

## Annotating Upstream: Digital Scholars, Art History, and the Interoperable Image

Matthew J. Westerby, National Gallery of Art, US, [m-westerby@nga.gov](mailto:m-westerby@nga.gov)

---

Written primarily from the position of an art historian engaged in digital research and data-intensive projects, this essay explores annotations on interoperable images as sites of knowledge production. Images and descriptive metadata can be presented and re-presented in any number of contexts on web pages, but annotations targeting the data sources for such presentations are contextual fragments of scholarly insights that do not translate easily across domains. While data models for web annotation are clearly defined in a technical sense, their implementation is socially motivated and constrained. This essay gives a very brief overview the ecosystem of IIIF annotations, sandbox projects, and Open Access initiatives at art museums and libraries as well as tools for digital scholars to reuse and remanifest IIIF content. I suggest that art historians should practice “annotating upstream” by maintaining the data that constitute the target and body of their annotations while also acknowledging the sociotechnical affordances and ephemerality of annotation spaces.

---



In October 2023, a slew of digital services hosted by the British Library became unavailable, the result of a widely reported cyberattack. Digital surrogates of books and manuscripts on the institution's web pages were suddenly inaccessible. This was partially a result of the loss of British Library systems for International Image Interoperability Framework (IIIF) services (IIIF, 2024). When the British Library's IIIF API endpoints became unresponsive, other 'downstream' uses of the institution's IIIF manifests and images were impacted. The loss of these IIIF endpoints threw light on a relatively recent yet fundamental change in the relationship between institutions and users of online collections, in which separate layers exist to serve images and also to present the digital surrogate of a real-world object, like a book. Recognizing this significant change in web publishing is now the urgent concern of scholars, as highlighted by Suzanne Paul, Benjamin Albritton, and other panelists in 'Now You See It, Now You Don't: Sustainable Access in a Digital Age', organized by *Digital Medievalist* and moderated by Laura Morreale (*Digital Medievalist*, 2024).

Although IIIF is widely implemented by large institutions, the field of art history is only starting to adjust to the paradigm of interoperability and reusability it represents. In a wider sense, and as a specialized field within the humanities, art historical practice is suddenly responding to the 'era of digital reproducibility' characterized by Béatrice Joyeux-Prunel (2024). In this paradigm, images and descriptive metadata can be reused and decentralized—assuming content is released under an open, share-alike license, which may not be the case for images of modern and contemporary art. In this essay I reflect on annotating interoperable images as art historical practice. Although self-intentioned implementations of IIIF are currently classified as the sole concern of digital scholars, specialists, and technologists, all scholars in the field of art history engage interoperable images when they locate sources online and view or annotate digital collections published by libraries, archives, and museums. IIIF Canvases, Manifests, and Collections, as the key interpretive fields of the IIIF API specifications, represent a great deal of knowledge and expertise.

Aside from technical specifications, the implementation of annotation spaces is defined by social circumstances. In this essay I briefly survey the ecosystem of IIIF annotations as outgrowths of open access initiatives led by institutions and of sandbox projects mostly led by individuals. I argue that digital scholars and art historians should be mindful of the architectures they engage and the situatedness of the annotations they author. Rather than relying on IIIF systems managed by institutions, scholars can strive to manage the annotations they author as both data services and research outputs, thereby 'annotating upstream'—which I define here as publishing and maintaining one's own annotations. This may also mean publishing and maintaining the digital

repository on which annotations are targeted, like IIIF manifests. I return to this notion through sections on reinscribing knowledge as a practice, the creation of meaningful manifests as a component of digital storytelling, and tools for social annotation.

### **Annotating Upstream**

Institutional authority is central to the perception of the accuracy of information about museum and library collections. Art historians will continue to look to collecting institutions as authoritative sources. At the same time, data provenance is an important factor in a decentralized Linked Open Data (LOD) context, and for this reason digital scholars will prefer to reuse IIIF manifests created by institutions, especially if a scholar does not have the means, ability, or time to generate their own manifests. The IIIF annotation layer, however, can be seen as a distinct content space made up of contextual fragments of scholarly insight and which comprise a semantic relationship that can be broken when taken out of context. Institutions may not include persistent identifiers that enable later reuse, such as UUIDs (Universally Unique Identifiers) for IIIF canvases. In the technical sense, the relationship of a IIIF annotation is defined by the Web Annotation Data Model (W3C, 2017), at least for the current IIIF Presentation API 3.0 specification (IIIF-C, 2020). Institutions could implement systems to track changes and make transparent the history of a remanifested object. Identifiers, like ARKs, can also be applied so that when an institution remanifests a IIIF object, a reference to the previous identifier is maintained, ensuring some degree of continuity and traceability.

