

Which Skills Debate Reaches the Public? Comparing Scientific Literature and Media Coverage of AI and LLM Skill Impacts (2022–2025)

Maud Reveilhac¹, Gerold Schneider², Bohdan Trembovelskyi³
Nisha Yadav¹, Vlada Druta³, Aurelia Tamò-Larrieux³, Simon Mayer⁴
Clement Guitton⁴, Joshua C. Yang⁵

¹LUT University (Finland); ²University of Zurich (Switzerland); ³University of Lausanne (Switzerland);

⁴University of St-Gallen (Switzerland); ⁵ETH Zurich (Switzerland)

Correspondence: maud.reveilhac@lut.fi

Abstract

As large language models (LLMs) rapidly enter classrooms and workplaces, the public discourse about their impact on human skills has significant consequences for public understanding, policy, and education. We analyze this discourse through a comparison of the scientific literature and Swiss multilingual media coverage in German and French on how LLMs affect skills in education and work from 2022 to 2025. Using a reproducible pipeline that combines conceptual mapping and BERTopic modeling and proposing and implementing a tripartite analytical framework – amplified, simplified, absent – for examining how scientific framings are recontextualized in public discourse, we find a sharp divergence in thematic structure. We use these three categories to show that media coverage amplifies concrete use cases, simplifies differentiated scientific debates about skills, and leaves several central scientific themes less visible. In education reviews, a single pedagogical topic accounts for the majority of the corpus; in workplace reviews, it remains dominant at 51.28%. In Swiss media, however, the dominant topic is a broad and generalized AI-skills discourse (52.99%), while the education-centered topic accounts for only 1.52%. Conceptual maps show in greater detail that media coverage foregrounds AI capabilities, job loss, and replacement, while giving limited attention to themes central in the literature, including AI literacy, reflective use, metacognition, and pedagogical integration.

1 Introduction

Since the public release of ChatGPT in late 2022, LLMs have become central to debates about education, work, and the future of human competence. In education, recent reviews describe a rapidly expanding research field that links LLM use to writing, feedback, assessment, AI literacy, critical thinking, and broader pedagogical change (Kasneji et al., 2023; Yan et al., 2024; Giannakos et al.,

2025). In workplace research, LLMs are discussed both as productivity-enhancing tools and as drivers of changing task structures, skill requirements, and job quality (Brynjolfsson et al., 2025; Gmyrek et al., 2023; Salari et al., 2025).

Yet, scientific knowledge does not enter the public sphere directly. For most citizens, emerging technologies are encountered through media discourse rather than through review articles or research syntheses. Media therefore matters not as a measure of public opinion itself, but as a site where issues are framed, simplified, dramatized, or normalized, thereby shaping the interpretive resources available to citizens when they form opinions about emerging technologies. Research on public attitudes toward AI shows that media use is associated with the kinds of interpretive frames people adopt when thinking about AI, while deliberative research suggests that public support for AI-related applications depends strongly on how issues are explained and contextualized (Brewer et al., 2022; Arnesen et al., 2025).

This matters especially for the notion of skills. In the scientific literature, “skills” is not a single, stable category. It can refer to cognitive skills, professional competences, (AI) literacy, metacognitive regulation, communication abilities, or new forms of human–AI collaboration (Acemoglu and Autor, 2011; Zhang et al., 2025; Sidra and Mason, 2024). In the public discourse, however, these distinctions may be compressed into narratives about replacement, efficiency, or the need to “adapt”.

We thus ask: which dimensions of skill transformation associated with LLMs in scientific literature reviews are amplified, simplified, or absent in the discourse on public media? To answer, we compare two corpora from the same period, 2022–2025: (1) literature reviews in English on the effects of LLMs on skills in education and workplace domains, and (2) Swiss German- and French-language media articles retrieved from Swissdox.

This paper makes four contributions. *Empirically*, we provide a comparison of scientific and media discourse on LLMs and skills in a multilingual Swiss context, drawing on 246 English-language review documents and 4,610 German- and French-language news passages from Swissdox covering the 2022–2025 period. *Methodologically*, we develop a reproducible pipeline that combines BERTopic modeling with conceptual mapping via a translation-based pivot, enabling cross-arena and cross-lingual comparison of how the same problem space is structured in different discursive settings. *Conceptually*, we propose a tripartite analytical framework - *amplified, simplified, absent* - for examining how scientific framings of skill transformation are recontextualized in public discourse. *Theoretically*, we argue that skill transformation due to AI should be treated not only as an educational or labor-market issue, but as a public-opinion issue.

2 Theoretical background

2.1 Skills, tasks, and technological change

To discuss expected skill transformation through LLMs and their use, we start by introducing a distinction between tasks, skills, and technologies. Acemoglu and Autor (2011) argue that technological change does not mainly replace occupations; rather it reorganizes task structures and alters the value of particular skills. This perspective remains relevant for LLMs, which intervene directly in language-based and knowledge-intensive tasks.

Recent research on generative AI extends this view in two directions. First, workplace studies emphasize that LLMs might augment performance unevenly, often benefiting less experienced workers while also reshaping expectations of expertise, autonomy, and evaluation (Brynjolfsson et al., 2025; Woodruff et al., 2024). Second, education research stresses that the relevant outcome is not simply performance improvement, but the transformation of competences such as writing, reflection, prompting, critical thinking, and digital literacy (Kasneci et al., 2023; Yan et al., 2024; Daniel et al., 2025).

This literature also shows that LLM-related skill change is normatively ambivalent. Some studies highlight productivity gains, scaffolding, personalization, and support for higher-order learning, while others warn against overreliance, cognitive offloading, shallow engagement, and unequal access to AI competence (Giannakos et al., 2025;

Wang and Fan, 2025; Wieczorek et al., 2025).

