

Visual Analysis of Cylinder Seal Impressions and Unwrappings: Annotation and Automation

Lara Bampfield

Faculty of Asian and Middle Eastern Studies, University of Oxford

Jacob L. Dahl

Faculty of Asian and Middle Eastern Studies, University of Oxford

Abhishek Dutta

Visual Geometry Group, Department of Engineering, University of Oxford

Andrew Zisserman

Visual Geometry Group, Department of Engineering, University of Oxford

Abstract: Cylinder seals are ancient artefacts whose impressions and digital unwrappings contain rich figurative scenes but lack spatial metadata identifying the motifs within them. We first describe the process of manually annotating motifs in a large digital collection of cylinder seal imagery using software tools, and then describe two applications enabled by these annotations. The first application is a quantitative analyses of motif layout and co-occurrence patterns, supported by a bespoke visualisation tool that aggregates motif distributions across seals. In the second application, we investigate the use of deep learning for automating motif detection. Using 2,491 manually annotated examples, we fine tune an EfficientDet object detector to recognise human related motifs (e.g. king, priest, deity, worshipper, dancer) in seal impressions and unwrapping. Our results demonstrate that combining structured manual annotations with modern computer vision tools can significantly enhance the curation, analysis, and accessibility of cylinder seal collections at scale.

Keywords: cylinder seal unwrapping, object detection, manual annotation, annotation visualisation

§1. Introduction

§1.1 Cylinder seals are ancient artefacts commonly made of stone, cylindrical in shape and containing written characters or figurative scenes or both. The CDLI has been collecting high resolution digital images of impressions and digital unwrappings of cylinder seals, in collaboration with museums worldwide, which allow researchers to study them as an archaeological artefact (see Dahl, Lafont, and Ouraghi 2019). These images have associated global metadata (e.g. excavation site, date) as discussed in Englund 2014 but lack spatial metadata specifying motifs (e.g.

king, animal, throne, crescent) within the scenes. Without such information, it is difficult to analyse the depicted motifs and their relationship. To address these limitations, this paper makes three contributions to cylinder seal research.

- Manual Annotation (Section 2): We provide software tools and describe a process for manually annotating visual content in cylinder seal images.
- Annotation Visualisation (Section 3): We present a custom tool for exploring the spatial distribution and co-occurrence of motifs

in seal impressions and unwrappings.

- Automatic Motif Detection (Section 4): We introduce a computer vision-based model that automatically detects "person" motifs (e.g. king, priest, deity, worshipper, and dancer).

§1.2 The software tools and online demos described in this paper are available as open source software (Section 5).

§2. Manual Annotation

§2.0.1 Manual annotation and curation of images showing cylinder seal unwrapping involves defining regions (e.g. rectangular boxes) and describing them with textual labels (e.g. motif names). These annotations provide valuable insight into the figurative content of the seals, including motif co-occurrence and spatial relationships as will be seen in Section 3.

§2.0.2 In this work, the VIA software tool described in Dutta and Zisserman 2019 was used to manually annotate 469 cylinder seal impressions from the British Museum, primarily dated to the Old Babylonian period (c. 1800-1600 BC). The manual annotation process was undertaken primarily by Lara Bampffield as part of her doctoral research on cylinder seal imagery, with methodological and technical support from Abhishek Dutta and Jacob Dahl. Annotation required specialist subject knowledge in order to distinguish motifs, iconographic variants, and damaged or ambiguous figures. While the principal annotation campaign was completed through approximately three weeks of concentrated full-time work, the wider process extended over an iterative research period involving vocabulary refinement, quality control, and subsequent model testing. In total, the broader annotation and development workflow represented approximately 200 hours of specialist labour. This demonstrates that, while the workflow is scalable, expert supervision remains important for producing high-quality training data. For each annotated region, three attributes were defined, with values drawn from the Getty AAT vocabulary defined by Osley, Savidge, and White 1991.

