Journal of Open Humanities Data

ARTigo: Data from Social Tagging with Art-historical Images

DATA PAPER

STEFANIE SCHNEIDER

ubiquity press

ABSTRACT

The ARTigo dataset, generated from over 10 million annotations, is a product of a citizen science project developed by the Institute of Art History and the Institute of Computer Science at LMU Munich. The project leverages *Games with a Purpose* (GWAPs) to foster a playful environment for tagging artworks. In these GWAPs, two anonymous players are given an image to annotate with textual or visual descriptors within a limited time frame. The annotations serve to improve the accessibility of art-historical images and offer vast research potential well beyond their utility as training datasets for *Computer Vision* (CV) algorithms.

CORRESPONDING AUTHOR:

Stefanie Schneider

Institute of Art History, LMU Munich, Munich, Germany

stefanie.schneider@itg.uni-muenchen.de

KEYWORDS:

crowdsourcing; social tagging; serious games; cultural heritage

TO CITE THIS ARTICLE:

Schneider, S. (2024). ARTigo: Data from Social Tagging with Art-historical Images. *Journal of Open Humanities Data*, 10: 54, pp. 1–6. DOI: https://doi.org/10.5334/johd.247

1 OVERVIEW

REPOSITORY LOCATION

Zenodo (https://doi.org/10.5281/zenodo.7763018).

CONTEXT

ARTigo[1] is a citizen science project developed jointly by the Institute of Art History and the Institute of Computer Science at LMU Munich. The project represents a playful approach to the collaborative tagging of artworks, leveraging the "wisdom of crowds" (Surowiecki, 2004) to collect art-historical data. It has two main objectives:

1. Through citizen science, the project intends to educate the public about art history and foster a sense of community and shared responsibility for cultural heritage. Its playful, inclusive approach helps to democratize art-historical knowledge and challenges the exclusivity traditionally associated with the field.
2. Despite recent advances in *Computer Vision* (CV), the automatic classification of art-historical objects remains unreliable. By harnessing human interpretation and observation, the project intends to provide more accurate metadata for those objects, thereby improving their accessibility and discoverability.

As of 1 September 2024, the project has collected extensive data, with 10,367,128 annotations on 66,465 images that were submitted by 56,345 players.[2] It is ongoing, with thousands of new tags being added every month. Unlike museum annotations, which are often leveraged in CV datasets and focus mainly on primary motifs (e.g., Mao et al., 2017), ARTigo's taggings extend to both the central and peripheral details of the image. For instance, while annotations such as "cupid," "venus," and "bees" are employed for Lucas Cranach the Elder's *Venus with Cupid the Honey Thief* (c. 1580–1620; Figure 2a) in the Metropolitan Museum of Art's *Tagging Initiative*,[3] ARTigo further enriches the artwork by the crowdsourced terms "grove," "castle," and "rock"—in addition to describing the figure of Venus with "veil," "incarnate," and "collar."[4]

The project collaborates with institutions dedicated to the dissemination of art and art history: in 2019, ARTigo was featured at the Austrian *Museum of Folk Life and Folk Art* with an installation of large-scale touchscreen walls. In addition, ARTigo has maintained a partnership with the *Belvedere Museum* in Vienna since 2023. For its sustained contributions to the field of citizen science, ARTigo received an honorary mention at the 2023 *European Union Prize for Citizen Science*.[5]

2 METHOD

STEPS

Data Collection

On the ARTigo platform, volunteers generate annotations by participating in *Games with a Purpose* (GWAPs). In these GWAPs, two anonymous players are given an image. Within a limited time frame, the player's task now is to annotate the image with textual or visual descriptors—points are awarded for *matching* annotations with those from previous sessions or with the current opponent's annotations. While the matching of textual annotations is based on *identity*, i.e., exact lexical matches between the words or phrases provided, the matching of visual annotations is based on the *overlap* of the selected image regions. The *Graphical User Interface* (GUI) of the GWAPs, as shown in Figure 1a, includes the image on the left and a sidebar for input on the right, topped by a progress bar displaying the time remaining per round. The sidebar emulates a conversational interface, where the opponent's input appears

1 https://www.artigo.org/ (accessed November 25, 2024).

2 See the appendix for detailed statistics of the dataset.

3 https://collectionapi.metmuseum.org/public/collection/v1/objects/459077 (accessed on November 25, 2024).

4 Tags have been translated from German. The original tags are: "hain," "burg," "felsen," "schleier," "inkarnat," and "halsband."

