

**Artificial Intelligence in Humanities and Arts Faculty Job Postings:
A Mixed-Methods Analysis of 1,085 HigherEdJobs.com Listings**

Elizabeth Elie

Department of Educational Technology, New Jersey City University

EDTC 813 — Assessment 3: AI in Higher Education Project

Dr. Christopher Shamburg

August 23, 2026

Introduction

Three years after generative large language models became publicly available, artificial intelligence (AI) has ceased to be the exclusive concern of computing departments. It now surfaces in the hiring documents of academic units with no historical relationship to computation: English, music, philosophy, graphic design, journalism, and library science.

Faculty job postings are a particularly informative site for observing this diffusion, because a posting is not a description of what a department currently does but a public declaration of what it intends to become. Unlike a syllabus revision or a workshop series, a tenure-track line commits budget, curriculum, and disciplinary identity for decades. When a music school writes AI into a permanent position, the institution is making a claim about the future shape of its field.

This study examines a corpus of 1,085 full-time faculty postings collected from HigherEdJobs.com that explicitly mention AI, and narrows its focus to the subset drawn from the humanities and the arts. These are the fields in which AI is least expected, and where its appearance therefore carries the most information about how the technology is being absorbed into disciplines that define themselves through interpretation and critique rather than computation.

Purpose and Research Questions

The purpose of this short mixed-methods study is to describe how AI is positioned within full-time faculty job postings in humanities and arts fields, and to determine whether that positioning differs in kind, and not merely in degree, from its positioning in computing-adjacent fields. Three research questions guide the inquiry.

First, how prevalent and how central is AI in humanities and arts postings relative to the full corpus? Second, what themes characterize the way AI is framed in these postings? Third, how does that framing compare with the computing, engineering, and science categories of the same corpus? Together these address a single underlying issue: whether AI enters the humanities and the arts as an imported technical skill set or as a distinctly disciplinary object of study.

Literature Review

Recent scholarship has established that faculty engagement with AI is uneven and that discipline is a meaningful source of that unevenness. Mah and Groß (2024) surveyed 122 faculty members and identified four latent profiles — optimistic, critical, critically reflected, and neutral — finding that the absence of AI literacy among both students and faculty was among the most frequently named challenges. In a subsequent study of 112 academic staff, Mah et al. (2025) reported that participants assigned the greatest importance to the ethics of AI in education, above instructional design and domain specificity, yet perceived relevance of ethics did not predict actual adoption of AI tools.

Kassorla et al. (2024) offer a complementary conceptual contribution, defining AI literacy in teaching and learning as a composite of technical understanding, evaluative judgment about tool outputs, and an ethical commitment to transparent use. What these studies share is a recognition that AI competence in higher education is not reducible to technical proficiency; it includes an interpretive and normative dimension that humanities and arts faculty are, in principle, well positioned to supply.

A second strand documents how AI is reshaping expectations within the creative disciplines. Fleischmann (2024) surveyed graphic design students and found strong interest in generative AI alongside significant uncertainty about how it should figure in professional

practice. Hajkowicz et al. (2023), in a bibliometric analysis of six decades of publications, showed that AI adoption has extended into the arts and humanities, though at a markedly lower volume and later onset than in the physical and life sciences.

Closest to the present study methodologically is Nastasa et al. (2025), who analyzed approximately 800 AI-related academic job postings from the European EURAXESS platform using text mining and topic modeling. They identified seven thematic clusters interpreted as distinct epistemic cultures, and observed that although AI research presents itself as interdisciplinary, computer science and engineering constitute a dominant core around which other disciplines are frequently oriented. Their corpus, however, was overwhelmingly STEM in composition. The present study takes up the same analytic premise — that postings are discursive claims rather than neutral descriptions — and applies it to the fields their sample could not reach.

Method

Data Source and Analytic Sample

The dataset is a spreadsheet of 1,085 full-time faculty job postings retrieved from HigherEdJobs.com, filtered from an initial population of 19,495 listings by whether the posting explicitly mentions artificial intelligence. Each row is one posting and contains the institution, position title, posting URL, disciplinary field, category, salary information where disclosed, and the full text of the description.

