušila žádná pravidla. […] Nepovažuji to za podvod, protože pouze lehce upravovala formulace. [As for the ethical aspects, I did not think about them much, because in principle I did not break any rules. […] I do not consider it cheating, because it only lightly adjusted the wording.] (Student 1, Czech)

The two substantive exceptions are instructive because neither is general. Student 2 identified a data-protection problem specific to the speech corpus her project proposed to use – registration would require supplying personal data to a foreign state, and the model's assurance that the data were anonymized did not dispel the perceived risk – and Student 7 rejected a topic on the grounds that researching stereotypes might unwittingly reproduce them. In both cases the ethical question arises from disciplinary judgement about a specific research design rather than from principles about AI use in the abstract.

v) Attitudes, Trust and Perceived Change

Ten students expressed an affective evaluation of the experience: six gave an overall positive appraisal, five reported relief at time saved, four reported frustration. More characteristic than either pole is ambivalence, expressed by five students in a single sentence that grants a benefit and immediately qualifies it:

> AI sice dokáže rychle generovat nápady a návrhy, ale nedokáže plně odhadnout, co přesně student očekává nebo jaké má konkrétní možnosti a zkušenosti. [AI can quickly generate ideas and suggestions, but it cannot fully judge what exactly a student expects or what their concrete possibilities and experience are.] (Student 8, Czech)

On trust, eight students articulated a stance of default distrust coupled with an obligation to verify – the most widely shared position in the corpus after error identification. Six framed the tool explicitly as an assistant rather than a substitute and three described deliberately limiting their own use of it to preserve the point of the exercise. Only two named over-reliance as a risk in the abstract, and only one admitted to expedient use. That admission comes from Student 10, whose reflection functions as a deviant case throughout the analysis. Having set out to produce a deliberately weak project, he described the outcome without mitigation:

> Jde pouze o snahu unaveného studenta vytvořit něco na poslední chvíli a o postupné přizpůsobování AI tomu, co se studentovi hodí. [It is simply the effort

> of a tired student to produce something at the last minute, and the gradual adaptation of the AI to what suits the student.] (Student 10, Czech)

Theirs is the only reflection that describes the mechanism by which uncritical use produces a plausible but empty artefact, and the only one in which sycophancy, expedient use and scepticism about AI's place in academic work appear together. It is reported here not as typical but because it makes visible the failure mode the other thirteen reflections describe themselves as avoiding.

On perceived change, twelve students reported some form of learning, but the shape of these reports is constrained by the reflection prompt, which did not ask about prior practice. Eight stated an intention for future practice and six recorded a metacognitive insight about their own process, while only Student 7 contrasted their practice before the course with her practice after it. Four students attributed a new awareness directly to the course. Because before-and-after comparison was not elicited, these are the only evidence of perceived change available, and the design permits no claim about actual change.

Disciplinary expertise as the ground of evaluation (RQ_2)

The evaluative work reported above is not generic critical thinking applied to an AI output. In nine of the fourteen reflections the criticism could not have been formulated without the student's own subject knowledge, and in each case the criterion comes from her specialisation: the design of a sociophonetic study of Russian-accented Czech (Student 2), the currency of Czech youth slang (Student 3), the placement of vocabulary at a CEFR level (Students 5 and 14), the naturalness and rhythm of dubbing dialogue (Student 6), research design in Slavonic sociolinguistics (Student 7), the currency of Russian internet slang (Student 9), Czech–Russian legal terminology (Student 11) and bibliographical accuracy in translation studies (Student 13).

The mechanism is visible in the way these students state their grounds. Student 9 rejected a lexical suggestion on the basis of her own knowledge of contemporary usage and then generalised the point to the volatility of the register:

> výraz „лул" se v současné ruštině podle mého vědomí příliš nepoužívá; spíše se používá jiná varianta. Jelikož se jedná o téma moderního slangu, který se velmi rychle vyvíjí, chyběla mi tam aktuálnost některých výrazů. [the expression "lul" is not much used in contemporary Russian as far as I am aware; another variant is used instead. Since the topic is modern slang, which develops very quickly, I missed the currency of some of the expressions.] (Student 9, Czech)

Student 6 grounded their judgement of a dubbing translation not in a rule but in a triangulation of sources:

> Opierala som sa pritom o vlastnú jazykovú intuiciu, poznatky zo štúdia a v niektorých prípadoch aj o porovnanie s existujúcimi dabingami. [In doing so I relied on my own linguistic intuition, knowledge from my studies, and in some cases also on comparison with existing dubbings.] (Student 6, Slovak)

Two features of this pattern matter. First, the type of error these students detect is one a general-purpose evaluator could not detect: an output that is fluent, grammatical and confidently presented, and wrong only with respect to register, level, currency or scholarly record. Second, their expressed knowledge of the model is not what enabled the detection – only one student connected the errors she found to how the model works, while the other eight identified them because they knew the subject matter and the output did not match it. Disciplinary knowledge and technical understanding of AI thus function here as independent resources, and only the former is well developed across the cohort.