- **Classification of region type** : classify the content of the region in one of the following three categories: (a) cuneiform sign, (b) inscription, or (c) motif.
- **Identification of region content** : describes the content of the region (e.g. animal, bird,

dog, king, lion, star, throne, etc.)

- **Description at fine-grain** : specifies properties of the region (ascending, female, kneeling, male, naked, sitting, small, standing, etc.)

§2.0.3 Whereas Getty AAT vocabulary terms are not always hierarchically structured in a way that fit our purpose we have entered them into a hierarchy of classifying, identifying, describing, and assessing; for a couple of terms that do not have Getty AAT terms we have used subject specific terminology (such as smiting). The three attributes provide a structured and interoperable framework for describing the visual and textual content of cylinder seal imagery. The *classification of region type* separates inscriptions, cuneiform signs, and motifs, reflecting the fundamental components that recur across seal unwrappings. Such distinction supports later processing steps where inscriptions and imagery may require different analytical treatment. The *identification of region content*, using controlled Getty AAT terminology, ensures that motifs are labelled with consistency across seals and collections. Standardised vocabulary is crucial for comparative study, particularly in large corpora where small inconsistencies in wording can obstruct quantitative analysis of motif frequency, distribution, and co-occurrence. The *fine-grain descriptive attributes* record specific morphological or behavioural features such as posture, action, gender, and scale. These descriptors capture nuances that are often significant for understanding workshop practice, compositional conventions, and variation within motif categories. Encoding these attributes in a structured manner allows them to be incorporated into subsequent spatial mapping, clustering, or machine-learning tasks. Together, the three attributes create a coherent annotation scheme that balances archaeological nuance with computational clarity, enabling reproducible and scalable analysis of cylinder seal imagery.

§2.0.4 Figure 1 shows a screenshot of the VIA annotation tool being used to define and describe regions containing various motifs (e.g. animal, throne) in a cylinder seal unwrapping.

§2.0.5 The manual annotations created using VIA are stored in a structured JSON file. Each annotation records the spatial coordinates in the field named "shape_attributes" and its descriptive metadata in "region_attributes". This structured