Acting upstream, and more specifically as the source of the annotation and target, foregrounds the digital scholar's location in a continuous process of art-historical interpretation. Quoting Emily Pugh's overview of domains and systems of information management through institutional change and scholarly practice, successive acts of interpretation fundamentally rely on 'how a book or document is described, classified, catalogued, and accessed', which 'has ripple effects that are felt downstream, at the point when a patron is searching a collection' (2020: 48). Considering a IIIF manifest as a data product in the same category as a search service accessed by a patron through a web page, this ripple effect extends to the point when a patron authors a web annotation about an item in the form of a IIIF manifest or collection.


For the field of art history, the practice of annotating upstream is also an intervention to shape systems of knowledge. By intentionally engaging the phenomenon of IIIF annotation space as a locus of discourse, art historians can avoid unknowingly adopting a data model designed by technologists who cannot possibly envision every potential future reuse or reinterpretation of an item. Following a

provocation later raised in Pugh’s essay on ‘Art History Now’, annotating upstream can also be seen as an imperative if art historians want ‘to continue to produce valuable and significant insights about artworks, artists, and ideas about art’ (2020: 56). These insights might take the form of creative exhibitions driven by IIIF annotations or content shared across online social annotation communities, as well as peer-reviewed essays and formal writing. They might also be outputs of computational analysis. I conclude by addressing a scenario in which a proliferation of user- and computer-generated annotations on a single canvas might be considered as a site for qualitative analysis, in which many layers of annotations target the same source, which I describe as a form of ‘thick data’.

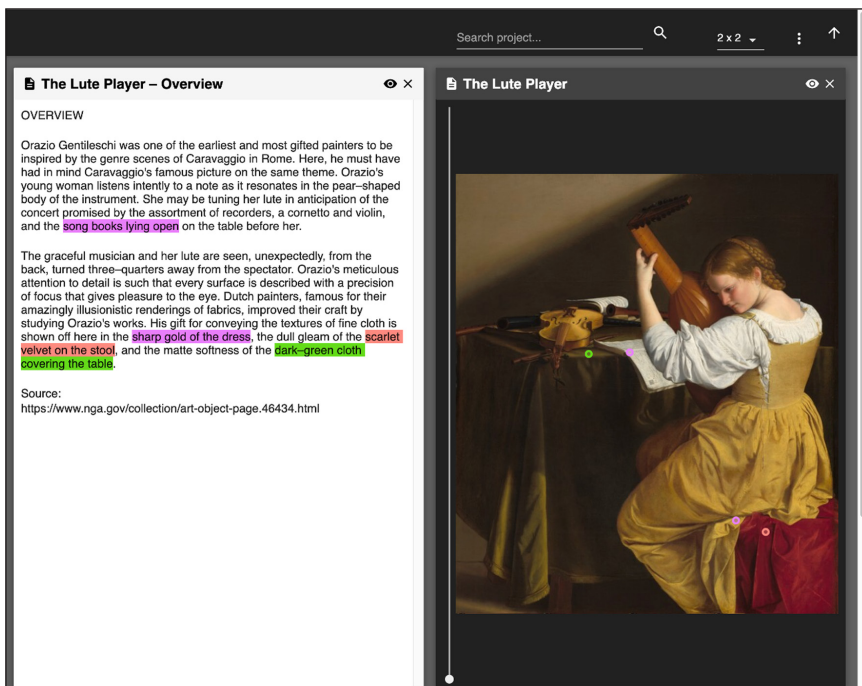
### Reinscribing Knowledge

A remark about an object can exist entirely independent from the object. One enduring practice in art history is to reinscribe written texts over or adjacent to images of objects. In this sense, descriptions, essays, and data systems are all means to represent reality. Considered as a system of pointers, they ‘provide a structured understanding of that reality within a particular bounded context’ (Langmead and Newbury, 2020: 367). A description might be reinscribed on a digital surrogate of a material object in the form of an annotation. This also applies to the original IIIF Shared Canvas Data Model (Sanderson and Albritton, 2013). One non-IIIF example reuses a translated text attributed to Suger, abbot of Saint-Denis (1081–1151 CE), describing colored bands of stone around an ancient bowl in the form of a digital annotation over a photogrammetric 3D model (**Figure 1**). As rendered in the Sketchfab viewer, the annotation relays how ‘the red sard of the stone stands out, in contrast to the blackness of onyx’ in context to a surrogate for the real-world object, over 800 years since this description was recorded (Cusimano and Whitmore, 2018: 115).

To demonstrate annotation spaces as sites for reinscribing texts with images of artworks, overviews written for web pages about museum objects can be reinscribed on an interoperable image. An example using a public domain description and image of Orazio Gentileschi’s *The Lute Player* at the National Gallery of Art shows how this can be done in Digital Mappa 2.0, a workspace for digital scholarship projects that lets users easily import a IIIF manifest URL and begin annotating collections with color-coded text highlights (**Figure 2**). As open-source software, the resources needed to spin up an instance of Digital Mappa are a relatively low barrier for creating a web-based annotation sandbox for research, teaching, or publication.