2.2 Framing, mediation, and public understanding

In this paper, we draw on a science-communication view of media as a site of recontextualization rather than mere transmission. Scientific reviews synthesize evidence and stabilize conceptual distinctions; journalism selects, condenses, and narrativizes. As a result, the issue is not whether media “covers” the same material as the literature in a one-to-one way, but how scientific problem definitions are transformed into public-facing frames.

Brewer et al. (2022) show that public attitudes toward AI are associated with different “frames in mind,” including optimistic and risk-oriented interpretations, and that news use predicts these framings. Arnesen et al. (2025) similarly show that expressed AI support is sensitive to deliberation and knowledge conditions, which underscores the importance of information quality for democratic opinion formation.

In the specific case of LLMs, one likely mechanism is product-centric framing: public discourse often treats a highly visible application such as ChatGPT as shorthand for AI more generally. From the standpoint of skills discourse, this matters because product-centered coverage may foreground spectacle, capability, and disruption while backgrounding slower-moving questions about pedagogy, competence formation, metacognition, or professional practice. Deeper scientific grounding of relevant public discourse is further hindered by polarization tendencies of several public media, specifically on social media platforms (cf. (Lorenz-Spreen et al., 2023; Nguyen and Hekman, 2024)).

2.3 Skills as a public issue

This leads to this paper’s core theoretical positioning: We argue that skill transformation should not be treated only as educational or labor-market issue, but as a public-opinion one. If a framed media discourse selectively highlights some skill dimensions and neglects others, the public debate around LLMs may be structured by partial visibility. For example, “job loss” may become publicly salient while “AI literacy” or “reflective human–AI collaboration” remains largely scientific vocabulary. The literature on AI literacy already suggests that the concept itself is heterogeneous and evolving, spanning technical understanding, evaluation, ethics, and metacognitive awareness (Zhang et al., 2025;

Sidra and Mason, 2024).

This positioning has two implications for our study. *Theoretically*, it reframes skill transformation as an object of public-opinion research rather than only an educational or labor-market concern: what counts as a “skill at risk” or a “skill to acquire” is partly constituted by the discursive frames through which the issue becomes publicly visible. *Empirically*, it motivates a comparative design in which the same problem space – the intersection of LLMs and skills – is examined across two arenas (scientific synthesis and public-facing media) in order to identify which dimensions travel, which are simplified, and which remain largely invisible.

3 Methods

3.1 Corpus design

The study compares two corpora covering the same time period (2022–2025): a corpus of scientific literature reviews in English and a corpus of bilingual Swiss news media articles (French and German). The comparison aims to analyze how the concept of skill transformation due to LLMs is represented across different discourse arenas.

The first corpus consists of English literature reviews on LLMs and skills in two domains: education and the workplace. Literature reviews were selected because they synthesize multiple empirical studies and therefore provide a condensed representation of the scientific debate. As such, they serve as a useful reference space for identifying the main dimensions through which LLM-related skill transformation is conceptualized in research. The literature corpus was collected using two domain-specific queries combining AI and skill. The exact query is in the Appendix A1. The query does not search only for ChatGPT or GPT, but includes broader generic terms for AI, generative AI, and large language models in English, German, and French. Nevertheless, because the terms GPT/ChatGPT are highly visible in public discourse, their inclusion may increase the visibility of OpenAI-related coverage.

The second corpus consists of Swiss news media articles retrieved from the Swissdox database.¹ The retrieval logic was aligned conceptually with the literature corpus by combining LLM-related and skill-related terms. The query includes English, German, and French variants to capture multilingual coverage, see the exact query in the Appendix

¹See <https://swissdox.ch/> (last accessed: 30.04.2026)

A1. In the Swissdox query language, this logic was implemented using token and sequence matching for multilingual expressions of both AI technologies and skill-related concepts. The extraction initially returned 6062 German and 3478 French articles. To focus on relevant passages only, and also because BERTopics has limited context, we chunk the articles into passages as follows: After every 200 tokens we split at the next full stop. The 200-token threshold was chosen as a pragmatic compromise: it preserves enough local context for semantic embedding while avoiding very long passages that may contain several unrelated article sections. This procedure does not fully preserve journalistic structure such as headlines, leads, quotations, or section boundaries. We therefore interpret the resulting units as local thematic passages rather than complete articles. Only news passages that corresponded to the search query were kept. This led to 2915 German and 1695 French, in total 4610 passages from Swissdox. This design does not eliminate genre differences, but it ensures that both datasets are sampled from the same semantic problem space: the intersection of AI and skills. The English literature reviews’ titles and abstracts are shorter, so we did not need to chunk them. The literature corpus contains 78 work- and 168 education-related documents.

3.2 Analytical logic

The analytical logic is asymmetric. The literature-review corpus is not treated as a gold standard in any normative sense, but as a structured representation of the scientific conceptual field. The media corpus is treated as the public mediation field in which some of these dimensions may be amplified, generalized, displaced, or omitted. This means the study is not asking whether the media reproduces the literature faithfully. Instead, it asks how the skill dimensions present in scientific discourse are recontextualized when they become part of public communication.

3.3 Pre-processing and representation

Swissdox XML files were converted into structured text data through a Python–R pipeline. Media texts in German and French were normalized for multilingual comparison with the English-language review corpus. We used the multilingual model paraphrase-multilingual-mpnet-base-v2 for

Topic Modeling (Reimers and Gurevych, 2019)². Because the analysis focuses on conceptual proximity rather than quotation-level interpretation, semantic text representations are especially important. Embedding-based approaches are appropriate here because they are able to capture latent patterns or semantic similarity beyond simple word overlap, which is crucial when comparing journalistic and scientific language that may refer to similar issues with different lexical choices (Reimers and Gurevych, 2019). We have also used a translated version (see below) and obtained similar Topic Modeling results.