One inference should be resisted. These errors are language-specific but not resource-specific: mis-levelled vocabulary, stale slang, an unnatural line of dubbing dialogue and a wrongly combined bibliographical record would occur in a high-resource language too, and would require the same disciplinary knowledge to detect there. What the data establish is that evaluative competence in this setting rests on disciplinary expertise; whether it rests on it more heavily in a lower-resource language than in English is a question this design, having no

Bringing the baseline and the reflections together (RQ_3)

The questionnaire shows that the cohort did not arrive as uncritical users: the highest-scoring items in the battery were evaluative, while understanding of how AI produces its output fell below the midpoint. The qualitative material is consistent with this profile at both ends and gives it content. The near-universality of error identification and external verification should therefore not be read as an effect of the intervention; the disposition to check was already declared at entry, and what the reflections add is that the cohort was able to act on it, in concrete and domain-specific ways, on outputs in their own field. The finding is one of enacted rather than newly acquired evaluative practice, and the design does not allow the two to be distinguished. At the low end, the convergence between a self-report measure and an unprompted written account on the same gap – mechanism – strengthens the inference that this is a real feature of the cohort's AI literacy rather than an artefact of either instrument.

The two sources diverge in one informative respect. Trust in AI sat at the scale midpoint at entry, yet the dominant stance in the reflections is default distrust coupled with obligatory verification. The most plausible reading is that the reflections were written immediately after a task designed to expose the model's weaknesses on material the students knew well, and that the stance they express is task-specific rather than dispositional. This is a limitation of post-hoc reflection as an instrument, and one reason the baseline is retained as context for the qualitative account rather than as a variable related to it.

DISCUSSION AND CONCLUSION

This cohort presents an academic paradox. Thirteen of the fourteen students identified specific errors in generative AI output, ten verified outputs against external sources, and eight named the criterion on which their judgement rested. Evaluative work filled 83 of the corpus's 197 segments and was the only category to which every student

contributed more than once. Yet across forty segments of claims about what AI is and does, only two concern how the technology produces what it produces, and on the questionnaire the item on understanding how AI generates its output fell below the scale midpoint (M = 3.55, 27% endorsement). High-level evaluative competence co-existed with an almost complete absence of declarative technical knowledge. This section models that phenomenon rather than restating it.

The inadequacy of linear models. Bloom's taxonomy and its revision order cognition so that understanding precedes application and application precedes evaluation (Krathwohl, 2002), and the AI-literacy framework adopted here inherits that ordering: *Know & understand AI* is the foundation on which "Use & apply AI" and then "Evaluate & create with AI" are built (Ng et al., 2021). If the ordering held, evaluative competence would be bounded above by technical understanding. In these data the ceiling is simply absent. The students who evaluated most sharply, the student who dismantled the design of a sociophonetic study, the student who rejected a vocabulary list miss-levelled for B2-C1 learners, the student who found the legal terminology unreliable – were not the students who knew most about language models. Of the nine students whose criticism rested on disciplinary grounds, only one also connected an error to a property of the model. The relationship between the two competencies in this corpus is therefore not weak, but rather orthogonal.

Domain-Anchored Evaluation: a decoupled model. We therefore propose that AI literacy in a disciplinary humanities setting be modelled not as a single ladder but as two independent dimensions, and we set out the model as four propositions.

Proposition 1: Decoupling. Evaluative competence with AI output and declarative knowledge of AI mechanism are separable dimensions. Neither presupposes the other, and a learner may be high on one and near zero on the other. AI literacy is not one construct but two.