**Figure 1:** Screenshot showing a single point digital annotation with text body over a user-generated photogrammetric 3D model of the Chalice of the Abbot Suger of Saint-Denis, 2nd/1st century BCE (cup); 1137–1140 (mounting), National Gallery of Art, 1942.9.277. Created with Polycam and viewed with Sketchfab viewer. Image: author.



**Figure 2:** Screenshot showing colored annotations in Digital Mappa 2.0 linking words describing elements of a painting to a IIIF canvas with an image of Orazio Gentileschi, *The Lute Player*, c. 1612/1620, National Gallery of Art, 1962.8.1. Image: author.

In this sense, annotations in IIIF content space reinscribe knowledge. Rather than relying on an institution to maintain a given data service (like a IIIF endpoint) that defines the content space, digital scholars can reuse Open Access images and metadata provided by an institution as a starting point to define their own annotation spaces. Tanya Clement and Liz Fischer describe this scenario, where manifests are ‘created or copied and reshared by readers who may wish to reorient how [an] object is presented online. By referencing or creating links to only the tiniest detail of an object such as an image or an audio file (the brightest star in Van Gogh’s “Starry Night” or one phrase in a poem spoken by Maya Angelou), that reader can create a manifest that reorients completely how an object is read or accessed’ (2021: 14). For annotations in IIIF space, a manifest produced by a downstream user would ideally use the same IIIF canvas ID generated by a museum or library for continuity of context. Even if a locator for this IIIF canvas ID is not provided, such as a DOI or ARK reference URL, the UUID assigned to that IIIF canvas could be matched or referenced in the future. This might be seen as a form of digital forensic art historical research.

Digital scholars that reuse IIIF manifests generated by institutions should now be aware that IIIF images, like websites, can disappear at any time. Users of these data services may not know exactly at what point they are entering an ecosystem of information. Although systems can be decentralized, users of cultural heritage data may prefer to reference institutional services because they are seen as authoritative, and in this sense the user is downstream of the institution. The value of any given IIIF manifest differs from web pages in that web pages do not provide a data service for an unforeseen downstream digital exhibition or other kind of web presentation (as does a manifest file for a client-side IIIF viewer). If a digital project relies on an API endpoint to reference annotation data and that endpoint disappears, the ability of the computer to render the semantic relationship is broken. Annotations authored by end users that target an institutional IIIF manifest are inherently risky if there is any expectation of permanence. For example, a museum might remanifest (that is, create a new manifest) for an object in their collection as a new digital object with a different IIIF canvas identifier. Here again a URL locator like a DOI or ARK for a IIIF canvas ID emerges as an important step towards linking interoperable content spaces.

Interoperable images and manifests generated by the Internet Archive may prove to be more durable over time than those generated by museums and libraries but similarly rely on institutional oversight. In 2023, the Internet Archive began to officially support the hosting of IIIF images uploaded by users, seeking to serve an average internet user who, in the words of Josh Hadro, ‘may not benefit from the same access to or affiliation with infrastructure offered by traditional research institutions’ (Internet Archive, 2023). As a comparison, Wikipedia does not yet



publish IIIF manifests. Overall, this means that most IIIF canvases and manifests are ephemeral, which implies that a given IIIF manifest has ‘no established authority as a definition of a content-space’, as described by Tom Crane, a contributing author of the IIIF technical specifications (2018). At the same time, the way in which museums and libraries expose IIIF manifests as features of IIIF viewers on web pages for online collections encourages their direct reuse in other applications. Aggregators of collections images and data are one approach to sharing, and I will turn later in this article to consider the example of the aggregator Europeana. Like the British Library cyberattack, the Internet Archive has faced repeated cyberattacks and was inaccessible in October 2024, at the time this essay was submitted for publication.

Institutions can also benefit by sharing and decentralizing digital collections. This may also be seen as a type of distributed collections security strategy. Following Dot Porter, collections sharing means that digital collections from one institution are shared and hosted by another institution, ‘in a way not dependent on the systems of the original institution’ (2024). Utilizing decentralized protocols like IPFS (InterPlanetary File System) for distributed storage and blockchains to provide unique identifiers is one strategy towards wider decentralization.

The desire to implement ideals like reusability for the continuity of context can run counter to the demands placed on students and scholars. Institutions supporting STEM fields may be more likely to adhere to FAIR data principles (Findable, Accessible, Interoperable, and Reusable). One example is the Distributed System of Scientific Collections (DiSSCo), built on the FAIR Digital Object Framework (Islam et al., 2023). In the domain of digital humanities and digital art history, a lack of funding or expertise to implement such ideals can relegate FAIR principles to the status of an optional feature, with the exception of large grant funded research projects which are required to draft and implement data strategies. Assigning UUIDs to IIIF canvases is one way to follow FAIR principles and is a practice I discuss later in this essay.

