

# AI Pedagogy and New Methods for Humanities Scholars: A Reflective Case Study

Chris Haffenden<sup>1</sup>, Justyna Sikora<sup>2,\*</sup>

<sup>1</sup>*KBLab, Kungliga biblioteket, Karlavägen 100, 115 26 Stockholm, Sweden*

<sup>2</sup>*KBLab, as above*

## Abstract

This paper presents a reflective case study on teaching AI methods to humanities scholars through the example of a workshop developed at KBLab, the AI and digital research lab at the National Library of Sweden. Conducted within the framework of the Swedish national research infrastructure Huminfra, the workshop introduced participants from heritage organisations and academia to multimodal topic modelling using CLIP-Topic and Google Colab. Drawing on this experience, we discuss pedagogical strategies for bridging the divide between large-scale computational analysis and the interpretive traditions of the humanities. We reflect on design choices such as how much code to expose, how to balance explanation and demonstration, and how to make domain-specific applications intelligible to non-technical users. The case illustrates how script-based, open formats can foster AI literacy and critical engagement, while highlighting the limits of one-off workshops and the need for sustained, cross-disciplinary collaboration.

## Keywords

AI pedagogy, AI literacy, multimodal topic modelling, digital humanities, cultural heritage collections,

## 1. Introduction

How can we bridge the divide between increasingly large-scale computational approaches to GLAM collections and the more qualitatively inclined perspectives of humanities scholars [1], while also raising awareness of the possibilities and limitations of AI-based search systems now entering the research landscape? What pedagogical strategies are available to those of us who work hands-on with digital research infrastructures and research services, seeking to connect these methodological worlds?

This paper draws on our experience at KBLab, the AI and digital research lab at the National Library of Sweden [2], and our outreach initiatives within the Swedish national research infrastructure Huminfra. We discuss the design and delivery of a new workshop developed to illustrate the potential of multimodal topic modelling for both heritage organisations with large collections of unlabelled images, and researchers working with visual culture.

The workshop builds on our earlier work with using **CLIP** (Contrastive Language–Image Pre-training) to enhance the searchability of the library’s postcard holdings [3] and with creating a user-friendly, script-based workshop for **BERTopic** [4]. In this paper, we reflect on our design choices in using Google Colab to demonstrate how image collections can be clustered according to topic—for instance, how much code to include and how many steps can realistically be explained. We also discuss the affordances and limitations of web-based workshops as a mode of methodological outreach.

By focusing on this concrete scenario, we highlight how AI literacy efforts in the humanities can be rooted in domain-specific needs and interpretive questions. We argue that effective outreach requires not only demystifying technical tools but also fostering dialogue around how these tools intersect with established scholarly practices. The case study illustrates the value of flexible, script-based formats that empower researchers to engage critically with emerging methods without requiring full technical

---

*Huminfra Conference 2025, Stockholm, 12–13 November 2025.*

\*Corresponding author.

✉ [chris.haffenden@kb.se](mailto:chris.haffenden@kb.se) (C. Haffenden); [justyna.sikora@kb.se](mailto:justyna.sikora@kb.se) (J. Sikora)

ORCID [0000-0002-5561-5163](https://orcid.org/0000-0002-5561-5163) (C. Haffenden)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

fluency. At the same time, we remain cautious about the transformative potential of one-off workshops, emphasising instead the importance of longer-term, transdisciplinary forms of methodological collaboration.

## 2. CLIP-Topic workshop: translating infrastructure into pedagogy

Before describing the workshop itself, we briefly place it within the broader context of digital research infrastructures. Initiatives such as KBLab and Huminfra increasingly act not only as technical enablers but also as pedagogical intermediaries: they are often where humanities scholars encounter AI methods in depth for the first time, bringing questions of usability, literacy and interpretation to the fore. In this sense, infrastructures have become *pressure points* for digital pedagogy—sites where the demands of large-scale computation meet the interpretive expectations of the humanities. We conceived of our CLIP-Topic workshop within this intersection, using the lab’s resources and expertise to translate an advanced AI method into an accessible, conceptually meaningful learning experience.

The AI method demonstrated in the workshop rests on a simple but powerful idea: *vector-based search*, which allows computers to measure similarity not through exact matches, but through proximity in meaning or appearance. In the practical setting of multi-modal heritage collections, this means that both words and images can be converted into numerical representations—known as *embeddings*—and compared within a shared vector space. Items that are close to one another in this space are interpreted as being semantically or visually related. This principle underpins the recommendation engines, image searches, and content-suggestion systems that most of us interact with daily—when a smartphone groups photos by theme, or when a streaming platform suggests films “you might also like”. The same technique enables us to cluster and explore large-scale heritage collections by their contents rather than their (often scant) metadata, offering new entry points into archives that have long remained opaque to search [3].