Proposition 2: Anchoring. Where mechanism knowledge is absent, evaluation is not criterion-free skepticism. It runs on criteria imported from the discipline. What these students did was not doubt the machine in general but match its output against a standard they already held: register, CEFR band, the currency of a slang form, phonetic method, bibliographical record. Epistemic vigilance is not a content-free disposition; it requires something to be vigilant with (Sperber et al., 2010), and in a disciplinary setting the discipline supplies it. Domain expertise is therefore the operative epistemic defense against flawed algorithmic output.

Proposition 3: Boundedness. Because the criteria are disciplinary, the competence is bounded by the discipline, and the boundary is invisible from the inside. Where a student holds no disciplinary standard, there is nothing for fluent output to fail against, and so no signal that judgement has run out. The model therefore predicts that students will rarely register the limits of their own capacity to evaluate. In this corpus only one student of fourteen did so, and she did it only because she happened to know that she had not read the text in question. Against thirteen students who confidently identified errors, that near-absence is not a marginal result; it is the model's signature.

Proposition 4: Differential visibility. The two dimensions cover different classes of failure, so the model predicts which errors are caught and which pass. Errors with a

disciplinary surface, such as miss-levelled word list, a stale slang form or a fabricated attribution, were caught by thirteen of the fourteen students. Errors with no disciplinary surface, such as accommodation to the user's stated position, silent omission or the probabilistic character of generation itself are detectable only from the mechanism side and were caught almost never. Sycophancy was named by one student, and he noticed it because he had deliberately engineered the conditions that produced it, not because he understood why models behave this way. Technical knowledge in this model is not a prerequisite for evaluation; it is a separate instrument that covers what disciplinary knowledge structurally cannot see. Figure 1 sets out the two dimensions and the four positions they define.

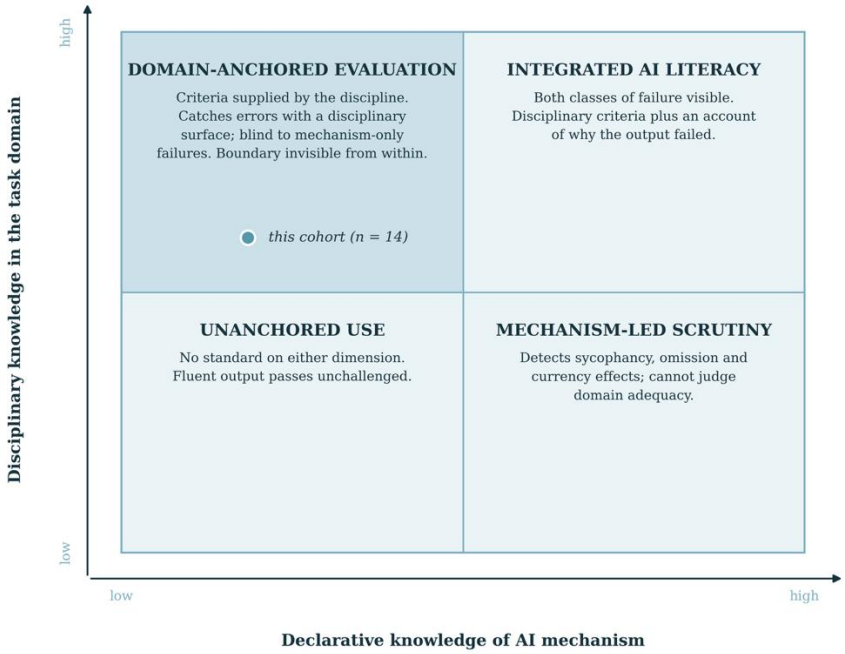

Figure 1. Two dimensions of AI literacy in a disciplinary setting. Evaluative competence with AI output is anchored in disciplinary knowledge (vertical axis); declarative knowledge of how the technology generates output is a separate dimension (horizontal axis)

Why this is a humanities finding. The claim that cognitive skills are substantially context-bound rather than general is long established (Perkins & Salomon, 1989), and the present study extends that finding to algorithmic output and to a discipline in which the object of the study and the medium of interactions are the same. It follows that the humanities curriculum is already, and largely unintentionally, an AI-literacy curriculum: close reading, source criticism, philological accuracy, sensitivity to register and level are the very criteria these students brought to bear. What the course supplied was not the competence but the occasion to exercise it on a new object. This also explains why a technical framing misdescribes what such students can do. Technological and content knowledge are better modelled as intersecting bodies than a nested levels (Mishra &

Koehler, 2006), and the present data suggest that AI literacy requires the same typology: in intersection, not a ladder.

