Lynn Rother • Fabio Mariani • Max Koss

INTERPRETING STRINGS, WEAVING THREADS: STRUCTURING PROVENANCE DATA WITH AI

The provenance of an artwork is the record of its ownership and socio-economic custody changes. Traditionally, provenance was recorded as texts in ledgers. More recently, it has shifted to free text fields in collection management systems in museums. The researching and writing of provenances have long been characterized by a high degree of complexity as well as fuzziness, both of which are now emerging as points of concern for the digitization of provenance. To begin with, there is the very heterogeneous and incomplete historical archives that source provenances. While new archival materials continue to be found and made available to researchers, more often than not, some information remains missing, leading to gaps in provenances. At the same time, the information researchers have at their disposal, previously used to compile and record provenances, is open to interpretation. For each generation of researchers and scholars producing provenance, the historical sources underpinning provenances can be reinterpreted with new perspectives, leading to the updating, rewriting, and reinterpretation of provenances in light of contemporary concerns and conventions. Provenances are generally texts without single authorship, amalgamations of the work of multiple authors active at different moments in time, working to more or less scientific standards.

Today, museums are called upon to structure their provenance records and, eventually, transform them into provenance-linked open data (PLOD), which is based on standards for publishing information as structured data on the web. This enables the interlinking and reusability of any such information and, consequently, the enhancement of shared knowledge. For example, linked open data allows using already vetted data of other institutions and querying datasets across institutions and repositories for complex research questions. While this process has yet to be widely adopted in the cultural heritage field, the advantages of such an approach are clear.

A PLOD approach would allow the identification of objects unlawfully appropriated during contexts of injustice, such as Nazi-era expropriation or colonial looting, serving restitution and decolonization efforts. Large-scale analysis of ownership and socio-economic custody changes may also be relevant for other research questions in such fields as art history, anthropology, sociology, and social and economic history.[1] PLOD will also make provenance as a knowledge practice more accessible, if not more democratic. Where the writing of provenance is still predominantly tied to institutions that often function as gatekeepers of knowledge, a digital provenance approach will transform provenance into a distributed and collaborative knowledge practice. Such an undertaking can potentially counteract the various historical biases, for example, sources or subjective interpretation. Last but not least, the digital future of provenance allows tackling the issue of the authority of provenance, as it provides the possibility to publish the provenance of provenance, clearly identifying authors and sources of every bit of data contained in a given digital provenance.

Today, the digital transformation of provenance faces two interrelated questions addressed in this paper: First, how can vast quantities of provenance texts be transformed into high-quality data, and second, what would such a process need to look like for the benefits of digital transformation to outweigh the efforts and costs required? This paper focuses on the use of artificial intelligence in the transformation of provenance. It lays out which AI techniques are particularly suited for the process and expands on the limitations of a technology-only approach. Indeed, given the complexity and fuzziness mentioned at the outset, it will become clear that the production of digital provenance will continue to require expert knowledge to make judgment calls where the machine cannot.

2. From Provenance Texts to Provenance Data

A look at the current state of provenances, especially in the United States of America, reveals that, ever since the publication of the Washington Conference Principles on Nazi-Confiscated Art in 1998, there have been rigorous research efforts to record the provenance of hundreds of thousands of works across numerous institutions.[2] Although no shared standard was established,

1 Lynn Rother, Max Koss, and Fabio Mariani: Taking Care of History: Toward a Politics of Provenance Linked Open Data in Museums, in: Perspectives on Data, ed. Emily Lew Fry and Erin Canning (Art Institute of Chicago, 2022); Lynn Rother, Fabio Mariani, and Max Koss: Hidden Value: Provenance as a Source for Social and Economic History, in: Economic History Yearbook, Special Issue: Digital Methods, vol. 64, no. 1 (May 2023).

2 U.S. Department of State: Washington Conference Principles on Nazi-Confiscated

Offered by the artist as a New Year's gift to "Mme X." [1] M and Mme Jules Féral, Paris, by 1932 until at least 1938.[2] Possibly (Galerie Charpentier, Paris) in 1951. [3] Capt. Edward H. Molyneux [1891-1974], Paris, by 1952;[4] sold 15 August 1955 to Ailsa Mellon Bruce [1901-1969], New York; bequest 1970 to NGA.

[1]According to Paul Jamot and Georges Wildenstein, *Manet*, Paris, 1932, no. 508.

[2]Lent by Féral to exhibitions in London in 1932 and Amsterdam in 1938. Eugène and Jules Féral [died c. 1949] acted as experts at sales at Hôtel Drout and elsewhere, the former betweeen c. 1876-1901 and the latter in the 1920s.