For art history and humanities, FAIR principles may not address the criteria most needed in the future to interrogate our current discourses about the past. CARE principles (Collective benefit, Authority to control, Responsibility, Ethics) are proposed for Indigenous data sovereignty (Carroll et al., 2020). Béatrice Joyeux-Prunel advocates for adding Ethics and Expertise, Source mention, and Timestamp to extend FAIR to FAIREST (2024: 2.5.2). Considered as more than mere metadata, a source mention and timestamp may help to translate across systems of reinscribed knowledge. In this sense, digital scholars place themselves and their interpretations in a time, place, and system of ethics. Wikidata can also be seen as a source of reinscribed knowledge in as much as Wikidata is not a source of the past but a model of discourse for how we discuss and conceive of the past (Joyeux-Prunel, 2024: 2.8),

where timestamps and usernames are tracked in a Wikidata record version history. A given user's 'talk' page may, for example, document the ethics and expertise this user has demonstrated. As a site for reinscribing knowledge and documenting a source and a time, digital scholars must consider the provenances of the IIF manifests they engage with.

### Meaningful Manifests and Storytelling

Counter to the argument for upstream annotation I have focused on so far, there are many reasons to directly reuse a IIF manifest generated by an institution. Cultural heritage institutions are also responding to a need for ways to discover manifests representing their collections. One example is the Simple IIF Discovery tool developed at the National Gallery, London (Padfield et al., 2022: 24). Museum manifest URLs can also be shared as LOD in open knowledge systems. The National Gallery of Art in Washington, for example, piloted a project in 2018 to donate Open Access images and data to Wikimedia Commons and Wikidata (Zweig, 2022), creating Wikidata statements for artworks that include IIF manifest URLs to make it more likely that digital scholars would discover and reuse the museum's manifests. Because the National Gallery of Art donated images and collections data to the Wikimedia Commons under a Creative Commons 0 (CC0) license, digital scholars can copy the structure into their own products and annotate these IIF manifests, as described by Clement and Fischer, where annotations exist in an 'unauthorized, distributed, and decentralized system' (2021: 4). Discoverability of museum and library manifests increases opportunities for digital scholars to practice upstream annotation.

The descriptive values—or metadata—within a IIF manifest often reproduces the same information about an artwork or object as a web page. To avoid duplication, web developers Mat Jordan and Mark Baggett designed a static-site generator Canopy IIF to build websites from IIF manifests and collections, envisioning the IIF manifest as the HTML '<main>', or the body of the web page (Jordan, 2022). This is one example of a shift towards creating more intentional manifests, not just more data outputs. According to Mat Jordan, the by-product of this focus on the front end is that if an institution's manifests are 'robust enough to represent our works on our own web presences, then we can have confidence that they are ready for in-the-wild consumption and use by researchers and the greater public' (2022). Designed for discovery-focused digital exhibits, Canopy IIF also contributes to a movement towards minimal computing in the humanities, exemplified in the Wax workflow for building digital exhibitions with static web pages developed by Marii Nyrop.



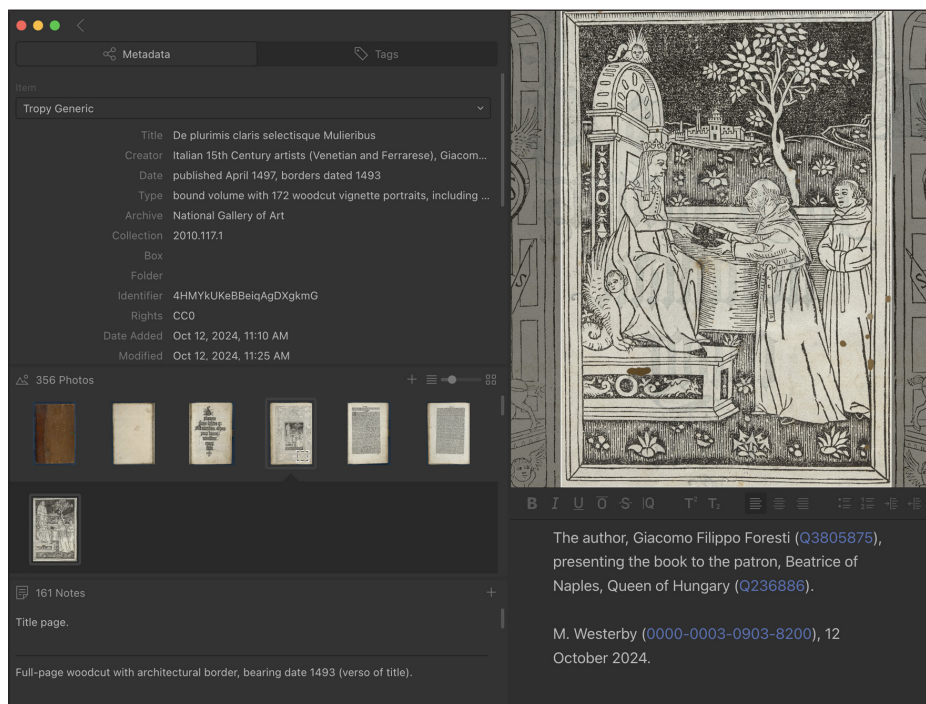
I have argued to this point that art historians and digital scholars should produce IIIF resources and not rely only on institutional data services. In other words, scholars should define a corpus and layer their interpretations in a continuous content space as other users whenever possible. Annotations can also be seen as socially produced data about interpretation, and their degree of currency is reliant on the ability of users to comment on the same source, or to comment on an earlier annotation. This locus of discourse in IIIF is the ‘annotation space’ (Crane, 2018), bearing pointers to the things represented by an interoperable image.