Implications for higher education practice

Four consequences follow, and they are institutional rather than merely didactic. Firstly, architecture. If evaluative criteria are disciplinary, evaluative practice cannot be taught centrally, but it must be embedded in disciplinary tasks, because only there do the criteria exist. If mechanism is what the discipline do not supply, it cannot be left to them either. The implication is a two-track provision – mechanism taught explicitly and generically, evaluation practised locally and assessed within the subject, rather than the single generic service course that institutions typically commission.

Secondly, measurement. Of the sixteen AI literacy instruments identified in a recent systematic review, thirteen are self-report and all are generic (Lintner, 2024). On a single-dimension instrument, a student's score is partly a function of whether the items happen to fall inside their domain, and a cohort like this one, with high evaluative competence, near-zero mechanism knowledge, averages into an unremarkable middle. The paradox reported here is invisible to any instrument that does not separate the two dimensions.

Thirdly, sequence and equity. If the defence is transferred disciplinary knowledge, its strength is a function of disciplinary maturity. Students earliest in their studies, and students working outside their specialisation, are the least protected – precisely the groups at which generic first-year AI literacy provision is usually aimed. An institution that front-loads AI literacy into the first year is relying on a transfer that has not yet happened.

Finally, task design. Proposition 3 predicts that students will not notice where their competence ends, so noticing must be taught rather than assumed. Tasks that deliberately cross a domain boundary, and that require students to state what they could not verify and why, address the one failure mode that disciplinary strength actively conceals.

To conclude, in a disciplinary humanities setting, AI literacy is therefore decoupled from technical literacy, as it operates as the transfer of deep domain knowledge into a new object domain, it is bounded by that knowledge, and it is blind in exactly the places where the domain offers no purchase. Domain expertise, not technical training, is the primary epistemic defence against flawed algorithmic output, and the limits of the discipline are the limits of the defence.

Limitations

The reflections discussed here came from assessed coursework written for a course taught by the first author, who also coded them, so students may have had an incentive to display the critical, verifying stance the analysis reports as dominant; the second coder coded independently, which mitigates this problem but does not eliminate it. Consent was sought after grades were issued. The questionnaire was anonymous and preceded the reflections, so the two strands cannot be linked at participant level and no claim about changes of attitude or opinion is possible. The reflection prompt partly determined which categories are prevalent – most visibly in „perceived change" where prior practice was not elicited. Because disagreements were reconciled in discussion, the analytic dataset is a consensus dataset, reported as negotiated agreement (Garrison et al., 2006). Fourteen of

eighteen students consented to the qualitative strand and eleven to the quantitative, all from one department in a single semester; and without a high-resource comparison condition, no effect of the languages' resource position can be isolated. Finally, the model set out above is generated from this cohort and is offered as a hypothesis rather than a validated framework.

REFERENCES

Blumer, H. (1954). What is Wrong with Social Theory? *American Sociological Review*, *19*(1), 3–10. https://doi.org/10.2307/2088165

Bowen, G. A. (2006). Grounded Theory and Sensitizing Concepts. *International Journal of Qualitative Methods*, *5*(3), 12–23. https://doi.org/10.1177/160940690600500304

Elo, S., & Kyngäs, H. (2008). The qualitative content analysis process. *Journal of advanced nursing*, *62*(1), 107–115. https://doi.org/10.1111/j.1365-2648.2007.04569.x

Feldt, L. S. (1965). The approximate sampling distribution of Kuder–Richardson reliability coefficient twenty. *Psychometrika*, *30*(3), 357–370. https://doi.org/10.1007/BF02289499

Fetters, M. D., Curry, L. A., & Creswell, J. W. (2013). Achieving integration in mixed methods designs—Principles and practices. *Health Services Research*, *48*(6, Pt. 2), 2134–2156. https://doi.org/10.1111/1475-6773.12117

Garrison, D. R., Cleveland-Innes, M., Koole, M., & Kappelman, J. (2006). Revisiting methodological issues in transcript analysis: Negotiated coding and reliability. *The Internet and Higher Education*, *9*(1), 1–8. https://doi.org/10.1016/j.iheduc.2005.11.001

Gupta, V., Pal Chowdhury, S., Zouhar, V., Rooein, D., & Sachan, M. (2025). Multilingual Performance Biases of Large Language Models in Education. ArXiv. https://doi.org/10.48550/arXiv.2504.17720