3.4 Conceptual mapping

From this structured text data, we constructed a visual semantic map by extracting salient terms from the corpora and positioning them in two-dimensional space according to their co-occurrence proximity (i.e., terms that frequently appear together are also rendered closer to one another). The result is a graph in which clusters correspond to conceptual neighborhoods, and the distance between nodes reflects the degree of lexical association. This enables an exploratory view of which concepts are corpus-specific, which are shared, and which form bridging zones between scientific and media discourse. This step follows a broader text-as-data approach in which texts are analyzed not only as containers of content but as structured signals of social and communicative processes (Stewart et al., 2022). It is especially useful here because the goal is to identify dimensions of framing rather than just keyword frequencies.

As this method does not support multilingual embeddings, we translated the passages automatically to English using OPUS-MT (Tiedemann et al., 2023)³. We then calculate term-term matrices with *textplot* (McClure, 2016). Finally we use the spring attraction algorithm *ForceAtlas2* (Jacomy et al., 2014) in the graph visualizer *gephi*.

3.5 Topic modeling

To validate and refine the conceptual map, we use topic modeling, with BERTopic as the main modeling framework. BERTopic is well suited to heterogeneous corpora because it combines

transformer-based embeddings, dimensionality reduction, clustering, and class-based TF-IDF topic representations (Grootendorst, 2022; Alammari and Grootendorst, 2024). Alternative topic-modeling paradigms such as Top2Vec provide useful methodological points of comparison because they also rely on semantic embeddings rather than purely bag-of-words assumptions (Angelov, 2020).

Topic modeling serves two purposes in our design. First, it checks whether the conceptual distinctions observed in the map persist at the level of document clusters. Second, it helps operationalize the paper’s core categories: i) *amplified*: themes proportionally more prominent in media than in the review corpus; ii) *simplified*: themes present in both corpora but realized through less differentiated vocabulary or narrower narrative associations in media; iii) *absent*: themes clearly established in the review corpus but weakly represented or missing in media.

3.6 Reproducibility

The study is conducted with aligned query logic, explicit corpus boundaries, scripted preprocessing, embedding-based semantic analysis, and topic-model validation. This is also relevant to the SwissText 2026 theme of reproducible NLP, because comparative discourse analysis can otherwise become difficult to audit when corpus construction and preprocessing decisions remain implicit.

4 Results

To ensure analytical clarity, the results are presented in two strictly separated parts: (1) topic modeling across corpora and (2) temporal-semantic patterns in the media corpus. No cross-references are made between these analytical layers.

4.1 Topic modeling across corpora

This subsection reports how the main BERTopic clusters are distributed across the three corpora and how they should be interpreted in relation to one another. We first clarify the topic labels, then compare their relative prominence across scientific education reviews, scientific workplace reviews, and Swiss media passages. In BERTopic, Topic -1 is a standard outlier or residual category: it contains documents that the clustering procedure does not assign to a more coherent topic. We therefore do not interpret Topic -1 as a substantive topic in the same way as Topics 0, 1, 2, or 3, but as an indicator of thematic heterogeneity and weak clustering.

²<https://huggingface.co/sentence-transformers/paraphrase-multilingual-mpnet-base-v2>

³<https://huggingface.co/Helsinki-NLP/opus-mt-roa-en> for French and Helsinki-NLP/opus-mt-de-en for German to English




Figure 1: Distribution of the main topics across corpora.

Topic 3 is the education-centered topic, characterized by terms such as students, learning, skills, and education. Topic 0 captures a broad AI-capabilities discourse, often centered on what AI systems can do and how surprising these abilities appear. Topic 1 and Topic 2 capture more applied or practice-oriented discussions, including uses, tasks, competences, and work-related implications. Topic -1, as noted above, is the residual category.

We first analyze how topics are distributed across the three corpora. In the scientific education corpus (*Sci_Education* in Figure 1), the distribution is highly concentrated: Topic 3 (students, learning, skills, education) accounts for the majority of all content. This indicates a unified and pedagogically anchored discourse, with minimal contribution from other topics. In the scientific workplace corpus (*Sci_Workplace*), the distribution is more heterogeneous. Topic 3 remains dominant (51.28%), but Topic -1 (26.92%) and Topic 2 (14.10%) contribute substantially. This reflects a broader thematic scope that combines general AI discourse with applied and practice-oriented perspectives on skills. In contrast, the media corpus (*SwissDox*) is markedly fragmented. Topic -1 dominates with 52.99%, followed by Topic 0 (9.70%), Topic 1 (7.22%), and Topic 2 (6.01%). Topic 3, i.e., the education-centered topic, accounts for only 1.52%, indicating that pedagogical perspectives are largely absent from media coverage. Topics -1 to 3 cover the vast mass of the textual data, but it is insightful to explore further topics, see the Appendix A2.

4.2 Temporal patterns in the media corpus

We now analyze the conceptual map of the media corpus over time (Figure 2). In this analysis, the publication year (2022–2025) functions as a meta-variable for comparing corpus slices and structur-

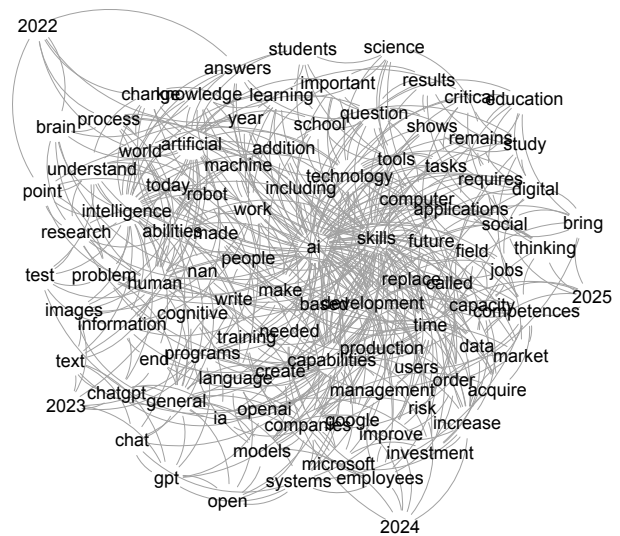


Figure 2: Conceptual map showing German and French translated news articles over time (2022–2025).

ing the semantic space. Figure 2 shows that media vocabulary shifts from early technological novelty in the year 2022 toward later labor-market and societal implications.