Story-driven digital exhibitions with interoperable images are also driven by annotations and contribute to the potential locus of discourse by framing regions of images and associating a written description as a narrative practice. One example is Exhibit, developed by Mnemoscene and based on Universal Viewer. Exhibit is an easy-to-use web app with many applications for writing visual narratives or creating interactive instructional materials, such as quizzes. However, unless a digital scholar or institution deploys a self-hosted instance of Exhibit, the annotations over IIIF canvases that drive the text and image relationships in a ‘scrollytelling’ story are stored to Mnemoscene’s Exhibit web servers.

Other web-based storytelling tools require editors to save these outputs of their intellectual work to a GitHub user account, and these decisions have deep implications. Annotate, for example, developed by Niqui O’Neill, enables scholars to deploy GitHub workspaces for IIIF annotations without needing to launch an IIIF annotation server. In an interview, O’Neill stated they ‘decided to use the annotation urls as a data source to make the tool usable on a wider range of images and to reduce user labor’ (2019). O’Neill further notes that ‘if the manifest had been used as the data source, external users would have had to create and host their own separate manifest’. Here again, the IIIF manifest emerges as a locus of discourse that remains obscured to many art historians.

Considering the issues raised it can be intimidating for scholars to envision practical methods to remanifest interoperable images as the ‘upstream’ source. Software plugins and helpful guides intended for digital scholars are one starting point. Plugins for the archival research tool Tropy allow users to import and export IIIF manifests. For example, a set of Open Access (CC0) images of a 15th-century book imported to Tropy on a digital scholar’s computer can be enriched with metadata and notes, which can then be mapped to IIIF Presentation and exported to a Github user profile to be published as a static web page. Notes written in the Tropy interface can be exported with the free and open-source plugin tropiiify, created by the ARKA digital studio. Starting local, an art historian (like myself) can identify subjects depicted in

the digital surrogate of a book, linking to Wikidata Q identifiers and to the ORCID identifier of the annotation author (**Figure 3**). Tropy notes are exported in HTML as annotations referenced inside the IIIF manifest exported with troppiify. Step-by-step guides and helpful context for using the troppiify plugin are shared in blog posts by Raffaele Messuti (2024) and Shawn Graham (2024).



**Figure 3:** Screenshot showing Tropy interface with a note identifying the subjects in a woodcut by name and Wikidata Q identifier, with a line of text identifying the name of the annotation author with ORCID identifier and date. Italian 15th Century artists (Venetian and Ferrarese) and Giacomo Filippo Foresti da Bergamo (author), *De plurimis claris selectisque Mulleribus*, published April 1497, borders dated 1493; bound volume with 172 woodcut vignette portraits. National Gallery of Art, 2010.117.1. Image: author.

If a IIIF manifest is imported to start this process, the creator/remanifestor should consider best practices for reusing images and metadata. Messuti and Graham both suggest steps to generate unique identifiers for each item in your Tropy collection that you intend to publish. If an imported manifest already provides a unique identifier or UUID for IIIF canvases, this identifier could be re-used or referenced. The IIIF Cookbook offers scenarios and examples for implementors, yet the re-use of manifests, or parts of manifests, is an evolving practice. Assuming the remanifestor uses Github to publish their new manifest, any changes to the manifest JSON file will also be tracked and discoverable.

## Social Annotation and Thick Data

Taking account of the range of IIF manifests and annotations from institutions to individual scholars, a comparative look at upstream annotation and social annotation platforms sheds light on the gap between the ideals of open annotation models and the demands of business models. Although the tools for individuals to manage annotations as data are relatively accessible, as in the previous example with Github, sociotechnical realities encourage most people on the web to create free user accounts on multiple platforms that generate closed content spaces. Social annotation is a feature of digital publishing platforms like Manifold, where readers can share highlights with reading groups which can be toggled on or off in the user's reading interface. Similar features are bundled into teaching and instruction platforms like Canvas. Encouraging conversation and discussion across groups of users is a common feature of these platforms.