5 https://ars.electronica.art/citizenscience/en/winners2023/ (accessed on November 25, 2024).

sequentially, with non-matching annotations initially blurred. Depending on the game mode, questions in the sidebar may encourage players to annotate not only the visible objects, but also the image's emotional context. Each session concludes with a comparative display of the results (Figure 1b).[6]

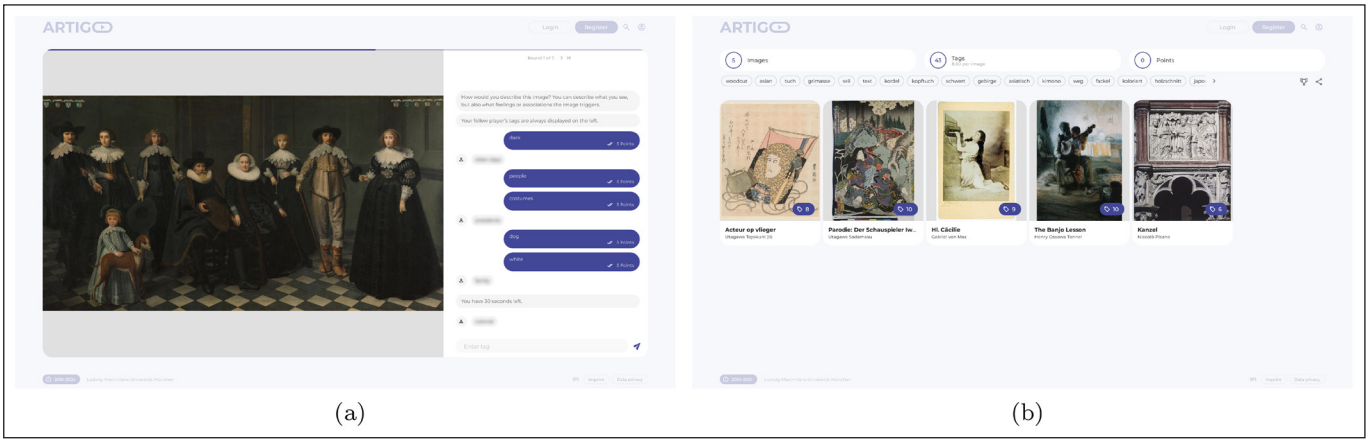

(a) (b)

Figure 1 Front-end design of ARTigo. **(a)** The GUI of the GWAP is divided into two sections, with the image for annotation on the left side and the sidebar on the right. **(b)** Each session concludes with a comparative display of the aggregated results.

Data Aggregation

For publication, the dataset is aggregated to retain only essential metadata (`creators`, `titles`, `location`, `institution`, and `source`) along with the annotations provided for each image, as depicted in Figure 2 by Cranach's *Venus with Cupid*; all personal identifiers are removed to preserve the anonymity of the players. Textual annotations are provided with details of the `language` used during their creation[7] and the `count` of times each tag has been applied. Visual annotations are provided as rectangular bounding boxes, described by the `x` and `y` coordinates of the bottom-left corner and the dimensions (`width` and `height`). That is, each bounding box is defined by normalized coordinates, with values ranging between 0 and 1, describing the relative position and dimensions of the box within the image.

(a) Image


```
{
    "id": 15741,
    "hash_id": "0d56860451ad3bf3ae902ac4ac91a67c",
    "titles": [
        {
            "id": 43513,
            "name": "Venus und Amor als Honigdieb"
        }
    ],
    "creators": [
        {
            "id": 6537,
            "name": "Lukas Cranach"
        }
    ],
    "location": "New York",
    "institution": "The Metropolitan Museum of Art",
    "source": {
        "id": 4,
        "name": "Kunsthistorisches Institut der Universitaet zu Koeln",
        "url": "http://khi.phil-fak.uni-koeln.de/"
    },
    "path": "https://api.artigo.org/media/0d/56/0d56860451ad3bf3ae902ac4ac91a67c.jpg",
    "tags": [
        {
            "id": 2,
            "name": "akt",
            "language": "de",
            "count": 3
        },
        ...
    ]
}
```


(b) Metadata

Figure 2 Each line in the JSONL file represents a single image **(a)** and its associated metadata **(b)**, as exemplified by Lucas Cranach the Elder's *Venus with Cupid the Honey Thief* (c. 1580–1620), with the key-value pairs shown.

Quality control

The dataset is provided "as is," including all spelling errors made by the participants during each gameround. Automatic correction has been deliberately avoided to prevent introducing additional errors that can easily occur when dealing with specialized art-historical vocabulary.[8]

6 See Schneider et al. (2023) for a detailed overview of the GUI.

7 This may not always correspond to the tag's content; e.g., English tags created in the German version of ARTigo are labelled as `German`.