The map shows a clear temporal progression of themes. In 2022, the full form "artificial intelligence" is still used frequently, and the hope that we can "understand" the "brain" due to AI "research" (terms appearing in the maps are in straight quotes). In 2022–2023, the discourse is centered on technological emergence and proliferation of use, with clusters around terms such as "chatgpt," "models," "language," and "openai". The "abilities" of AI are surprising. This phase reflects the introduction and initial public engagement with generative AI.

In 2023–2024, the semantics shift toward use and application, with terms such as "tools," "applications," "tasks," "skills," and "users" becoming central. This indicates a transition from novelty to practical experimentation and integration.

In 2024–2025, the discourse turns more oriented toward societal and economic implications, with prominent clusters including "jobs," "market," "risk," "investment," and "competences." The debate focusses on labor-market transformation and future relevance of skills, particularly in 2025, where "social" impact, "jobs" and the question if AI will "replace" us moves into the center, and also which "skills" are required.

Across these phases, we have identified three distinct patterns in this progression:

First, we observe that the media discourse evolves from abstract technological concepts to more concrete socio-economic concerns. A first pattern is therefore *amplification through concrete use cases*. By this we mean that media coverage tends to make LLM-related skill transformation visible through recognizable examples: specific professions, companies, products, risks, public figures, or dramatic scenarios. This concreteness makes the issue more accessible, but it also amplifies those aspects of the debate that are easy to narrate, such as job loss, replacement, or spectacular AI capabilities. The media discourse appears to amplify the visible and dramatizing dimensions of skill transformation, especially those that can be narrated through recognizable actors, professions, or risks. This amplification intensifies in the later period (2024–2025), where labor-market and risk-related terms become more prominent.

The second pattern is *simplification through lexical compression*. Preliminary results indicate that scientific distinctions among skills, competences, literacy, reflection, and higher-order thinking are often collapsed in media discourse into broader and less specialized categories. Within the semantic map, this is reflected in densely connected clusters where heterogeneous terms co-occur without forming clearly separated conceptual domains. In this sense, the media does not simply omit scientific vocabulary; it often translates differentiated scientific debates into broader narratives about adaptation, employability, or disruption.

A third pattern is *absence or weak visibility*. Several dimensions that are prominent in the reviewed scientific literature currently seem underrepresented in media discourse: AI literacy as a structured competence, metacognitive awareness, reflective use of AI, pedagogical integration, and nuanced accounts of augmentation versus substitution. These dimensions do not appear as distinct or central clusters in the semantic structure of the

media corpus. We interpret this as weak public visibility: these concepts may appear occasionally, yet they do not organize the media discourse in the way they organize parts of the scientific literature.

These observations remain preliminary and will need to be tested against topic-model outputs and robustness checks, but they already point to a significant difference between the scientific conceptualization of skill transformation and the public framing available through media discourse.

4.3 Semantic patterns among the corpora including the meta-information

The semantic maps including the different corpora (see Figure 3 in Appendix A3) visualizes a dense, interconnected landscape centered on the transformative impact of artificial intelligence, specifically generative models, on education and the professional world, in distinction to the Swissdax news passages (see "French" and "German" on the left). The news are framing "people" and persons (e.g., Sam Altman and Lukas Walker), express surprise about the "abilities" of the models, often shortcutting methods and frameworks to individual products and companies ("gpt", "openai"). There are aspects relating to the local settings, "swiss" is of course on the side of Swiss news, while scientific concepts like "experimental," "methods," "significant," and "effects" are due to scientific practice.

The core of the map is dominated by technological terms such as "AI," "language models," "automation," and "code," which radiate outward to touch upon critical societal domains. Technological concepts do not exist in isolation; their proximity and repeated connections to terms such as "learning," "teaching," "students," "teachers," and "workplace" indicate that the mapped discourse frequently connects AI technologies with educational and professional activities. The graph indicates a research field that is actively mapping how tools like chatbots and LMMs are reshaping pedagogical methods, assessment strategies, and the very nature of "skills" required in the modern economy. The fact that "skills," "research" and "university" are central in the map indicates that also news media is reporting on the scientific discussion, and universities are seen in a bridging function.

A big portion of the map is dedicated to ethical and practical challenges arising from this integration. Terms like "ethical," "risk," "concerns," and "responsibly" are woven throughout the connections between AI and human systems, highlighting

a pervasive concern for the safe and fair deployment of these technologies. The map suggests that the conversation is not only about efficiency or innovation but equally on potentials for harm, the need for regulation, and the preservation of human agency. This is emphasized by the presence of words like “security,” “replace,” “affect” and “lack”, close to the constant, ambiguous need for “change,” “innovation” and “automation,” counterweights to more optimistic terms like “support,” “facilitate,” “promote.” All these terms are between media and science, but the positive ones closer to science, the negative ones closer to media.

The specific highlighting of “sciwork” and “sci-edu” points to a specialized focus on scientific work and science education, implying that the data underlying this map may originate from academic literature or research projects specifically targeting these domains. The dense clustering around “literature,” “research,” “study,” and “framework” confirms that this visualization likely represents a synthesis of scholarly inquiry rather than general public opinion. The intricate web of connections between “methods,” “evaluation,” “assessment,” and “feedback” suggests a rigorous approach to understanding how AI can be integrated into educational curricula and professional development without compromising academic integrity or professional standards. Ultimately, the map portrays a complex, evolving ecosystem where technology, education, and work are inextricably linked, requiring continuous adaptation and careful consideration of ethical implications.

5 Discussion

The discussion follows the three main findings identified above: amplification through concrete use cases, simplification of the science, and omission of key scientific themes. The preliminary findings support a view of media discourse not as a diluted copy of science, but as a selective reframing of scientific and technological developments for public communication. The analysis does not directly measure public opinion; rather, it identifies the media frames through which members of the public are likely to encounter debates about LLM-related skills.