Other web technologies open the content space to any domain on the internet by implementing the Web Annotation Data Model (WADM) (W3C, 2017). This is the goal of the social annotation tool called Hypothesis, supported by the public benefit corporation Anno, which aims to 'bring content and conversation together with a layer of social annotation' (Hypothesis, 2023). With a browser extension and a single user account, users can annotate what they see and read on the internet and share annotations with defined groups.

Assuming digital scholars share annotations in the widest possible sense, one can envision a scenario where a proliferation of user- and computer-generated annotations referring to a single IIF canvas can be analyzed as 'thick data'. This term is adapted for researchers engaging big data for ethnographical study, recognizing the idea of 'thick description' in the work of anthropologist Clifford Geertz. Thick description refers to 'an account that interprets, rather than describes' (Moore, 2018: 56 citing Geertz, 1977). Elaborated by Paul Moore, a 'thick data' approach shows that the ways in which data are used is a cultural rather than a technological problem, emphasizing that 'all technologies are ultimately subject not only to the needs of the user but also to the context in which they are being used' (2018: 52). Here Moore's insights parallel Alison Langmead and David Newbury's foundational interpretation of pointer- and proxy-based models of the phenomenal world and the socially situated use of computer systems, noting that 'tacit and explicit contexts [...] cannot be moved to other contexts without damage to their intellectual coherence and subsequent impact' (2020: 367).

In practice, this kind of 'thick data' for IIF annotations needs a centralized platform. The ability to analyze large groups of human and computer-generated

annotations as an interpretive practice relies on users contributing to a collective annotation space that can be referenced, queried, or mapped. The potential for these repositories to address art historical research questions with the ‘thick data’ of annotation spaces as a locus of discourse only beginning to emerge beyond special interest groups. Like the Hypothesis web annotation tool, and similarly based on the Web Annotation Data Model, the Europeana Annotations API is one annotation space that allows for creating, retrieving, and managing annotations targeted on Europeana objects. One can imagine future research questions that consider the number or type of annotations on a given IIF manifest in a quantitative or qualitative approach, and a narrative about the cultural significance of museum objects could be framed in this way. This parallels emerging approaches to analyzing published articles as data, or as machine-readable writing about art (Klinke, 2023).

The recent explosion of content created by Generative AI may be seen to blur the line between ‘thick data’ and ‘data pollution’. The IIF community is beginning to respond to this question. Allison Sherrick and Diego Pino Navarro have proposed including ‘no-AI’ or ‘regulated AI’ use license tags in IIF manifests to prevent AI bot harvesting; they also suggest tagging AI-generated outputs, like image descriptions or annotations pointing to image regions, to identify their method of creation (2024). In art historical discourse, ‘data pollution’ is a familiar phenomenon, admittedly on a far more limited scale. One example are scholarly attributions for makers of artworks following the Western European tradition of connoisseurship that posits notnames for anonymous masters. The notnames upheld by the scholarly consensus of curators and researchers are used as data points in vocabularies like the Getty Union List of Artist Names (ULAN), but notnames cannot always be resolved with observable facts. For example, in the Wikidata model, ‘anonymous master’ (Q474968) may have the conflicting statements ‘instance of’ (P31) the entity ‘human’ (Q5) and ‘instance of’ the class ‘notname’ (Q1747829). Unlike the observable fact that a human once lived and acted, the names of anonymous masters are generated by and for academic discourses about styles and schools and may not be grounded in observable or reproducible fact. As such, they represent a form of circular scholarly discussion that generates conflicting statements—or statements about statements—that are obscured further as they are repeated in scholarly literature. AI products powered by LLMs (Large Language Models) trained on data scraped from the internet, when they are used to generate annotations within the same content space, represent a different but relatable phenomenon.

Although verbosity is avoided in data science, it is nonetheless a common feature of knowledge production in art history. CRMaaa (Art and Architectural Argumentation Conceptual Reference Model), an unofficial extension to the CIDOC CRM standard, is

designed to represent historical data which is not otherwise easily modeled (ORDEA, 2023). CRMaaa introduces ‘institutional fact’ and ‘speech act’ to represent socially grounded actions as historical facts. A similar approach might be used to represent computer-generated content published in the form of IIIF annotations. As grounded in scholarly discussion, this type of model introduces near infinite complexity and may not be suited for implementation by institutions, where the integrity and authority of metadata is the primary mission. To address computer-generated annotations, institutions may instead implement strategies that anticipate or respond to risks beyond their control, such as providing tags to identify authors, creation dates, and desired uses for the content and annotations they publish, as suggested in the extension of FAIR to FAIREST (Joyeux-Prunel, 2024: 2.5.2). In this way, institutions and digital scholars manifesting and remanifesting images and metadata take an ethical stance in the content spaces they create (cf. Sherrick and Pino Navarro, 2024), even if they lack agency to restrict or police malicious uses of their interoperable collections (such as IIIF manifests and canvases).