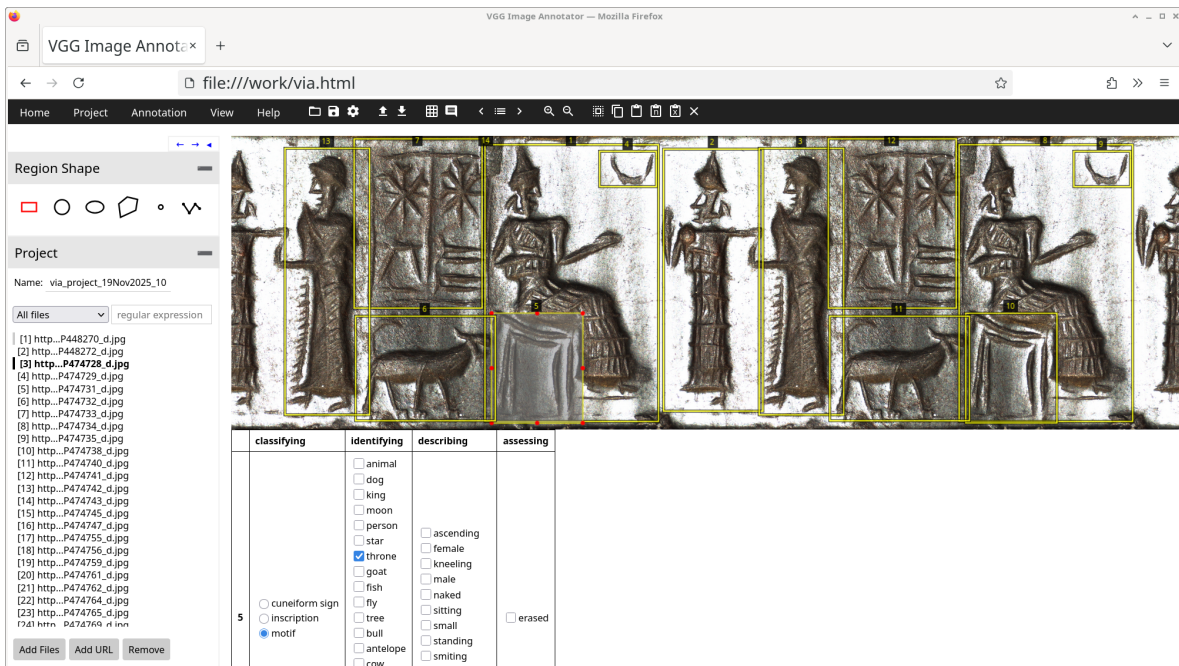


Figure 1: Manual annotation of the unwrapping of P474728 a cylinder seal in the British Museum using the VIA software tool. Regions (yellow boxes) can be drawn by clicking and dragging the mouse button. Region attributes (e.g. motif name or "identifying" attribute) can be defined by selecting a region (e.g. region 5 with red dots in the boundary) and selecting an appropriate Getty AAT vocabulary term (e.g. throne) to describe the content of the selected region .

format facilitates interoperability with other software tools as well as contributing to human understanding.

§2.0.6 As an example, the following snippet shows the JSON metadata for the seal unwrapping image P474728_d.jpg. It defines a rectangular region of size 585 × 932 pixels with its top-left corner at coordinates (856, 29). The region is tagged with metadata motif and people for the "classifying" and "identifying" attributes, respectively.

```
"_via_img_metadata": {
  "1": {
    "filename": "https://cdli.mpiwg-berlin.m
    → pg.de/dl/photo/P474728_d.jpg",
    "size": -1,
    "regions": [
      {
        "shape_attributes": {
          "name": "rect",
          "x": 856,
          "y": 29,
          "width": 585,
          "height": 932
        },
        "region_attributes": {
          "classifying": "http://vocab.getty
          → .edu/aat/300009700",
          "identifying": {
```

```
"http://vocab.getty.edu/aat/3003
→ 43850": true,
"http://vocab.getty.edu/aat/3000
→ 24979": true
},
"describing": {
  "http://vocab.getty.edu/aat/3002
→ 63970": true
}
},
// ... (other regions in this image)
}
// ... (other images in this project)
}
```

§2.1. Motif Spatial Distribution Analysis

§2.1.1 Manual annotations of cylinder seal unwrapping images enable analysis of motif spatial distributions across the dataset, revealing preferred layout patterns for specific motifs. To demonstrate the potential of such analysis, we examine the spatial distribution of six motifs: disc, crescent, throne, ball-staff, pot, and footrest.

§2.1.2 Because cylinder seals are continuous objects, unwrappings have no natural beginning or end. Any quadrant-based analysis therefore, re-

quires the imposition of a start coordinate. This was determined according to observable features in the scene rather than randomly. Where figures were shown approaching a seated deity or king, the start point was positioned so that the scene unfolded towards that focal figure. In most cases, scenes were further aligned to read from left to right to ensure consistent comparison across the corpus; the position of accompanying inscriptions, which frequently appear at the right-hand edge, often supported this orientation. In scenes without a clear focal break, such as repetitive or circular designs, a fixed boundary was applied consistently across the dataset. All start points were assigned manually and are treated as methodological constructs rather than inherent features of the artefact.