8 See Levushkina (2014) for techniques to correct and normalize the ARTigo taggings.

DOI: 10.5334/johd.247

Furthermore, it was considered valuable to preserve these spelling errors for research purposes, as their analysis could provide insights into tagging behavior that automatic corrections would otherwise obscure. For use cases where high annotation accuracy is required, we recommend setting a threshold and only considering those annotations that meet or exceed the specified `count` value.

3 DATASET DESCRIPTION

REPOSITORY NAME

Zenodo

OBJECT NAME

`data.jsonl` and `media.zip`

FORMAT NAMES AND VERSIONS

The images are stored in a ZIP file structured into directories named by the first two characters of each image's `hash_id`. Within these directories, subfolders named after the next two characters of the `hash_id` contain the image files, which are named using their full `hash_id` with a .jpg extension. The annotation data is provided in a JSONL file, where each line encodes metadata for a single image (Figure 2b).

CREATION DATES

Since 2008-01-01

DATASET CREATORS

The dataset is generated from annotations provided by both anonymous and registered volunteers via the ARTigo platform. Since 2008, the development of the platform has been driven by the following contributors: Matthias Becker, Martin Bogner, Fabian Bross, François Bry, Caterina Campanella, Laura Commare, Silvia Cramerotti, Katharina Jakob, Martin Josko, Hubertus Kohle, Fabian Kneißl, Thomas Krefeld, Maximilian Kristen, Elena Levushkina, Stephan Lücke, Alessandra Puglisi, Anke Regner, Christian Riepl, Clemens Schefels, Corina Schemainda, Eva Schmidt, Stefanie Schneider, Gerhard Schön, Klaus Schulz, Franz Siglmüller, Bartholomäus Steinmayr, Florian Störkle, Iris Teske, Ricarda Vollmer, and Christoph Wieser. All contributors were associated with the LMU Munich (Munich, Germany) during the development phase of the project.

LANGUAGE

German, English, and French

LICENSE

CC BY-SA 4.0

PUBLICATION DATE

Since 2021-07-27. The data are updated monthly.

4 REUSE POTENTIAL

The ARTigo dataset provides extensive research opportunities for training CV algorithms. By enabling automated tagging of image collections, it facilitates natural-language search, transforming access to and interaction with large-scale image repositories (Springstein et al., 2021). Efforts to improve the accessibility of collections based on ARTigo tags have been pursued by several researchers: Tänzel (2017) enhanced artwork retrieval systems by merging textual annotations with low-level visual features, while Moosburger (2017) automatically extracted color palettes. In other machine-learning applications, Schneider & Kohle (2017) applied

clustering techniques to categorize tagged artworks beyond traditional genre boundaries. Poetis et al. (2022) found that misattributions in paintings associated with Claude Monet and William Turner correspond to discernible stages in their artistic careers. Both Schneider & Kohle (2017) and Poetis et al. (2022) demonstrate that, although citizen science initiatives such as ARTigo primarily yield so-called *surface tags* (Wieser et al., 2013), the compilation of many of these overly descriptive tags often reveals underlying art-historical themes. Thus, even superficial annotations can be valuable in retrieval contexts where detailed art-historical expertise is not required, only the ability to identify visible components.

The dataset holds untapped potential for several fields of art-historical research. The annotations could be leveraged to develop machine-learning models capable of detecting and classifying the emotional content of artworks, thereby providing an understanding of how artists convey emotion—and how these expressions have evolved over time. In addition, the dataset is well-suited for analyzing temporal and cultural trends. By utilizing the ARTigo taggings, researchers could explore the evolution of artistic themes and motifs, revealing shifts in visual elements that occur across different cultures and historical periods. With a publicly available RESTful *Application Programming Interface* (API) conforming to the *OpenAPI 3.0* specification (OpenAPI Initiative, 2017), ARTigo could also be leveraged independently of the ARTigo front-end, allowing institutions to customize the GWAPs. Moreover, the platform's source code is freely distributed under the *GNU General Public License* (GPL),[9] so that the methods employed are made transparent and reproducible.

APPENDIX

As of 1 September 2024, the project has gathered a total of 10,367,128 annotations from 56,345 players, with 6.79 % registered, across 873,471 gamerounds. The majority of the annotations—99.79 %—are textual, consisting of words or phrases. The remaining 0.21 % are visual annotations of selected image regions, which have only been collected from 2022 onwards in designated GWAPs. The project's image repository, as elaborated in Schneider & Kohle (2017), contains 66,465 digital reproductions, spanning artworks from the 15th century BC—such as Thutmosis III's *Karnak, Temple of Amun*—to modern works like Franz Marc's *Fighting Forms* (1914). The selection is influenced by the academic priorities of LMU Munich's faculty, with a focus on 17th- and 19th-century art. Each artwork in this repository is assigned a median of 8 gamerounds, receiving 91 taggings and 58 unique tags; 7.09 % of artworks have fewer than 11 taggings. The tagging process is expressed as a triple (*u, r, t*), where a player *u* annotates a resource *r*, the artwork presented in the game, with a tag *t*.