## Conclusion

Art historians and digital scholars can take ownership of their insights and storytelling outputs by practicing upstream annotation, preferring to target IIIF canvases they create and host or to target images uploaded to free-to-use centralized platforms like the Internet Archive or decentralized but locatable by persistent identifiers. Adopting best practices, like timestamping the layers of interpretations, can further contribute to valuable uses of shared canvases and meaningful manifests. This practice will not in itself lead to better preservation of scholarly annotation content but it may influence digital scholars to be mindful of their outputs. IIIF manifests created by museums and libraries with open access content can become starting points for an ongoing social process of reinscribing and generating knowledge, rather than existing as the products of a given institution’s API endpoint.

---

## Acknowledgments

This essay was developed from an outline of thoughts I presented at the roundtable discussion 'The Quotable Text in the Digital Age: Images, Words, and Music' at the 2023 Renaissance Society of America conference. I am deeply grateful to Angela Dressen for organizing the RSA roundtable and for editing this special collection. I am also indebted to the reviewers whose comments and suggestions significantly improved the form and content of this essay and to many colleagues in person and online for collaborations and conversations around IIF, including the IIF Community #annotations and #museum Slack channels.

## Competing Interests

The author has no competing interests to declare.

---

## References

- Annotate.** <https://dnoneill.github.io/> [Last Accessed 14 July 2024].
- ARK.** Archival Resource Key. <https://arks.org/> [Last Accessed 19 October 2024].
- Canopy IIF.** <https://canopy-iif.github.io/docs/> [Last Accessed 14 July 2024].
- Canvas.** <https://www.instructure.com/canvas> [Last Accessed 19 October 2024].
- Carroll, S R, Garba, I, Figueroa-Rodríguez, O L, Holbrook, J, Lovett, R, Materechera, S, Parsons, M, Raseroka, K, Rodriguez-Lonebear, D, Rowe, R, Sara, R, Walker, J D, Anderson, J, and Hudson, M** 2020 The CARE Principles for Indigenous Data Governance. *Data Science Journal*, 19: XX, pp. 1–12. <https://doi.org/10.5334/dsj-2020-043>
- CIDOC CRM.** International Council for Documentation Conceptual Reference Model. <https://www.cidoc-crm.org/> [Last Accessed 19 October 2024].
- Clement, T E and Fischer, L** 2021 Audiated Annotation from the Middle Ages to the Open Web. *Digital Humanities Quarterly*, 15(1). <https://www.digitalhumanities.org/dhq/vol/15/1/000512/000512.html> [Last Accessed 6 April 2024].
- Crane, T** 2018 Beyond the Viewer: fragments and links in annotation space [*Medium*]. 19 June. <https://medium.com/digirati-ch/beyond-the-viewer-fragments-and-links-in-annotation-space-b3284e25f34> [Last Accessed 7 April 2024].
- CRMaaa.** <https://github.com/swiss-art-research-net/crmaaa> [Last Accessed 19 October 2024].
- Cusimano, R and Whitmore, E** 2018 *Selected Works of Abbot Suger of Saint Denis*. Catholic University of America Press. <https://doi.org/10.2307/j.ctt1zkjxb9>
- Digital Mappa.** <https://www.digitalmappa.org/> [Last Accessed 14 July 2024].
- Digital Medievalist** 2024 'Now You See It, Now You Don't: Sustainable Access in a Digital Age' [*Bluesky*]. 31 January. <https://bsky.app/profile/digitalmedievalist.bsky.social/post/3kkbmukocc52f> [Last Accessed 6 April 2024].
- Europeana.** <https://www.europeana.eu/> [Last Accessed 14 July 2024].
- Exhibit.** <https://www.exhibit.so/> [Last Accessed 14 July 2024].