[3] A 1949 sale of objects from Jules Féral's collection held at the Galerie Charpentier did not include the NGA picture. However the picture was included in an 1951 exhibition held at Charpentier, with no owner listed, and was probably sold to Charpentier by Mme Féral by that time.

Figure 1: The provenance of Édouard Manet's *Flowers in a Crystal Vase*, as published on the website of the National Gallery of Art, Washington, DC (https://www.nga.gov/collection/art-object-page.52181.html, 3 February 2023).

the publication of the American Alliance of Museums (AAM) guidelines on how to write provenance has resulted in many U.S. institutions adopting a similar approach to documenting provenance.[3] In addition, institutions outside the U.S. have adopted comparable guidelines, proposed by the International Foundation for Art Research (IFAR) recommended.[4] Figure 1 shows the provenance of Édouard Manet's *Flowers in a Crystal Vase*, as published on the National Gallery of Art website in Washington, DC.

Recording provenance according to the AAM guidelines involves compiling the chain of provenance events in chronological order up to the acquisition by the current owner. In our example, the first event in the history of any object – the creation of the painting – is omitted. The first recorded event is the gift of the object by its creator Édouard Manet to an anonymous »Mme X«; the last recorded event is the bequest by Alisa Mellon Bruce to the National Gallery of Art, the work's current owner. Each event corresponds to a sentence in the text divided from the previous one by a semicolon when the transfer between the parties was direct. »If a direct transfer did not occur or is not known to have occurred,« then the AAM guidelines suggest dividing the events by a period.[5] For example, this type of gap appears between the ownership of »Mme X« and that of »M and Mme Jules Féral,« the two

Art (https://www.state.gov/washington-conference-principles-on-nazi-confiscated-art/, accessed 3 February 2023).

3 Nancy H. Yeide, Amy L. Walsh, and Konstantin Akinsha: *The AAM Guide to Provenance Research* (Washington, DC: American Association of Museums, 2001).

4 International Foundation for Art Research. Provenance Guide (https://www.ifar.org/Provenance_Guide.pdf, accessed 3 February 2023).

5 Yeide, Walsh, and Akinsha, The AAM Guide, p. 33.

parties named in the first and second event of the provenance text in Figure 1. Footnotes should be used to document historical sources and clarify uncertain events. Terms such as »probably« and »possibly« can be used to indicate hypotheses about events not entirely accounted for. For example, the provenance text in Figure 1 indicates that the ownership of the Galerie Charpentier is considered possible; a footnote explains the reason for this uncertainty, namely, that this work was on view at the gallery in 1951, which does not necessarily imply ownership.

While many institutions have now adopted the AAM guidelines to record the provenance of thousands of objects, a stricter standardization of writing provenance has yet to be achieved. In light of this continued heterogeneity, we cannot consider provenance compiled according to the AAM guidelines to be structured, machine-readable knowledge. At the same time, AAM-compliant provenances can be considered a foundation for employing advanced knowledge extraction techniques to streamline the process of creating and publishing PLOD.

3. The Role of AI in Structuring Provenance Texts

Although provenance texts written according to the AAM guidelines are unstructured, one can use artificial intelligence (AI) to automatically extract information from the text and structure it in a machine-readable format. Because this is a challenge involving texts, the research area for this process is Natural Language Processing (NLP), which develops computational methods that automatically process human language to solve specific problems. One such problem is the extraction of events from a text. In our case, the chronological nature, lining up event after event, helps to extract information from provenance texts.

We have successfully experimented with event extraction from provenance texts by approaching the problem with two NLP tasks.[6] The first task is sentence boundary detection (or disambiguation, SBD). SBD aims to identify and disambiguate punctuation marks that separate sentences in a text. As discussed earlier, events in a provenance text may be separated by a semicolon or a period, depending on whether the change of ownership is direct or not. However, characters such as a period can be ambiguous. For example, a period indicating an abbreviation may or may not mark the end of an event.

6 The experiment is discussed in Rother, Mariani, and Koss, »Hidden Value.«

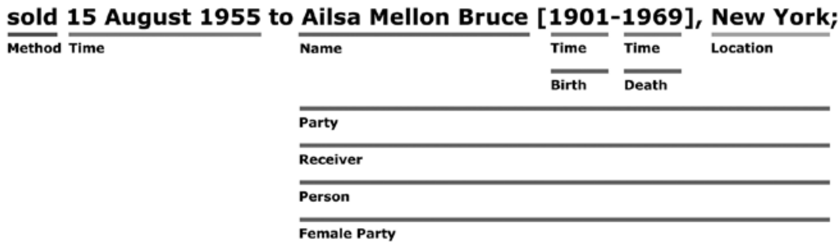

Figure 2: Example of span categorization. The event is taken from the provenance of Édouard Manet's *Flowers in a Crystal Vase*, as published on the website of the National Gallery of Art, Washington, DC (https://www.nga.gov/collection/art-object-page.52181.html, 3 February 2023).