§2.1.3 The start point defines the top left corner of the unique rectangular image region contained in an image showing two complete (*i.e.* 720 degree) rotation of a cylinder seal. The height of this unique region covers the full height of the seal image. The width of this unique region can be determined computationally as follows. A source patch region (*e.g.* 31 × 31 pixels) in the image is compared against other target patches of same size in the horizontal direction of the seal image. Similarity between source and target patches is quantified by taking the sum of absolute difference in pixel values. The target patch region with largest similarity with the source patch is used to define a candidate width of the unique region. More candidate widths can be obtained by selecting source patches at other locations. A majority voting method is used to decide the most likely width of the unique region in a cylinder seal image. The unique image region contained in each seal image is divided into four quadrants as shown in Figure 2 using the image centre coordinates to define the boundary between these quadrants. Table 1 reports the number of motif instances whose bounding box centres fall within each quadrant. The results show that footrest, ball-staff, and throne motifs predominantly appear in the lower quadrants (Quadrant 3 and 4), whereas disc, crescent, and pot motifs occur mainly in the upper quadrants (Quadrant 1 and 2).

§2.1.4 Analysing the spatial distribution of motifs, especially when connected to archaeological metadata, allows specialists a new and refined overview of this very large dataset. It alleviates bias (especially so-called connoisseurship,

still common in studies of ancient Mesopotamian cylinder seals) by allowing for sampling very large datasets with minimal subjective impact and at speed. Finally, it introduces enhanced proof and reproducibility into the study of motifs in cylinder seals.

§3. Annotation Visualisation

§3.1 The *Annotation Explorer* is a software tool developed to visualise the spatial distribution of motifs and their relationship with the spatial distribution of co-occurring motifs. To enable comparative visualisation, all the cylinder seal unwrapping images and their annotations are resized to a common width while preserving their aspect ratio, aligning them in a shared two-dimensional space. Although the seals vary in height, they share the same width because the unwrapping represents a full 360° rotation around the seal's central axis. For example, Figure 3 shows the spatial positions (blue circles) of bounding box centres corresponding to the motif "deity" across multiple seal unwrappings. This visualisation indicates that the "deity" motif commonly appears in the vertical central region, with frequent occurrences also on the left and right sides. Hovering the mouse cursor over any blue circle displays the original cylinder seal unwrapping image along with the manually drawn bounding box around the corresponding motif, as shown in Figure 4.

§3.2 The *Annotation Explorer* can also display multiple motifs simultaneously, using colour to distinguish between them. This visualisation helps to identify whether certain object categories consistently appear in specific regions of a cylinder seal as compared to other motifs. For example, Figure 5 shows the spatial locations of the "animal" (blue) and "deity" (green) motifs. This visualisation indicates that the "deity" motif typically appears in the vertically central region of the seal, whereas the "animal" motif can occur across a wider range of positions.

§3.3 The visualisations produced by the *Annotation Explorer* facilitate a scale of analysis that has traditionally been difficult to achieve in glyptic studies, where research often relies on close reading of individual seals or stylistically defined sub-groups. By aggregating spatial coordinates across hundreds of annotated unwrappings, the tool makes visible recurrent compositional structures that might otherwise remain implicit. For example, clear horizontal and vertical patterning is observable in motifs such as the disc, crescent,




Figure 2: The unwrapped cylinder seal is split into four quadrants using horizontal and a vertical lines passing through the midpoint of the unwrapped cylinder seal image.

Motif	Quadrant 1	Quadrant 2	Quadrant 3	Quadrant 4
disc	9	102	2	0
crescent	21	179	7	3
throne	1	0	70	2
ball-staff	1	1	30	39
pot	35	20	3	3
footrest	0	0	22	0

Table 1: Spatial location of various objects (e.g. disc, throne) contained in the unique region of unwrapped images. The table gives the number of instances of each motif (or object) that were manually annotated and that appears in one of the four quadrants, as defined in Figure 2.




Figure 3: **The annotation explorer view 1.** Spatial locations (shown using blue circles) of all manually annotated instances of the object category "deity". All the cylinder seal unwrapping images and their annotations are resized to a common width while preserving their aspect ratio thereby allowing visualisation of all instances of a motif in a common two dimensional space.