Hsieh, H. F., & Shannon, S. E. (2005). Three approaches to qualitative content analysis. *Qualitative health research*, *15*(9), 1277–1288. https://doi.org/10.1177/1049732305276687

Krathwohl, D. R. (2002). A revision of Bloom's taxonomy: An overview. *Theory Into Practice*, 41(4), 212–218. https://doi.org/10.1207/s15430421tip4104_2

Lintner, T. (2024). A systematic review of AI literacy scales. *Npj Science of Learning*, *9*(1). https://doi.org/10.1038/s41539-024-00264-4

Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–16). ACM. https://doi.org/10.1145/3313831.3376727

Loock, R., & Léchauguette, S. (2021). Machine translation literacy and undergraduate students in applied languages: Report on an exploratory study. *Tradumàtica: tecnologies de la traducció*, *19*, 204–225. https://doi.org/10.5565/rev/tradumatica.281

Luketić, D., & Diković, M. (2026). Artificial intelligence in higher education teaching: Usage, attitudes and trust among Croatian educators. *Sodobna pedagogika/Journal of Contemporary Educational Studies*, *77*(143), 70–90. https://doi.org/10.63384/spB61z867a

Mishra, P., & Koehler, M. J. (2006). Technological pedagogical content knowledge: A framework for teacher knowledge. *Teachers College Record*, 108(6), 1017–1054. https://doi.org/10.1111/j.1467-9620.2006.00684.x

Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, *2*, 100041. https://doi.org/10.1016/j.caeai.2021.100041

Ng, D. T. K., Leung, J., Su, M., Yim, I., Qiao, M., & Chu, S. (2022). *AI literacy in K-16 classrooms*. Springer. https://doi.org/10.1007/978-3-031-18880-0

Ng, D. T. K., Leung, J. K. L., Su, J., Ng, R. C. W., & Chu, S. K. W. (2023). Teachers' AI digital competencies and twenty-first century skills in the postpandemic world. *Educational Technology Research and Development*, *71*(1), 137–161. https://doi.org/10.1007/s11423-023-10203-6

OECD. (2025). *Empowering learners for the age of AI: An AI literacy framework for primary and secondary education* (Review draft). https://ailiteracyframework.org

Perkins, D. N., & Salomon, G. (1989). Are cognitive skills context-bound? *Educational Researcher*, 18(1), 16–25. https://doi.org/10.3102/0013189X018001016

Piskorski, J., Přibáň, P., Nakov, P., Yangarber, R., & Marcinczuk, M. (Eds.). (2025). *Proceedings of the 10th Workshop on Slavic Natural Language Processing (Slavic NLP 2025)* (pp. 254–275). Association for Computational Linguistics. https://doi.org/10.18653/v1/2025.bsnlp-1

Șandru, I. (2025). Embracing the Future of Artificial Intelligence in the Classroom: Transforming Education through Intelligent Technology. *Euromentor Journal. 16*(3), 52–65. https://euromentor.ucdc.ro/euromentor_september_02102025.pdf

Schei, O. M., Møgelvang, A., & Ludvigsen, K. (2024). Perceptions and use of AI chatbots among students in higher education: A scoping review of empirical studies. *Education Sciences, 14*(8), 922. https://doi.org/10.1016/j.caeai.2024.100259

Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origgi, G., & Wilson, D. (2010). Epistemic vigilance. *Mind & Language*, 25(4), 359–393. https://doi.org/10.1111/j.1468-0017.2010.01394.x

Tai, J., Ajjawi, R., Boud, D., Dawson, P., & Panadero, E. (2017). Developing evaluative judgement: enabling students to make decisions about the quality of work. *Higher Education*, *76*(3), 467–481. https://doi.org/10.1007/s10734-017-0220-3

Zhang, S., Ganapathy Prasad, P., & Schroeder, N. L. (2025). Learning About AI: A Systematic Review of Reviews on AI Literacy. *Journal of Educational Computing Research*, *63*(5), 1292–1322. https://doi.org/10.1177/07356331251342081