5.1 Amplification through concrete use cases

First, media discourse amplifies LLM-related skill transformation by making it concrete through rec-

ognizable use cases, actors, products, professions, and risks. What becomes publicly salient is not the full conceptual architecture of the literature, but a narrower subset of dimensions that are easier to narrate: spectacular capability, labor-market disruption, expertise, and competition. This concreteness makes an abstract technological debate accessible, but it also gives disproportionate visibility to dramatic or easily personalized examples, such as ChatGPT, prominent companies, job loss, and replacement scenarios.

5.2 Simplification of the science

Second, media discourse simplifies the science by compressing differentiated concepts into broader and less specialized categories. This is visible in the media’s tendency to collapse distinct pedagogical concepts into single, high-frequency nodes, whereas the scientific networks maintain clearer separations between “learning,” “assessment,” “literacy,” and related concepts. This matters because the scientific literature treats skill transformation as a complex process involving learning design, literacy, metacognition, ethical judgment, and evolving professional competence (Yan et al., 2024; Zhang et al., 2025; Clear et al., 2025). In the scientific maps, terms such as “ethical,” “concerns,” and “responsible” connect AI to human agency; in the media corpus, these distinctions are more often absorbed into generalized narratives about adaptation, risk, disruption, or employability. Deskilling therefore tends to appear in media as a broad anxiety frame, while the literature differentiates between competences that may atrophy under overreliance and those that may become newly important, including prompt design, source evaluation, reflective monitoring, and human–AI coordination.

5.3 Omission of key scientific themes

Third, several key themes from the scientific literature are weakly visible or largely absent in media discourse. These include AI literacy as a structured competence, metacognitive awareness, reflective use of AI, pedagogical integration, and nuanced accounts of augmentation rather than simple substitution. Their weak visibility matters because these concepts would help the public understand what skillful adaptation to LLMs entails.

5.4 Implications for science communication research

These findings extend existing work on AI framing and public attitudes (Brewer et al., 2022; Arnesen et al., 2025) in three ways. First, the analysis quantifies the asymmetry between scientific and public framings within the same semantic problem space: a single pedagogical topic accounts for the majority of the education-review corpus but only 1.52% of the media corpus. Second, the temporal analysis suggests that media framing evolves from technological novelty (2022–2023), through use and application (2023–2024), to labor-market and societal implications (2024–2025), providing an empirical reference point for studies of how emerging-technology discourse stabilizes over time. Third, the persistent synecdochal use of “ChatGPT” as shorthand for AI more generally indicates a product-centered framing that may shape policy and educational debates by tying abstract questions about competence to a single commercial artifact.

From a public-opinion perspective, this is consequential. If citizens encounter LLM-related skills mainly through product-centric and disruption-oriented narratives, democratic debate about education and work may be shaped by a partial repertoire of interpretations. Rather than encountering the full range of scientific concepts used to understand learning, work, and human–AI collaboration, the public is more likely to encounter a narrower discourse centered on capability, disruption, and risk. Methodologically, the study shows the value of comparing media and review literature through semantic and topic-based methods rather than direct content matching: the goal is not literal equivalence, but a structured account of recontextualization across discourse arenas.

6 Conclusion

This study shows that scientific and media discourses construct different public meanings of LLM-related skill transformation. While the scientific literature emphasizes differentiated questions of AI literacy, cognitive skills, ethical control, metacognition, and human–AI collaboration (Brynjolfsson et al., 2025; Wang and Fan, 2025; Zhang et al., 2025), the Swiss media corpus more often foregrounds concrete and product-centered narratives of capability, disruption, replacement, and economic risk. Media discourse therefore does not simply transmit scientific knowledge; it recontextu-

alizes it into public-facing frames that shape which skill futures become publicly visible (Brewer et al., 2022; Arnesen et al., 2025).

Our main contribution is to make this recontextualization empirically visible through the framework of amplification, simplification, and absence. Media coverage amplifies concrete use cases, visible actors, and dramatic risks; simplifies scientific distinctions among skills, competences, literacy, reflection, and professional adaptation; and leaves key scientific themes such as pedagogical integration, metacognitive regulation, and reflective human–AI collaboration weakly visible. The frequent use of “ChatGPT” as shorthand for AI further reinforces a product-centered understanding.

Practically, the findings suggest that public debate about AI and skills may become reactive if it is organized mainly around job loss, technological spectacle, and replacement. Such a framing risks narrowing democratic deliberation and policy imagination: it makes risks visible, but less often shows how skills can be developed, protected, or transformed. More balanced science communication should therefore make visible not only disruption, but also the forms of learning, judgment, AI literacy, and human–AI coordination needed for skillful adaptation in education and work.

The study remains limited by corpus comparability, multilingual semantic alignment, its focus on reviews and Swiss media, and the absence of direct public-opinion or behavioral measures. The year-based analysis also does not yet use formal temporal models such as dynamic topic modeling or change-point detection. Future research should extend the comparison to empirical studies, policy documents, and additional media contexts, and should test how specific media frames shape public attitudes toward AI in education and work. It should also examine the conditions under which more differentiated scientific framings enter public discourse, and whether media coverage can move beyond the binary of “human versus machine” toward a more precise account of augmentation, adaptation, and responsible skill development.

Acknowledgments

The TA-SWISS Foundation is funding the project entitled "Sprachmodelle und Kompetenzverschiebungen," within the framework of which this study is being conducted (see: <https://www.ta-swiss.ch/grosse-sprachmodelle>).