Once we have divided provenance texts into events based on punctuation, we implement the second task identified for event extraction: span categorization (or classification). This task involves identifying and categorizing portions of texts (spans) by assigning them to a category.[7] We can apply span categorization to any event previously extracted by SBD. Figure 2 shows an example of span categorization applied to a provenance event extracted from the provenance of Édouard Manet's *Flowers in a Crystal Vase* discussed previously. The different spans of the text assigned to a category are highlighted. For example, the span »sold« corresponds to the method of transfer used in the event, so »method« is the category assigned to the span. Similarly, the text portion »15 August 1955« represents the »time« when the event occurred.

A distinctive feature of the span categorization task is that spans may overlap. For example, the span »Ailsa Mellon Bruce [1901-1969], New York« can be categorized as »party,« in the form of a »person,« and with its role as »receiver,« in other words, the party who receives the object in this event. In addition, we can indicate the presumed gender of the party by assigning the category »female party.« Finally, within the span, we can find additional spans. In the event annotated in Figure 2, the span »Ailsa Mellon Bruce« is the »name« of the party, »1901« is the date of birth, to which we assign both the »birth« and »time« categories, while »1969« is the date of death, to which

7 For the documentation of the provenance specific annotation scheme, see Fabio Mariani, Lynn Rother, Max Koss: Teaching Provenance to AI: An Annotation Scheme for Museum Data, in: AI in Museums: Reflections, Perspectives and Applications, ed. Sonja Thiel, Johannes Bernhardt (Bielefeld: transcript, 2023), pp. 167-176.

we instead assign both the »death« and »time« categories. Lastly, the span »New York« is the party's location.

We trained two deep learning models to address the SBD task and the span categorization task through AI. Both models performed remarkably well in our experiment on the provenance texts published by the Art Institute of Chicago. The SBD model achieved an F1 score of 0.99, while the span categorization model scored an F1 score of 0.94.[8] Given these results, we can automatically extract information from provenance texts written according to the AAM guidelines with high degrees of accuracy.[9]

The use of AI to extract knowledge from provenance texts reveals promising scenarios for the fast publication of large amounts of data. However, introducing a heuristic process, such as deep learning models, in dealing with historical information requires a critical awareness of the technology used and how it shapes its results. Indeed, behind the output of the AI's black box lie substantial human interventions that influence the heuristic process. Therefore, we must not be tempted by AI's lure of objectivity to accept its results but rather maintain a critical approach in supervising them.[10]

4. Interpreting Strings, Weaving Threads

Despite the satisfactory results that AI models achieve in extracting information from provenance texts, the production of PLOD cannot be considered complete with these computational methods. In fact, two main issues require human intervention when using AI to structure provenance texts on a large scale: First, despite good test performance, techniques such as span categorization are not error-free. Although a low percentage of errors is not statistically significant when analyzing large amounts of data (distant reading), each error becomes noteworthy when analyzing individual provenances published in LOD (close reading).[11] Given the accuracy of the tests, ignoring the low

8 The F1 score is a measure to assess the accuracy of an AI model. Its value is between 0 and 1.

9 We refer to Rother, Mariani, and Koss, »Hidden Value,« for a comprehensive description of the models' implementation and training.

10 The »lure of objectivity« of computational methods is one of the five challenges of the digital humanities presented in Bernhard Rieder and Theo Röhle: Digital Methods: Five Challenges, in: Understanding Digital Humanities, ed. David M. Berry (London: Palgrave Macmillan UK, 2012), pp. 67-84.

11 For a focus on the concepts of close and distant reading in art history, see: Harald Klinke: The Digital Transformation of Art History, in: The Routledge Companion

error rate of AI can be a concern should any of these errors involve legally and ethically problematic provenance events. Consider, for example, potential errors in data extraction for events involving looting or confiscation. Such neglect would go against the principles of transparency and accountability that museums aim to uphold, not least by publishing PLOD. Therefore, it is essential to always monitor the output of AI models to prevent the publication of erroneous historical information.