The analytic sample for this study is every posting the dataset assigns to the Humanities or Arts categories. That yields 108 postings, 9.95% of the corpus, spanning 25 fields and 69 institutions. This subset was chosen because it is the disciplinary space furthest from computing,

and because the category boundary is a property of the dataset rather than a judgment introduced by the researcher, which makes the sampling decision transparent and replicable.

Quantitative Procedure

The quantitative component is descriptive. It does not test a hypothesis and reports no inferential statistics; its purpose is to establish how much AI language each posting contains and how prominently that language is placed, so that the humanities and arts subset can be compared against the other seven categories. The procedure was carried out in five steps.

Step 1. Count AI vocabulary in each description. For every one of the 1,085 postings, the description text was searched for four expressions: artificial intelligence, machine learning, AI, and LLM. Matching was case-insensitive, and word-boundary constraints were applied so that the letters ai occurring inside other words, such as available or chair, were not counted. The result is a single number per posting, ai_mentions, ranging from 0 to 26.

Step 2. Flag AI in the position title. Each posting was separately checked for whether the same vocabulary appears in the job title rather than the body, producing a yes-or-no variable, ai_in_title. Title placement is treated as distinct from frequency because a title is the shortest and most contested text in a posting; language reaching it has survived institutional negotiation.

Step 3. Assign a centrality band. The two measures above were combined into one three-level variable. A posting was banded Central if AI appeared in the title or the description contained five or more mentions, Integrated if the description contained two to four mentions, and Peripheral if it contained one mention or none. The thresholds are descriptive conventions rather than established cut

points: two mentions distinguishes a posting that returns to AI from one that names it once in a list, and five approximates the corpus mean for Computer Science, the most AI-saturated category. Because the bands are conventions, both the underlying count and the band are reported so a reader can apply different thresholds.

Step 4. Compute the same measures for all nine categories. Every statistic produced for the humanities and arts subset was also produced for the full corpus, so the subset is interpreted against a baseline rather than in isolation.

Step 5. Tabulate contextual variables. Rates of salary disclosure, institutional distribution, and position rank were counted for the subset and for the corpus.

The variables produced by this procedure are summarized in Table 1.

Table 1

Variables Produced by the Quantitative Procedure

Variable	Values	What it records
ai_mentions	Count (0–26)	Number of times AI vocabulary appears in the posting description.
ai_in_title	Yes / No	Whether the position title itself contains AI vocabulary.
centrality	Central / Integrated / Peripheral	Three-level band combining the two measures above.
category	9 categories	Disciplinary category assigned by the dataset.
salary_disclosed	Yes / No	Whether a salary or range appears in the posting.
rank	6 ranks	Position rank stated in the title.

Qualitative Procedure

The qualitative component is an inductive thematic analysis of the 108 postings in the subset. Every sentence containing an AI reference was extracted, producing 298 AI-bearing

sentences across 103 postings; five postings contained no AI language in their description at all and could not be coded.

These sentences were read in full and grouped according to the function AI performed — what AI was being asked to be, rather than which words surrounded it. Five themes emerged from that reading. Each posting was then coded twice against those themes: once for which themes appear anywhere in it, since most postings blend framings, and once for which single theme is primary, judged by reading the posting as a whole. Reporting both is what allows the finding that dominance and presence diverge, as they do for creative practice, which is primary in 18.5% of postings but present in 41.7%.

Limitations

Three limitations apply. The categorical assignment of fields, including the placement of psychology within Humanities, is inherited from the dataset and would differ under other taxonomies. Nine postings originate outside the United States despite the corpus being framed as a study of U.S. institutions. And posting language reflects institutional intention rather than realized practice, so nothing here should be read as evidence of what these departments actually do once a hire is made.