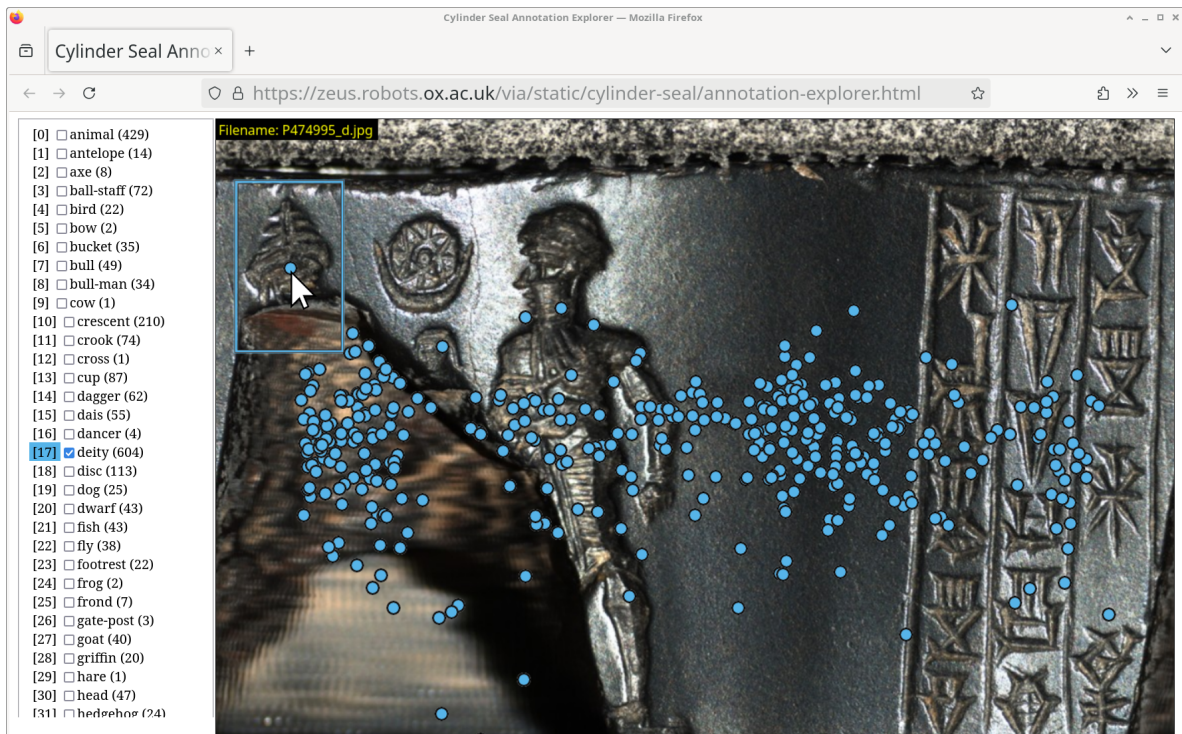


Figure 4: The annotation explorer view 2. The original unwrapped image of the cylinder seal P474995 as well as the location of manual annotations are revealed when the mouse cursor is pointed to any instance location represented by a blue circle.




Figure 5: The annotation explorer view 3. The annotation explorer can show the spatial location of more than one object category while using colour as a way to distinguish between them. For example, the instances of "animal" (blue) and "deity" (green) object categories are shown in a common two dimensional space.

throne, ball-staff, pot, and footrest, which demonstrate strong preferences for particular quadrants within the seal's layout. These alignments appear consistently across large datasets and correspond with the results shown in Table 1, indicating stable conventions in motif placement. Equally significant is the ability of the tool to highlight anomalies or secondary clusters. Such outliers can be indicative of several factors, including regional variation, workshop practice, recarving, or differences in the chosen start-point of the unwrapping. Visual clustering of annotations therefore provides a foundation for identifying seals whose layout diverges from wider patterns and may warrant further interpretative or technical examination. The comparative mode, which allows multiple motifs to be viewed simultaneously, also supports new forms of iconographic analysis. It becomes possible to test relationships often assumed in the literature, such as the pairing of certain ritual objects or the alignment of divine and human figures, by examining their spatial correlation across a large corpus. In this way, the visualisations offer a reproducible method for evaluating both expected patterns and unexpected compositional arrangements, enriching traditional art-historical approaches with quantitative insight. Overall, the tool provides a transparent, corpus-wide perspective on motif placement and co-occurrence that complements traditional stylistic and contextual analysis, while enabling new research questions to be formulated around consistency, variation, and spatial organisation in cylinder seal imagery.