The second reason for human intervention in AI-extracted data concerns certain types of historical information that require expert interpretation to be recorded and published in LOD in a manner commensurate with their complexity. One can divide this information into four categories: vague, incomplete, subjective, and uncertain. To emphasize that this information requires human supervision, we grouped the four categories under the acronym VISU, from the Latin *de visu*, which translates as ›with your own eyes‹.[12] Vague information concerns approximations of spatial and temporal data, examples of which are expressions such as »near Florence« or »by 1932.« The expert's task in such cases is to evaluate the vague information in the provenance data and reconstruct the information as accurately as possible.

Incompleteness refers to the lack of provenance information, which may occur as a gap when the transfer between two owners is not known to have been direct. As indicated earlier, the AAM guidelines recommend recording such gaps by separating the events with a period.[13] In dealing with such gaps, experts may formulate new hypotheses for what may have occurred by interpreting available historical sources or analyzing already structured provenance data. Indeed, data analysis can support this process, revealing patterns and insights that can help suggest new hypotheses.[14] However, this machine intervention in the historian's hermeneutic approach should not be understood as an automatic process in which the machine generates new hy-

to Digital Humanities and Art History, ed. Kathryn Brown, Routledge Art History and Visual Studies Companions (London: Routledge, 2020), pp. 32-42.

12 Fabio Mariani: Introducing VISU: Vagueness, Incompleteness, Subjectivity and Uncertainty in Art Provenance Data, in: Proceedings of the Workshop on Computational Methods in the Humanities 2022 (forthcoming).

13 Yeide, Walsh, and Akinsha, The AAM Guide.

14 An experiment on the use digital methods and analysis for reconstructing missing art market information is presented in: Matthew Lincoln and Sandra van Ginhoven: Modeling a Fragmented Archive: A Missing Data Case Study from Provenance Research, in: Digital Humanities 2018: Book of Abstracts, ed. Jonathan Girón Palau and Isabel Galina Russell (Mexico City: Red de Humanidades Digitales A.C., 2018), pp. 428-432.

potheses. Rather, data analysis becomes a new research tool for the historian, albeit not exempt from the source criticism required by historiographical methods.[15] Finally, further incomplete information may exist in the components of a provenance event, such as biases in the representation of female parties or minorities. Indeed, it is not uncommon to find female parties in the text recorded with the husband's surname or even with the husband's first and last name, as in the example of »Mme Jules Féral.« In this case, the bias is explicit in the text, and the information is propagated in the data without appropriate human intervention.

Writing a provenance text and supervising the data extracted by AI are both hermeneutic processes that call on domain experts to formulate hypotheses. Individual scholars create information that – while following scientific criteria – is subjective insofar as any use of historical sources is an act of individual interpretation by the domain expert. Any provenance event is recorded following such standard historical practice of interpreting sources; however, as discussed in the Introduction, current provenance writing usually neglects to identify authorship and, to a lesser extent, sources. The intervention of a domain expert on the data extracted by AI enables the reconstruction of the history of an object parallel to the history of documenting its provenance, by whom it was conducted, when, and with what sources.

Documenting this information means recording, in addition to the provenance, the provenance of provenance.[16] This approach enables a further step in the process of professionalizing and raising the scientific profile of provenance research. In particular, the provenance of provenance meets the need for transparency and accountability in museum documentation. It is imperative to record the author, date, and sources used to formulate each piece of information. As discussed earlier, AI and data analysis can assist historians in producing and enhancing provenance data without replacing their role as experts and critics. The provenance of provenance also records the use of computational methods to ensure a transparent account of how provenance

15 Joris J. van Zundert: Screwmeneutics and Hermenumericals: The Computationality of Hermeneutics, in: A New Companion to Digital Humanities, ed. Susan Schreibman, Ray Siemens, and John Unsworth (Chichester, UK: John Wiley & Sons, Ltd, 2015), pp. 331-347.

16 Susanne Al-Eryani, Gudrun Bucher, Stefanie Rühle: Ein Metadatenmodell für gemischte Sammlungen, in: Bibliotheksdienst 52 (2018), pp. 548-564; Christian Huemer: The Provenance of Provenances, in Collecting and Provenance: A Multidisciplinary Approach, ed. Jane Milosch and Nick Pearce (Lanham: Rowman & Littlefield Publishers, 2020), pp. 2-15; David Newbury and Louise Lippincott: Provenance in 2050, in: Collecting and Provenance, pp. 101-109.

data were generated and by whom (or by what AI). Finally, documenting the provenance of provenance enables recording contradictory hypotheses. For example, two historians may disagree on the interpretation of a particular source, drawing different conclusions regarding the life trajectory of an object. In this way, institutions can publish provenance data without discarding one hypothesis in favor of another, recording both hypotheses and documenting their relative provenance of provenance.