### 2.1. Workshop aims

Our aim with the workshop in broad terms was to showcase how the same AI techniques we encounter in everyday online contexts can be harnessed to improve the accessibility and research potential of image collections. Beyond introducing the conceptual foundations of such techniques, we also wanted to convey the practical mechanics of working programmatically with “collections as data”—that is, through a hands-on immersion in running code.

The target audience consisted of two groups, partially overlapping but principally divergent: heritage professionals with responsibility for digital collections management, and academic scholars in the humanities and social sciences working with digital images. The learning objectives for these groups differed slightly. In both cases, we wanted to demonstrate a sense of the possibilities of using vector search to cluster and organize a large volume of documents—in this case, images—according to their thematic contents rather than preexisting (and often limiting) categories [5]. More specifically, for heritage professionals we aimed to offer tangible inspiration for applying AI-based search to their own collections, and for researchers, to provide a concrete sense of how such techniques can serve as a starting point for exploring and analyzing visual research data.

### 2.2. Designing the workshop

To enable participation on a national scale, we chose to make the workshop web-based, ensuring that heritage professionals and researchers outside the Stockholm area could take part. Conceived as a pedagogical activity delivered via Zoom, the workshop followed a three-part structure, which we outline below.

The first part consisted of a *conceptual introduction*—a twenty-minute presentation that familiarised participants with two key components of the workshop. First, we outlined the particular challenges facing heritage institutions with large collections of unlabelled images, to which CLIP-Topic offers a potential

solution. The aim here was to anchor what might otherwise seem a highly technical method within a concrete practical context, underlining that such tools are valuable insofar as they help us address the real challenges of working with digital culture. Second, we provided a brief overview of the conceptual architecture of the method, beginning with an explanation of OpenAI's CLIP model [6] and followed by a sketch of the logic behind topic modelling [7]. The goal was to give participants a basic understanding of multimodal topic modelling without confusing or alienating non-technical attendees with excessive detail; accessibility and broad brushstrokes were deliberately prioritised over depth. (For those interested in exploring the method in more depth, we included a slide with links to further reading.)

The second and main part of the workshop consisted of a *live demonstration*: a guided walkthrough of a Google Colab notebook we created to show how CLIP-Topic can cluster and visualise images by topic [8]. We chose Colab as the pedagogical platform because it allows participants to run pre-written Python notebooks directly in the browser, while also providing free access to the GPUs needed to fit the models. This approach removes typical barriers of software installation and environment setup, while keeping the process transparent—every cell of code can be read, edited and executed. We had previously used Colab for workshops on text-based topic modelling with BERTopic [4], where participants particularly appreciated being able to revisit and rerun the notebooks at their own pace. The same logic guided this design.

To ensure it could be published as an open resource—something participants could return to and others discover and use on their own terms—we based the workshop on openly available Swedish heritage data. The notebook uses a dataset of 1,508 images downloaded from DigitaltMuseum. We selected these as situated examples: by choosing heritage images likely to be familiar to participants in both style and sentiment, we hoped to anchor the more technical discussion of AI models in the professional realities of those taking part. In addition to emphasizing that the notebook can be revisited for future learning, we also highlighted its adaptability—users are encouraged to insert their own data into the script and run the analysis on their specific collections.

The notebook itself was organised to alternate between short explanatory sections and executable code cells, allowing participants to see each step of the process before running it. This structure was intended to make the workflow intelligible even to those without prior programming experience, while still offering enough detail for more technically confident users to experiment further.

The third and final part centered on *applications and questions*. After completing the Colab walkthrough, we shifted focus from the mechanics of the method to what it can reveal in practice. Using examples drawn from ongoing experiments with image search and description at both KBLab and Stanford University Library, we demonstrated how vector-based techniques can surface new thematic connections within heritage collections and support visual exploration at scale [9]. This segment also served as an open discussion space: participants were invited to reflect on how such approaches might translate to their own institutional or research contexts, and to raise questions about, for instance, interpretability. In pedagogical terms, the aim was to re-connect the technical and conceptual threads of the workshop—moving from doing to thinking, and from method to meaning—thereby reinforcing AI literacy as both a practical and interpretive competence.

This three-part format—combining conceptual framing, guided demonstration and open discussion—was designed to balance accessibility and depth. It aimed to help participants connect technical operations with interpretive reflection and imagine how AI methods might be meaningfully adapted to their own disciplinary and institutional settings.