Findings

Prevalence and Density

The humanities and arts are underrepresented in the corpus relative to their disciplinary breadth. The 108 postings constitute 9.95% of all listings while accounting for 25 of the 114 fields present (21.9%), meaning these disciplines contribute roughly a fifth of the corpus's field diversity but a tenth of its volume.

Within the subset, AI mentions per posting averaged 3.72 (Mdn = 2.00, SD = 4.19, range 0–26). This aggregate conceals a meaningful internal split. Considered separately, Humanities averaged 4.01 mentions per posting — the second highest of the nine categories, exceeded only by Computer Science at 5.92 — while Arts averaged 2.76. Humanities postings are therefore not marginal participants in AI discourse by volume; they are among its densest sites outside computing. Salary was disclosed in 47.2% of subset postings against 35.6% corpus-wide, and assistant professor was the most common rank (n = 47).

Table 2

AI Mention Density and Title Prominence by Category (N = 1,085)

Category	n	M mentions	Mdn	AI in title (%)
Computer Science	356	5.92	4.0	39.6
Humanities	83	4.01	2.0	14.5
Science	146	3.38	2.0	15.1
Engineering	178	3.20	2.0	9.0
Education	32	2.84	2.0	12.5
Arts	25	2.76	2.0	16.0
Biology	133	2.49	2.0	11.3
Business	119	2.34	2.0	3.4
Vocational	13	2.31	2.0	0.0
Full corpus	1,085	3.96	2.0	20.1

Centrality

Centrality distinguishes these fields sharply from computing. AI appeared in the position title of 17 subset postings (15.7%), against 39.6% in Computer Science and 20.1% corpus-wide. Applying the bands, 31 postings (28.7%) were Central, 42 (38.9%) Integrated, and 35 (32.4%) Peripheral. Six postings (5.6%) mentioned AI nowhere in the description, a discrepancy attributable to search matching on metadata rather than posting text.

The distribution varies substantially by field. Library and Information Science, the single largest field in the subset ($n = 15$), produced eight Central and seven Integrated postings and no Peripheral ones — the only field in which AI was role-defining in every listing. Psychology ($n = 14$), Media and Communication Studies ($n = 11$), and Digital Arts ($n = 7$) each contained a mix of all three levels, while Art History, Music, and several smaller fields skewed Peripheral. Nine postings mentioned AI only in reference to the hiring process itself, either disclosing that AI tools might be used in candidate review or, in one state system, affirming that automated decision systems are not used in selection.

Themes

The thematic analysis identified five ways AI functions in these postings. The most common primary theme was Pedagogy and AI Literacy (32 postings, 29.6%), in which AI figures as something students must be taught to use and evaluate. Critical and Ethical Inquiry (27, 25.0%) positions AI as an object of humanistic analysis rather than a tool. Creative and Professional Practice (20, 18.5%) treats AI as a medium or workflow component. AI as Research Method (14, 13.0%) invokes AI as instrumentation for disciplinary scholarship. Institutional Context Only (10, 9.3%) captures postings in which AI appears solely in descriptions of campus initiatives or hiring procedures.

Table 3

Thematic Distribution of Humanities and Arts Postings ($n = 108$)

Theme	Primary	% primary	% present (any)
Pedagogy and AI literacy	32	29.6	52.8
Critical and ethical inquiry	27	25.0	39.8
Creative and professional practice	20	18.5	41.7
AI as research method	14	13.0	32.4
Institutional context only	10	9.3	29.6

Theme	Primary	% primary	% present (any)
Uncodable (no AI in text)	5	4.6	—

Measured as presence rather than dominance, pedagogy appeared somewhere in 52.8% of postings, creative practice in 41.7%, critical inquiry in 39.8%, research method in 32.4%, and institutional context in 29.6% — confirming that most postings blend framings rather than adopting one exclusively.