References

- Daron Acemoglu and David Autor. 2011. **Skills, tasks and technologies: Implications for employment and earnings**. In David Card and Orley Ashenfelter, editors, *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier.
- Jay Alammar and Maarten Grootendorst. 2024. *Hands-On Large Language Models: Language Understanding and Generation*. O’Reilly Media.
- Dimo Angelov. 2020. **Top2Vec: Distributed representations of topics**. *Computing Research Repository*, arXiv:2008.09470.
- Svein Arnesen, Tanja S. Broderstad, James S. Fishkin, Mads P. Johannesson, and Alice Siu. 2025. **Knowledge and support for AI in the public sector: A deliberative poll experiment**. *AI & Society*, 40(5):3573–3589.
- Paul R. Brewer, John Bingaman, Albert Paintsil, David C. Wilson, and Will Dawson. 2022. **Media use, interpersonal communication, and attitudes toward artificial intelligence**. *Science Communication*, 44(5):632–658.
- Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. 2025. **Generative AI at work**. *The Quarterly Journal of Economics*, 140(2):889–942.
- Tony Clear, Åsa Cajander, Alison Clear, Roger McDermott, Mats Daniels, Anna Fjellkner, Arto Hellas, Päivi Kinnunen, Juho Leinonen, James Prather, Arnold Pears, Emma Saarinen, and Claudia Szabo. 2025. **AI integration in the IT professional workplace: A scoping review and interview study with implications for education and professional competencies**. In *Proceedings of the 2024 Working Group Reports on Innovation and Technology in Computer Science Education*, pages 34–67. Association for Computing Machinery.
- Kibona Daniel, Mussa M. Msambwa, and Zhe Wen. 2025. **Can generative AI revolutionise academic skills development in higher education? a systematic literature review**. *European Journal of Education*. Advance online publication.
- Michail Giannakos, Roger Azevedo, Peter Brusilovsky, Dragan Gašević, Inge Molenaar, Xavier Ochoa, and 1 others. 2025. **The promise and challenges of generative AI in education**. *Behaviour & Information Technology*, 44(11):2518–2544.
- Pawel Gmyrek, Janine Berg, and Daniel Bescond. 2023. **Generative AI and jobs: A global analysis of potential effects on job quantity and quality**. Working Paper 96, International Labour Organization.
- Maarten Grootendorst. 2022. **BERTopic: Neural topic modeling with a class-based TF-IDF procedure**. *Computing Research Repository*, arXiv:2203.05794.
- Mathieu Jacomy, Tommaso Venturini, Sebastien Heymann, and Mathieu Bastian. 2014. **Forceatlas2, a continuous graph layout algorithm for handy network visualization designed for the gephi software**. *PLoS ONE*, 9(6):e98679.
- Enkelejda Kasneci, Katharina Sessler, Stefan Küchermann, Maria Bannert, Darya Dementieva, Frank Fischer, Ulrike Gasser, Georg Groh, Stephan Günemann, Eyke Hüllermeier, Stephan Krusche, Gitta Kutyniok, Tomer Michaeli, Christian Nerdel, Friedrich Pfeiffer, Oleksandra Poquet, Michael Sailer, Albrecht Schmidt, Tina Seidel, and 2 others. 2023. **ChatGPT for good? on opportunities and challenges of large language models for education**. *Learning and Individual Differences*, 103:102274.
- Philipp Lorenz-Spreen, Lisa Oswald, Stephan Lewandowsky, and Ralph Hertwig. 2023. **A systematic review of worldwide causal and correlational evidence on digital media and democracy**. *Nature Human Behaviour*, 7:74–101.
- David McClure. 2016. **textplot**. <https://github.com/davidmcclure/textplot>. GitHub repository.
- Dong Nguyen and Ekaterina Hekman. 2024. **The news framing of artificial intelligence: a critical exploration of how media discourses make sense of automation**. *AI & Society*, 39:437–451.
- Nils Reimers and Iryna Gurevych. 2019. **Sentence-BERT: Sentence embeddings using siamese BERT-networks**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 3982–3992. Association for Computational Linguistics.
- Nader Salari, Mojtaba Beiromvand, Ali Hosseinian-Far, Masoumeh Mohammadi, Bahareh Fatahi, and Siti Shohaimi. 2025. **Impacts of generative artificial intelligence on the future of labor market: A systematic review**. *Smart Analytics, Artificial Intelligence and Sustainable Performance*, 1(1):100067.
- Sidra Sidra and Christine Mason. 2024. **Reconceptualizing AI literacy: The importance of metacognitive thinking in an artificial intelligence-enabled workforce**. In *2024 IEEE International Conference on Teaching, Assessment and Learning for Engineering (TALE)*. IEEE.
- Brandon M. Stewart, Justin Grimmer, and Margaret E. Roberts. 2022. *Text as Data: A New Framework for Machine Learning and the Social Sciences*. Princeton University Press, Princeton, NJ.
- Jörg Tiedemann, Mikko Aulamo, Daria Bakshandaeva, Michele Boggia, Stig-Arne Grönroos, Tommi Nieminen, Alessandro Raganato Yves Scherrer, Raul Vazquez, and Sami Virpioja. 2023. **Democratizing neural machine translation with OPUS-MT**. *Language Resources and Evaluation*, (58):713–755.

Jiaqi Wang and Wei Fan. 2025. The effect of ChatGPT on students' learning performance, learning perception, and higher-order thinking: Insights from a meta-analysis. *Humanities and Social Sciences Communications*, 12:744.

Marta Wiczorek, Mohammad Hosseini, and Bert Gordijn. 2025. Unpacking the ethics of using AI in primary and secondary education: A systematic literature review. *AI and Ethics*. Advance online publication.

Allison Woodruff, Renee Shelby, Patrick G. Kelley, Shiri Rousso-Schindler, Jeffrey Smith-Loud, and Lauren Wilcox. 2024. How knowledge workers think generative AI will (not) transform their industries. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pages 1–26. Association for Computing Machinery.