FUNDING INFORMATION

This project was partly funded by the German Research Foundation (DFG) under project number 161260864.

COMPETING INTERESTS

The author has no competing interests to declare.

AUTHOR CONTRIBUTIONS

Stefanie Schneider: Data Curation, Formal Analysis, Investigation, Methodology, Project Administration, Writing – original draft.

AUTHOR AFFILIATIONS

Stefanie Schneider orcid.org/0000-0003-4915-6949
Institute of Art History, LMU Munich, Munich, Germany

9 https://github.com/arthist-lmu/artigo (accessed on November 25, 2024).

REFERENCES

Levushkina, E. (2014). *Computerlinguistische Methoden in Community-basierten Anwendungen*. Hamburg: Dr. Kovač.

Mao, H., Cheung, M., & **She, J.** (2017). DeepArt: Learning joint representations of visual arts. In Q. Liu, et al. (Eds.), *Proceedings of the 2017 ACM on multimedia conference, MM 2017* (pp. 1183–1191). New York: ACM. https://doi.org/10.1145/3123266.3123405

Moosburger, M. (2017). *Colour labelling of art images using colour palette recognition* (Bachelor's Thesis, LMU Munich). Retrieved November 25, 2024, from https://www.en.pms.ifi.lmu.de/publications/projektarbeiten/Matthias.Moosburger/PA_Matthias.Moosburger-high-comp.pdf

OpenAPI Initiative. (2017). *OpenAPI*. Retrieved November 25, 2024, from https://spec.openapis.org/oas/v3.0.0 (Version 3.0.0).

Poetis, P., Radmacher, E., Smiatek, K., & **Schneider, S.** (2022). "Wertlose" Taggings und ihr Nutzen für die Kunstgeschichte. In M. Geierhos (Ed.), *8. Tagung des Verbands Digital Humanities im deutschsprachigen Raum, DHd 2022* (pp. 225–229). https://doi.org/10.5281/zenodo.6304590

Schneider, S., & **Kohle, H.** (2017). The computer as filter machine: A clustering approach to categorize artworks based on a social tagging network. *Artl@s Bulletin*, *6*(3), 81–89. https://doi.org/10.5282/ubm/epub.41319

Schneider, S., Kristen, M., & **Vollmer, R.** (2023). Re: ARTigo: Neuentwurf eines Social-Tagging-Frameworks aus funktionalen Programmbausteinen. In A. Busch & P. Trilcke (Eds.), *9. Tagung des Verbands Digital Humanities im deutschsprachigen Raum, DHd 2023* (pp. 173–178). https://doi.org/10.5281/zenodo.7688632

Springstein, M., Schneider, S., Rahnama, J., Hüllermeier, E., Kohle, H., & **Ewerth, R.** (2021). iART: A search engine for art-historical images to support research in the humanities. In H. T. Shen, et al. (Eds.), *Proceedings of the 2021 ACM on multimedia conference, MM 2021* (pp. 2801–2803). New York: ACM. https://doi.org/10.1145/3474085.3478564

Surowiecki, J. (2004). *The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business, economies, societies and nations*. New York: Doubleday.

Tänzel, T. (2017). *Measuring similarity of artworks using multidimensional data* (Master's thesis, LMU Munich). Retrieved November 25, 2024, from https://www.en.pms.ifi.lmu.de/publications/diplomarbeiten/Tobias.Taenzel/MA_Tobias.Taenzel.pdf

Wieser, C., Bry, F., Bérard, A., & **Lagrange, R.** (2013). Building an artwork search engine with games and higher-order latent semantic analysis. In *Proceedings of the first AAAI conference on human computation and crowdsourcing* (pp. 15–20). Menlo Park, CA: AAAI Press. https://doi.org/10.1609/hcomp.v1i1.13060

TO CITE THIS ARTICLE:
Schneider, S. (2024). ARTigo: Data from Social Tagging with Art-historical Images. *Journal of Open Humanities Data*, 10: 54, pp. 1–6. DOI: https://doi.org/10.5334/johd.247

Submitted: 30 September 2024
Accepted: 20 November 2024
Published: 10 December 2024

Journal of Open Humanities Data is a peer-reviewed open access journal published by Ubiquity Press.