Lastly, uncertainty relates to the interpretability of provenance information discussed above. A scholar most certainly has varying degrees of confidence in formulating different hypotheses, which the AAM guidelines also reflect. They suggest using terms such as »possibly« (more confident) or »probably« (less confident), depending on the degree of certainty with which the statement can be made.[17] Historians must therefore evaluate the uncertain information after structuring the provenance data using AI. The certainty of a hypothesis, which is related to the interpretability of provenance information, is thus additional information to be included in the provenance of provenance.

The example of the provenance of *Flowers in a Crystal Vase* is enlightening in demonstrating the importance of human supervision of AI-extracted data, mainly VISU information. Regarding vague information, we note several approximations of dates. For example, the text records that Capt. Edward H. Molyneux had acquired the work »by 1952.« In this case, we need to turn to the last known date before the event in question, which delimits a time interval for locating the vague acquisition date. Since the earlier date is »1951,« we can infer that Molyneux acquired the work between 1951 and 1952. Nevertheless, to validate this inference, we must first consider the incomplete, subjective, and uncertain information in the text: Some of the names of the previous owners are unknown, and there is no record of the name of »Mme X« nor of Jules Féral's spouse, represented as »Mme Jules Féral.« This incomplete information coincides with gaps in the provenance. Indeed, we do not know what happened to the painting once it was given to »Mme X.« The object, created circa 1882, reappears 50 years later as the property of Jules Féral and his anonymized spouse, after which a further gap occurs. In trying to fill this gap, the editor of the provenance text formulated the hypothesis that the Galerie Charpentier might have owned the object. In fact, through a note in the provenance text, we learn that the gallery presented the painting in a 1951 exhibition. The editor of the provenance speculates that Jules Féral's spouse sold the object to Galerie Charpentier by that year. Since

17 Yeide, Walsh, and Akinsha, The AAM Guide, p. 33.

the AI models we introduced previously do not involve extracting knowledge from notes, this valuable information would have been lost without human intervention. However, in this instance, in reconstructing the provenance of provenance and assessing the reliability of such a hypothesis, we must note the omission of an appropriate reference to the historical sources used to formulate a hypothesis.

Through the term »possibly,« the provenance text indicates the uncertainty of the hypothesis that Galerie Charpentier owned the object. The historian supervising the extracted data may accept this uncertainty or engage in further research. For example, an archival search might turn up new documents related to purchases made by Galerie Charpentier, filling the gap. Further help for the historian might come from analyzing other provenance data. For example, one could analyze the data and identify the main parties who sold objects to Galerie Charpentier, particularly whether there were other instances of Jules Féral's widow selling objects. However, we need an appropriate record of the parties' names to facilitate this analysis. In the case of »Mme Jules Féral,« it is necessary to record this person's name for her proper historical representation and consistency across provenance data. Indeed, we might have cases where »Mme Jules Féral« is recorded as »Mme Féral« or »Mrs. Jules Féral.« The analysis could be even more arduous if, for instance, Jules Féral had more than one wife.

Based on the considerations and decisions we introduced, a historian could finally assess whether Molyneux acquired the work between 1951 and 1952, thus accepting that Galerie Charpentier acquired the object in 1951. Otherwise, discarding this hypothesis, one might infer that Molyneux could have acquired the object between 1938, the last date when Jules Féral and his spouse owned the object with any certainty, and 1952 when there is the certainty that the object was already in his possession.

5. Conclusion

This paper discusses how the provenance of museum objects can be (semiautomatically) structured and published. By leveraging the power of artificial intelligence, in particular deep learning models, we can process large quantities of data relatively quickly. Nonetheless, when dealing with the qualitative nature of much historical information, we found it necessary to consider human intervention to monitor the results of AI. This approach is essential for error correction and appropriately handling VISU information. Thus, we developed a two-step, two-speed digitization process: fast digitization, enabled by AI and its quantitative benefits, followed by and combined with slow digitization, performed by domain experts, evaluating and ensuring the scientific quality of the data.

The domain expert is not replaced by technology but becomes an essential factor in the digitization process. The historian need not participate in the time-consuming data structuring process, which AI can successfully perform. Instead, the expert is involved in critiquing sources and formulating historical hypotheses. This demarcates a precise boundary between the tasks delegated to AI and the tasks appropriate for domain experts. After all, history is written by humans, not machines.