Lixiang Yan, Lei Sha, Ling Zhao, Yi Li, Roberto Martinez-Maldonado, Guanliang Chen, Xiaoming Li, Yiqi Jin, and Dragan Gašević. 2024. Practical and ethical challenges of large language models in education: A systematic scoping review. *British Journal of Educational Technology*, 55(1):90–112.

Shuo Zhang, P. Ganapathy Prasad, and Noah L. Schroeder. 2025. Learning about AI: A systematic review of reviews on AI literacy. *Journal of Educational Computing Research*. Advance online publication.

A Appendices

A.1 Search-queries used to retrieve the documents

For the education domain, the exact query was: ("large language models" OR "LLM" OR "LLMs" OR "AI" OR "GenAI" OR "artificial intelligence") AND ("skills" OR "competences" OR "upskilling" OR "up-skilling" OR "deskilling" OR "de-skilling") AND ("students" OR "pupils") AND ("education" OR "school"). For the workplace domain, the query was: ("large language models" OR "LLM" OR "LLMs" OR "AI" OR "GenAI" OR "artificial intelligence") AND ("skills" OR "competences" OR "upskilling" OR "up-skilling" OR "deskilling" OR "de-skilling") AND ("workplace" OR "job" OR "workforce" OR "labor" OR "labour" OR "market").

The conceptual query used for Swissdax corresponds to the following structure: ("LLM" OR "LLMs" OR "GPT" OR "AI" OR "KI" OR "IA" OR "GenAI" OR "large language models" OR "grosse Sprachmodelle" OR "grands modèles de langage" OR "artificial intelligence" OR "künstliche Intelligenz" OR "intelligence artificielle") AND ("skills" OR "compétences" OR "capacités" OR "Fähigkeiten" OR "Kompetenzen" OR "deskilling" OR "Entqualifizierung" OR "déqualification" OR "de-skilling" OR "De-Qualifizierung" OR "dé-qualification" OR "upskilling" OR "Weiterqualifizierung" OR "requalification" OR "up-skilling" OR "Up-Skilling").

A.2 Description of the topics 4 to 10

Topics -1 to 3 cover the vast mass of the textual data, but it is insightful to explore further topics. The keywords of topics 4 to 10, followed by our interpretation and an excerpt from the most prototypical passage are:

Topic 4: ki, Or0, fähigkeiten, intelligenz, ia, menschliche, intelligence artificielle, artificielle, capacités, künstliche. Topic 4 focusses on how similar AI is to human intelligence. “Früher als gedacht könnten die Fähigkeiten von künstlicher Intelligenz denjenigen eines menschlichen Gehirns entsprechen.”

Topic 5: capacités, capacités calcul, calcul, ia, données, développement, microsoft, centres, ki fähigkeiten, plus. Topic 5 is about the computing centres and their energy consumption “... l’IA générative, qui nécessite des capacités de calcul colossales pour traiter les informations accu-

mulées”

Topic 6: lernen, müssen, wurden, kompetenzen, compétences, führt, denken, gpt, wichtiger, daran. Topic 6 takes up the discussion on skills needed in society, it is close to topic 3, but with concrete framing. “Diese Kompetenzen wurden schon immer gebraucht. Aber sie wurden wichtiger. Gerade jetzt, wo die künstliche Intelligenz hinzukommt, führt kein Weg daran vorbei, das Lernen an sich zu thematisieren”

Topic 7: capacités, scientifique, intelligence artificielle, artificielle, ki, Fähigkeiten, zurich, avancée, intelligence, dix. Topic 7 takes up the scientific discourse and fulfills a bridging function. “Une approche en collaboration avec de nombreuses institutions dont les Universités de Genève et Zurich et les deux Écoles polytechniques fédérales.”

Topic 8: gar, eigenen, Intelligenz, apple, co, kognitiven Fähigkeiten, kognitiven, weitere, chatgpt, vielmehr. Topic 8 features philosophical discourses which are often absent in scientific literature. “Was uns bei ChatGPT als künstliche Intelligenz erscheint, ist demnach lediglich ein Spiegel, der unsere eigenen Wünsche, Bedürfnisse und kognitiven Fähigkeiten zurückwirft.”

Topic 9: Intelligenz, künstliche, künstliche Intelligenz, menschliche Fähigkeiten, imitieren, Fähigkeit, menschliche, 00, Fähigkeiten, Kreativität. Topic 9 seems as unspecific as Topic -1. What may distinguish it is its pedagogical focus. “Dieser Artikel stammt aus der Schweizer Familie Was genau ist künstliche Intelligenz?”

Topic 10: générale, sam, sam altman, altman, humaines, intelligence artificielle, artificielle, cognitives humaines, intelligence, artificielle générale. Topic 10 shows the personalization which is typical for news reporting. “Son patron, Sam Altman, a récemment expliqué travailler désormais vers l’intelligence artificielle dite «générale» ...”