### 3. Pedagogical reflections: Successes and Challenges

When we first ran the workshop, over 60 people had registered in advance and more than 40 attended the two-hour online session. The fact that participants came from such a wide range of institutions—from major national museums to small local archives, and from universities stretching from Florence and Cork to Luleå—attests both to the topicality of the subject and the benefits of circulating the invitation through diverse scholarly and professional networks. In what follows, we offer some reflections on what worked

well with the workshop as a pedagogical event, and what proved more challenging.

### 3.1. What worked

**Transparency and legibility of code.** In an era of chat-based LLMs offering researchers immediate results but little accountability for how those results are produced, one of the key advantages of a step-by-step, script-based workshop was the *sense of process* it provided. By including every stage—from loading data and producing embeddings, to fitting the model and visualising the results—our CLIP-Topic notebook helped participants grasp the logic of the entire workflow. Including each line of executable code also conveyed that this was about the operation of a legible method, rather than the “silver-bullet” magic of AI. Through combining explanatory comments with code cells, we could walk participants through three modes of topic modelling images—using only textual descriptions, only the images themselves, and a combination of the two—accompanied by a running commentary on the distinct effects and merits of each approach.

**Advanced access as icebreaker.** From previous experience, we have learned the value of sharing the notebook a few days before the workshop. This gives participants the opportunity to explore the code and familiarise themselves with the method in advance. Such early access can be particularly beneficial for non-technical users who might otherwise feel overwhelmed by lines of Python. The advantage is that participants can then focus on the live explanations during the session, rather than feeling stressed about running code cells.

This approach aligns with our goal of keeping the event accessible to all, regardless of prior programming experience: the technical threshold is deliberately low. It also underlines the notebook’s sustainability as a resource that participants can return to and adapt after the workshop.

**Situated examples and relevance.** Following the script demonstration with examples of domain-specific applications—vector-based image search and VLM-generated image descriptions—helped situate the method within the realities of heritage collections. Showing how we had already integrated CLIP into an image search demo at the National Library provided a concrete reminder that such approaches are feasible in practice, rather than speculative.

It also prompted relevant questions from heritage professionals during the final discussion. That said, such a “shopwindow” presentation of applications carries a risk of oversimplifying the scale of development work involved—a point we return to below.

### 3.2. What proved difficult

**Lack of interactivity and real-time feedback.** The principal drawback of offering the workshop to a large, technically diverse audience online was that participants tended to take part fairly anonymously, risking becoming passive listeners rather than active learners. The online setting creates a challenging pedagogical situation: talking through and explaining a series of code cells can at times feel like lecturing a brick wall. Given that the schedule was tightly packed, there were perhaps too few opportunities to check in with participants during the script to gauge whether they had questions or reflections.

We encouraged the use of the chat for questions, but this is not equivalent to a classroom experience where instructors can identify and respond to moments of confusion in real time. The lack of interaction was probably accentuated by the fact that participants may not always feel comfortable raising questions in such settings, particularly when the group spans a wide range of skill levels and institutional contexts.

**Limits of single-session workshop.** This challenge of limited interactivity is, to a significant extent, an effect of the format itself. There are strict limits to what can be achieved in a two-hour session—especially one that starts from scratch. The ambition to showcase both the conceptual and practical possibilities of the method necessarily reduced the time available for dialogue and experimentation.

In previous workshops, we have addressed this issue by dividing the material into two sessions, with a “homework” task in between: a first session devoted to theory, an individual exercise in between, and a second session for reflection and discussion. In future iterations, we may consider introducing a similar

structure for this workshop as well—though the trade-off is that it becomes harder to attract participants to commit to multiple dates.

**Risk of appearing “plug and play”.** A final challenge concerns the risk of oversimplifying what such methods entail. The self-playing quality of a Colab notebook can give the impression that multimodal topic modelling is a “plug-and-play” solution: press run, and AI does the rest. In reality, deploying these methods within heritage organisations is rarely frictionless. Adapting them to local infrastructures and workflows requires sustained technical maintenance and coordination across organisational silos. In this sense, the workshop risked presenting AI as a normalised, readily available technology when, in practice, it remains dependent on scarce expertise, time, and institutional support.

A key pedagogical takeaway, therefore, was to demystify not only the models themselves but also the organisational conditions that shape their use in practice. We raised this point in response to a question, but the exchange underscored the importance of addressing it explicitly as a condition of possibility within the session itself.

**Broader reflections.** Taken together, these challenges point to a broader insight: effective AI pedagogy in the humanities depends as much on institutional capacity, infrastructural support, and collaboration as on individual technical skill or curiosity. It is less a matter of training humanities scholars to become AI experts, and more about fostering critical and interpretive literacy—while strengthening the conditions for sustained interaction between researchers, data scientists, and the institutions that mediate between them.