Cross-Category Comparison

Ethics language co-occurring with AI appeared in 16.0% of Arts postings and 13.3% of Humanities postings, against 7.0% in Computer Science, 5.3% in Biology, 2.8% in Engineering, and 2.1% in Science. The phrase AI literacy, or an explicit variant, appeared in 6.0% of Humanities postings and in less than 1% of Computer Science postings; in Engineering, Science, Biology, and the Arts it did not appear at all. The pattern is consistent across both measures: computing fields mention AI far more often and far more prominently, but humanities fields are markedly more likely to attach normative and instructional framing when they do. Where a computing posting names a technique, a humanities posting more often names a responsibility.

Discussion

These findings both confirm and complicate the existing literature. They align with Mah et al. (2025) on the disproportionate importance faculty assign to AI ethics, and extend that finding from self-reported attitudes to institutional hiring language, which is a harder and more consequential form of commitment. Where survey research captures what individual faculty say they value, a posting captures what a department was willing to write into a permanent line.

The findings extend Nastasa et al. (2025) in one respect and complicate their conclusion in another. Their observation that AI-related hiring is organized around a computational core is

supported by the centrality data. But their inference that other disciplines are consequently oriented around that core, rather than partnered with it, is not fully borne out. Humanities is the second most AI-saturated category in the corpus, and when humanities departments engage AI they most often do so through frames native to their own epistemic traditions rather than borrowed from computing. Library and Information Science is the clearest case: every one of its fifteen postings treated AI as role-defining, with framings running to information literacy, attribution, copyright, and responsible use rather than model development. This looks less like orbit than like annexation.

Two findings complicate any straightforwardly optimistic reading. First, the substantial Peripheral tier — nearly a third of the subset — indicates that a large share of what registers as AI hiring in the humanities is boilerplate. Aggregate counts of AI-mentioning postings will therefore overstate genuine disciplinary uptake unless centrality is measured. Second, the nine postings that mention AI only in relation to candidate review are a reminder that AI now shapes academic hiring as a process, not only as a subject, and that faculty in these fields may encounter algorithmic systems as applicants before they encounter them as scholars.

Conclusion

Artificial intelligence enters humanities and arts faculty hiring less often than it enters computing, but when it does, it enters differently. Across 108 postings from 69 institutions, AI appeared most frequently as something to teach students to evaluate and as something to subject to ethical and critical analysis, and only secondarily as a research instrument or creative tool. Humanities postings proved the second densest in AI language of any category while remaining far less likely than computing postings to name AI in the position title, and they attached ethics and literacy framing at rates two to six times those of the STEM categories.

The picture that emerges is not one of non-traditional fields catching up to a technical agenda set elsewhere, but of disciplines translating AI into their own terms and hiring accordingly — while a substantial peripheral tier cautions that institutional AI language often runs ahead of institutional AI commitment. Future work should track whether the Central postings identified here result in sustained curricular change, and should extend the same centrality-sensitive method to the health professions.

References

- Fleischmann, K. (2024). Generative artificial intelligence in graphic design education: A student perspective. *Cogent Education*, 11(1), 2287891.
- Hajkowicz, S., Sanderson, C., Karimi, S., Bratanova, A., & Naughtin, C. (2023). Artificial intelligence adoption in the physical sciences, natural sciences, life sciences, social sciences and the arts and humanities: A bibliometric analysis of research publications from 1960–2021. *Technology in Society*, 74, 102260.
- Kassorla, M., Georgieva, M., & Trach, E. (2024). *AI literacy in teaching and learning*. EDUCAUSE.
- Mah, D.-K., & Groß, N. (2024). Artificial intelligence in higher education: Exploring faculty use, self-efficacy, distinct profiles, and professional development needs. *International Journal of Educational Technology in Higher Education*, 21, 58.
- Mah, D.-K., Groß, N., & Kolodziej, R. (2025). Ethics of artificial intelligence in education: Perspectives of academic staff. *International Journal of Educational Technology in Higher Education*.
- Nastasa, A., Popescu, A., & Ionescu, L. (2025). Epistemic cultures in AI research: A topic-modeling analysis of academic job postings. *Journal of Documentation*.