A.3 Additional conceptual maps

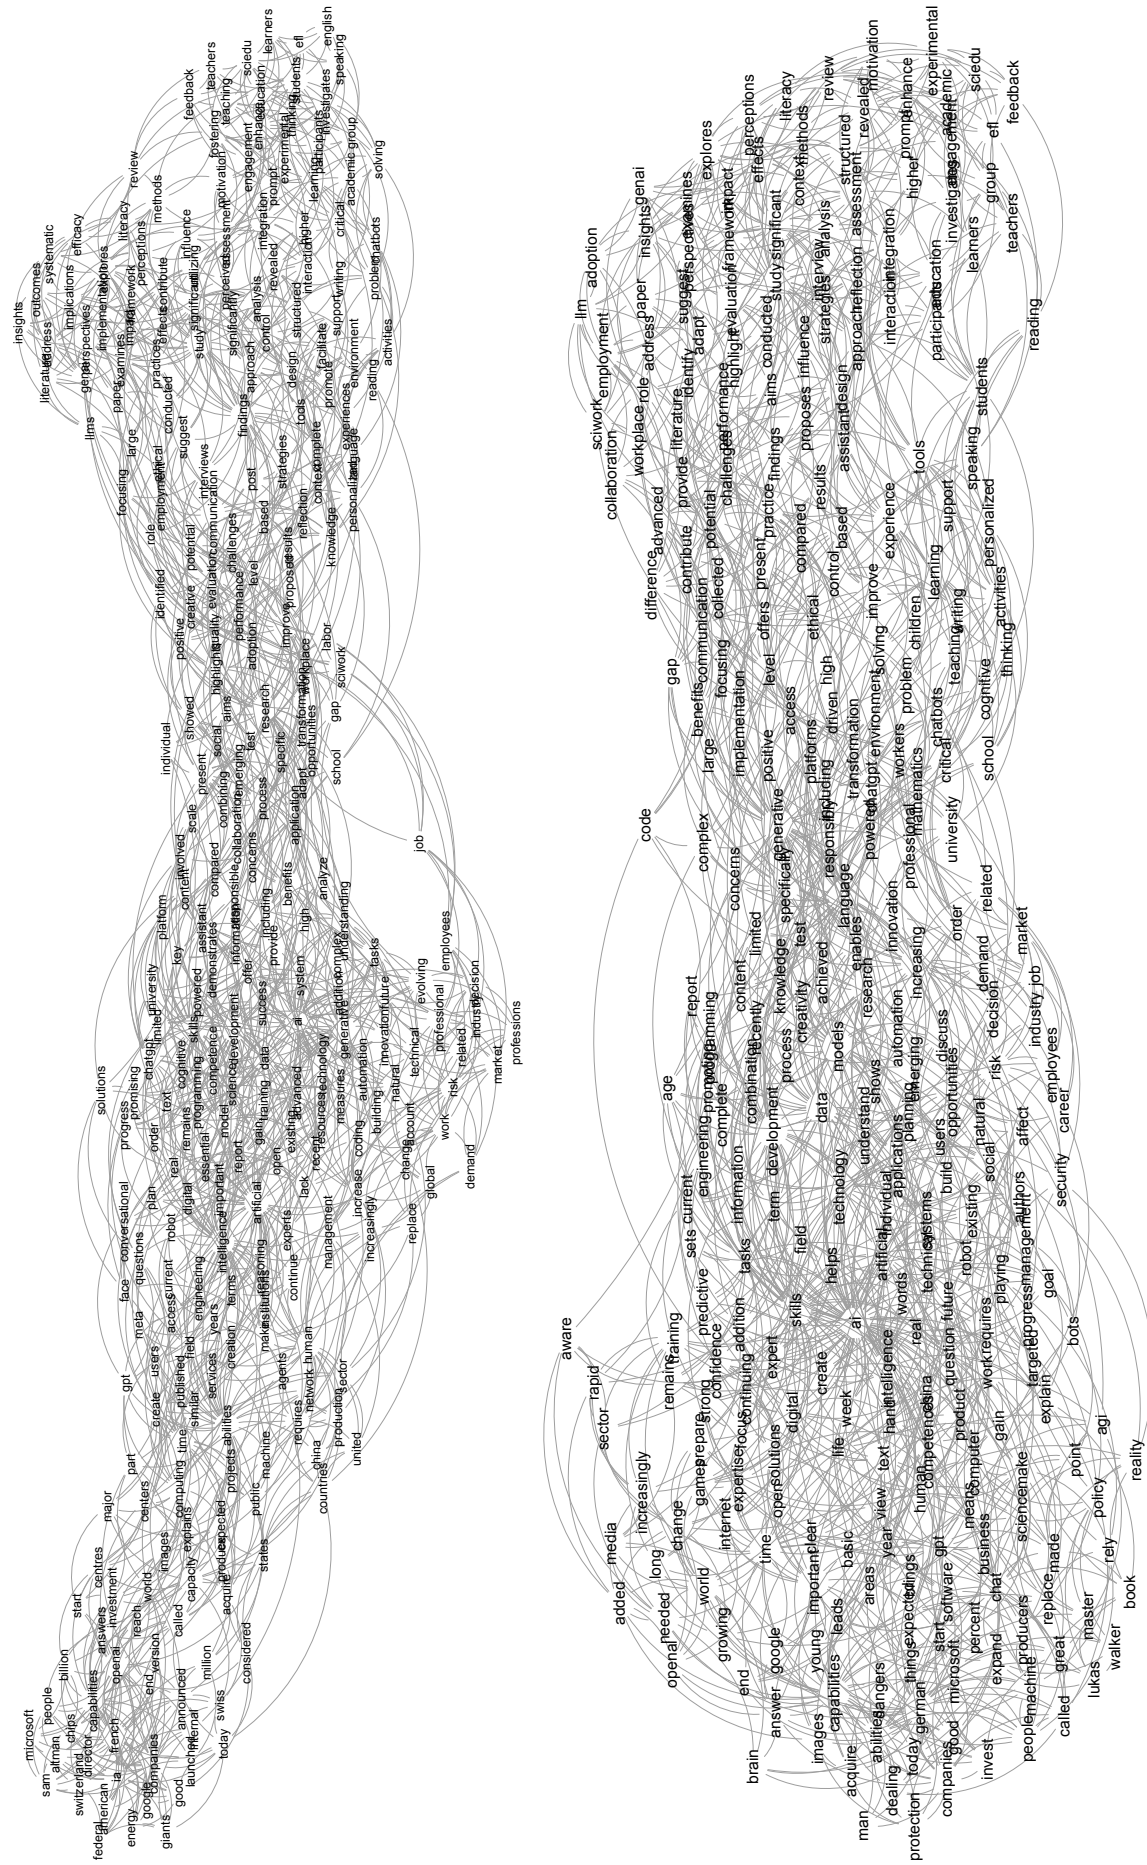


Figure 3: Conceptual maps of French (“french,” left on the top graph) and German (“german”, left at the bottom graph) news articles versus the scientific literature (“sciwork” and “sciedu”, on the right of each graph).