#### 4. Conclusion and further work: towards AI literacy

Designing and running this type of introductory workshop involves a considerable degree of pedagogical risk. Participants arrive with widely varying levels of technical experience; even a minimal Colab notebook can cause cognitive overload for those new to programming; and the very convenience of the format can create an illusion of understanding without deeper comprehension of the underlying mechanics. Acknowledging these risks from the outset, we sought to balance transparency and “data realism” with accessibility and usability.

Our experience affirms a broader insight within the digital humanities learning community: workshops are powerful entry points but weak sustainers. They can ignite curiosity and provide conceptual orientation, yet they remain far removed from the pedagogical structures of sustained practice, repetition, and peer learning required to develop technical proficiency. Nor should we expect them to deliver miracles or shortcuts to understanding. If we instead regard workshops as gateways to longer processes of discovery and engagement, their value becomes clearer: they equip participants with a new vocabulary, methodological awareness, and sense of possibility.

In terms of further work, we are planning several adjustments for the next iteration of the workshop. First, we intend to experiment with a smaller, more interactive format—either as a limited online group with greater opportunities for exchange, or as an in-person teaching event, for instance within a master’s programme in Digital Humanities. Second, we will explore adding short reflective checkpoints within the Colab notebook to prompt active engagement during the session. Building on this, we are also considering a follow-up meeting to provide more time for discussion and cross-participant exchange. Finally, we are considering introducing a post-workshop feedback form to gather more systematic insights into participants’ experiences.

The broader task ahead lies in embedding AI literacy within the interpretive habits of the humanities. This means treating computational methods not as external aids but as evolving objects of inquiry in their own right—tools that both enable and transform the ways we encounter and understand cultural heritage collections. Workshops like this represent a modest but necessary step toward that goal—and toward a more reflective, literate engagement with AI in scholarly and heritage practice.

## References

- [1] L. Jaillant, K. Aske, Are Users of Digital Archives Ready for the AI Era? Obstacles to the Application of Computational Research Methods and New Opportunities, *J. Comput. Cult. Herit.* 16 (2024) 87:1–87:16. URL: <https://dl.acm.org/doi/10.1145/3631125>. doi:10.1145/3631125.
- [2] L. Börjeson, C. Haffenden, M. Malmsten, F. Klingwall, E. Rende, R. Kurtz, F. Rekathati, H. Hägglöf, J. Sikora, Transfiguring the Library as Digital Research Infrastructure: Making KBLab at the National Library of Sweden, *College & Research Libraries* 85 (2024) 564. URL: <https://crl.acrl.org/index.php/crl/article/view/26325>. doi:10.5860/crl.85.4.564, number: 4.
- [3] C. Haffenden, F. Rekathati, E. Rende, Unearthing forgotten images with the help of AI – The KBLab Blog, 2023. URL: <https://kb-labb.github.io/posts/2023-10-20-unearting-forgotten-images-with-the-help-of-ai/>.
- [4] F. Rekathati, BERTopic Workshop: Analyzing Swedish Parliamentary Motions, 2022. URL: <https://colab.research.google.com/drive/10kB3wfoHSfZE48vEKmznIw-ff36uR8gs?usp=sharing>.
- [5] M. Malmsten, V. Lundborg, E. Fano, C. Haffenden, F. Klingwall, R. Kurtz, N. Lindström, F. Rekathati, L. Börjeson, 13 Without Heading? Automatic Creation of a Linked Subject System, in: E. Balnaves, L. Bultrini, A. Cox, R. Uzwysyn (Eds.), *New Horizons in Artificial Intelligence in Libraries*, De Gruyter Saur, Berlin, Boston, 2025, pp. 179–198. URL: <https://doi.org/10.1515/9783111336435-014>. doi:doi:10.1515/9783111336435-014.
- [6] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, I. Sutskever, Learning Transferable Visual Models From Natural Language Supervision, 2021. URL: <https://arxiv.org/abs/2103.00020v1>.
- [7] D. M. Blei, Probabilistic topic models, *Communications of the ACM* 55 (2012) 77–84. URL: <https://dl.acm.org/doi/10.1145/2133806.2133826>. doi:10.1145/2133806.2133826.
- [8] J. Sikora, AI for Image Collections: A Hands-On Workshop with CLIPtopic, 2024. URL: [https://colab.research.google.com/drive/1CMJ0ko8kWA4Zjf\\_qkBA0AN\\_TwDHeRV8t?usp](https://colab.research.google.com/drive/1CMJ0ko8kWA4Zjf_qkBA0AN_TwDHeRV8t?usp).
- [9] P. Broadwell, L. King, Beyond "This Image May Contain...", 2025. URL: <https://thegogglesdonothing.com/projects/cni2025/>.