- Geertz, C** 1977 *The Interpretation of Cultures*. London: Basic Books.
- Graham, S**, 2024 Make Your Tropy Collection and Annotations Available With Canopy. <https://electricarchaeology.ca/2024/07/25/make-your-tropy-collection-and-annotations-available-with-canopy/> [Last Accessed 19 October 2024].
- Hypothesis** 2023 About. <https://web.hypothes.is/about/> [Last Accessed 8 April 2024].
- IFPS**. InterPlanetary File System. <https://ipfs.tech/> [Last Accessed 19 October 2024].
- IIIF** 2024. International Image Interoperability Framework. <https://iiif.io/> [Last Accessed 14 July 2024].
- IIIF-C** 2020 IIIF Presentation API 3.0. <https://iiif.io/api/presentation/3.0/> [Last Accessed 8 April 2024].
- IIIF Cookbook**. <https://iiif.io/api/cookbook/> [Last Accessed 19 October 2024].
- Internet Archive** 2023 'Making IIIF Official at the Internet Archive'. 18 September. <https://blog.archive.org/2023/09/18/making-iiif-official-at-the-internet-archive/> [Last Accessed 12 November 2024].
- Islam, S, Beach, J, Ellwood, E R, Fortes, J, Lannom, L, Nelson, G, and Plale, B** 2023 Assessing the FAIR Digital Object Framework for Global Biodiversity Research. *Research Ideas and Outcomes*, 9. <https://doi.org/10.3897/rio.9.e108808>
- Jordan, M** 2022 The (IIIF) Front End Theory: Envisioning the IIIF Manifest as the HTML <main>. [Medium]. 24 June. <https://medium.com/@mathewrjordan/the-front-end-theory-e35abc580f14> [Last Accessed 6 April 2024].
- Joyeux-Prunel, B** 2024 Digital humanities in the era of digital reproducibility: towards a fairest and post-computational framework. *International Journal of Digital Humanities*. <https://doi.org/10.1007/s42803-023-00079-6>
- Karpeles, M** 2023 [Internet Archive Blogs] Making IIIF Official at the Internet Archive. 18 September. <https://blog.archive.org/2023/09/18/making-iiif-official-at-the-internet-archive/> [Last Accessed 8 April 2024].
- Klinke, H** 2023 Articles As Data. Machine-Readable Writing. [Medium]. 2 May. <https://medium.com/@HxxxKxxx/articles-as-data-machine-readable-writing-55c4110adac9> [Last Accessed 7 April 2024].
- Langmead, A and Newbury, D** 2020 Pointers and Proxies: Thoughts on the Computational Modeling of the Phenomenal World. In: K Brown (ed.) *The Routledge Companion to Digital Humanities and Art History*. New York and London: Routledge. pp. 58–373. <https://doi.org/10.4324/9780429505188-31>
- Manifold**. <https://manifoldapp.org/> [Last Accessed 14 July 2024].
- Messuti, R**, 2024 Building a simple IIIF digital library with Tropy, Tropiify and Canopy. <https://literarymachin.es/iiif-tropy-canopy/> [Last Accessed 19 October 2024].
- Moore, P** 2018 From Big Data to Thick Data: Theory and Practice. In G. Schiuma and D. Garlucci (eds.) *Big Data in the Arts and Humanities*. New York: Routledge. pp. 51–61. <https://doi.org/10.1201/b19744>

- O'Neill, N** 2019 Open Source Annotations on IIF resources: An Interview with Niqui O'Neill. [Digital Humanities Collaborative of North Carolina]. 30 July. <https://dhcnc.org/project/open-source-annotations-on-iif-resources-an-interview-with-niqui-oneill/> [Last Accessed 9 April 2024].
- ORDEA**, 2023 Art and Architectural Argumentation Conceptual Reference Model (CRMAaa). <https://www.sari.uzh.ch/en/ordea/crmaaa.html> [Last Accessed 19 October 2024].
- Padfield, J, Bolland, C, Fitzgerald, N, McLaughlin, A, Robson, G, and Terras, M** 2022 Practical applications of IIF as a building block towards a digital National Collection. <https://doi.org/10.5281/zenodo.6884884>
- Porter, D** 2024 Collection Sharing as a Form of Digital Collections Security [Dot Porter Digital]. 19 March. <https://www.dotporterdigital.org/collection-sharing-as-a-form-of-digital-collections-security/> [Last Accessed 10 April 2024].
- Pugh, E** 2020 Art History Now: Institutional Change and Scholarly Practice. *International Journal for Digital Art History*, 4: 47–58. <https://doi.org/10.11588/dah.2019.4.63448>
- Sanderson, R, and Albritton, B** 2013 Shared Canvas Data Model. <https://iif.io/api/model/shared-canvas/1.0/> [Last Accessed 19 October 2024].
- Sherrick, A K and Pino Navarro, D A** 2024 Creating a better balance: the need for tools and practices to combat AI harvests and resource flooding in repository environments. Open Repositories 2024 (OR2024), Gothenburg, Sweden. Zenodo. <https://zenodo.org/records/12579304> [Last Accessed 19 October 2024].
- Sketchfab**. <https://sketchfab.com/> [Last Accessed 14 July 2024].
- Tropiify**. <https://github.com/arkalab/tropiify> [Last Accessed 19 October 2024].
- Tropy**. <https://tropy.org/> [Last Accessed 19 October 2024].
- Universal Viewer**. <https://universalviewer.io/> [Last Accessed 14 July 14 2024].
- W3C** 2017 Web Annotation Data Model (WADM). 23 February. <https://www.w3.org/TR/annotation-model/> [Last Accessed 6 April 2024].
- Wax**. <https://minicomp.github.io/wax/> [Last Accessed 19 October 2024].
- Zweig, B** 2022 Opening up: The National Gallery of Art's Wikimedia project. *Journal of Digital Media Management*, 10(4): 358–367. <https://www.ingentaconnect.com/content/hsp/jdmm/2022/0000010/0000004/art00009> [Last Accessed 19 October 2024].