Appendix A

#	Item (abbreviated)	M	SD	Mdn	Range	% ≥ 5	*r*_{it}
	AI-mediated interaction (α = .80)						
6	Follow AI's advice when planning work	2.55	1.44	2.0	1–5	18	.87
7	Adjust work according to AI's recommendations	2.64	1.29	2.0	1–5	9	.54
18	Refine prompts (context, examples, criteria)	4.55	0.69	5.0	3–5	64	.61
28	Understand how AI produces feedback	3.55	1.13	3.0	2–5	27	.56
	Over-reliance on AI (α = .93)						
37	Trust AI so much that critical appraisal feels unnecessary	2.64	2.29	1.0	1–7	27	.85
38	Follow AI's suggestions even when unsure they are right	3.18	2.14	2.0	1–6	36	.89
47	Rarely adapt an AI-provided plan to own preferences	2.64	1.63	2.0	1–6	18	.84
49	Rely on AI's recommendation over own judgement	2.55	1.29	2.0	1–4	0	.77
57	Barely check AI output before using it	2.18	1.40	2.0	1–6	9	.69
67	Choose AI's suggestion over own idea	3.36	2.73	2.0	1–7	36	.90
	Reflection on AI feedback (α = .74)						
39	Compare own solution with AI's before deciding	5.00	1.26	5.0	3–7	73	.55
40	Consider how using AI affects own learning	4.45	2.07	4.0	2–7	45	.48
46	Use AI's comments to plan future similar tasks	2.82	2.18	2.0	1–7	18	.66
54	Critically appraise AI feedback to find own error	4.91	1.04	5.0	4–7	55	.38
65	Revisit own plan after receiving AI feedback	3.18	1.17	3.0	2–6	9	.61
	Trust in AI (α = .88)						
35	Know what to expect from AI	4.64	1.36	4.0	3–7	36	−.08
36	AI acts in my interest when I learn	5.00	1.34	5.0	3–7	55	.53
41	AI behaves consistently on similar tasks	4.18	1.94	4.0	1–7	36	.84
42	AI is open about how it arrives at a solution	4.18	1.54	4.0	2–7	27	.55

44	Trust AI even without understanding how it works	4.18	2.40	4.0	1–7	45	.80
48	AI is capable of completing set tasks	4.73	1.35	5.0	2–6	64	.71
51	Willing to use AI again after faulty feedback	4.82	1.66	5.0	1–7	55	.81
53	The way AI works makes sense and is predictable	4.64	1.75	5.0	2–7	55	.76
56	Feel confident letting AI decide for me	2.64	1.86	2.0	1–6	27	.93
58	AI always works in the same way	2.73	1.49	3.0	1–5	9	.01
59	AI acts honestly and ethically	3.64	1.43	4.0	1–5	36	.23
60	Even if AI errs, it can correct itself and remain useful	3.82	1.60	4.0	1–6	36	.33
61	AI provides accurate and correct feedback	3.91	1.58	4.0	1–6	36	.63
62	AI explains its suggestions comprehensibly	4.73	1.49	5.0	2–7	55	.44
	Emotional attachment to AI (α = .71)						
43	Interacting with AI improves my mood	4.27	2.49	4.0	1–7	45	.77
50	Feel better when I receive feedback from AI	4.36	1.57	5.0	1–6	55	.21
55	Feel emotionally attached to the AI tool I use	3.09	2.43	2.0	1–7	36	.47
63	Miss its guidance when I stop using AI	2.64	1.86	2.0	1–6	27	.60
	Separation anxiety (α = .90)						
45	Postpone work until AI is available again	2.91	2.02	2.0	1–7	18	.79
52	Feel anxious when I need something but cannot use AI	2.64	1.75	2.0	1–7	9	.54
64	Worry about my progress when AI is unavailable	2.73	2.37	2.0	1–7	27	.91
66	Feel uneasy without AI at my disposal	2.82	2.27	2.0	1–7	27	.94

СВЕДЕНИЯ ОБ АВТОРАХ / THE AUTHORS

Тереза Поледникова,
tereza.polednikova@phil.muni.cz
ORCID 0009-0009-8231-8904

Моника Ричицова, ricicova@phil.muni.cz
ORCID 0009-0002-2379-8062

Павел Пилх, pilch@phil.muni.cz
ORCID 0000-0002-3432-5298

Tereza Polednikova,
tereza.polednikova@phil.muni.cz
ORCID 0009-0009-8231-8904

Monika Ricicova, ricicova@phil.muni.cz
ORCID 0009-0002-2379-8062

Pavel Pilch, pilch@phil.muni.cz
ORCID 0000-0002-3432-5298

Статья поступила 5 июля 2026
одобрена после рецензирования 14 августа 2026
принята к публикации 21 сентября 2026

Received: 5 July 2026
Revised: 14 August 2026
Accepted: 21 September 2026