§4. Automatic Motif Detection

§4.1 The manual annotation of cylinder seal unwrappings described in the previous section is practical for a limited number of images (e.g. a few hundred) but becomes labour-intensive and error-prone when applied to thousands of images of cylinder seals held by museums worldwide. To address this challenge, we employ computer vision and artificial intelligence (AI) techniques to automate parts of the annotation and curation process, assisting expert curators in managing large digital collections of cylinder seal impression and unwrapping.

§4.2 Recent advances in deep learning have enabled the development of a family of deep neural networks that can automatically recognise common objects like apple, cat, bus, aeroplane, tree in a large collection of images, accurately and at speed. These deep networks are called object de-

tectors. Since these object detectors are already very good at detecting common object categories, they can easily learn to detect novel object categories by training on a small number of exemplar images. Object detectors have been successfully used for automatically detecting a diverse set of objects categories. For example, Dutta, Bergel, and Zisserman 2021 have used object detectors to automatically detect printed illustrations in books. Dutta, Pérez-Campanero, et al. 2025 used an object detector for automatically detecting fish and Bain et al. 2021 used it for detecting a chimpanzee in videos captured as part of behavioural experiments. In this work, we adapt object detectors to automatically identify visual depictions of king, priest, deity, worshipper, and dancer motifs in cylinder seals imagery, as illustrated in Figure 6. Collectively, we refer to these categories as "person" motifs, encompassing all human representations depicted on cylinder seals.

§4.3 Tan, Pang, and Le 2020 introduced the EfficientDet family of neural networks for object detection that attains very high accuracy while being light in weight (both model size and computational cost). Pre-trained EfficientDet models are available under a permissive licence that allows unrestricted use in academic research projects and commercial products. These pre-trained models are trained to detect 80 common and diverse objects categories. This model has been trained on natural images, which are very different to unwrapped cylinder seals. Therefore, we need to adapt the model to this new domain by fine-tuning on a small number of images from unwrapped cylinder seals.

§4.4 Our manually annotated dataset contains 2491 instances of the "person" category. As a proof of concept, we fine-tuned the EfficientDet object detector to automatically identify "person" instances in cylinder seal impressions and unwrappings. The manually annotated dataset was split into three sets as shown in Table 2. The training set was used to fine-tune a pre-trained EfficientDet-d2 model. The fine-tuning process involves showing a cylinder seal impression to the object detector model and asking it to predict the location of "person" as a bounding box. The difference between predicted location and manually annotated location of instances of the "person" class is called prediction error which is used to automatically adjust the model such that the prediction error is lower next time the same image is encountered by the model. This process is repeated



Figure 6: Automatic detection of person (which includes king, priest, deity, worshipper, dancer) in the unwrapping of an image of [P475175](#) a cylinder seal in the British Museum.

over all the images in the training set. To evaluate the learning progress during the fine-tuning process, the object detector model needs to be tested on images that are not included in the training set, in order to provide a more fair estimate of the progress, and assess generalisation away from the training images. The images contained in the validation set are used to guide the training procedure and obtain a fine-tuned model with the best possible automatic detection accuracy.

§4.5 The automatic detection performance of the final fine-tuned “person” detector model is evaluated using the images contained in the test set. None of the images in the test set are contained either in the training or the validation set. An automatic detection is considered accurate if the detected region has more than 50% overlap with the region manually annotated by a human expert. Our fine tuned “person” detector model achieved a detection performance of 94%. In other words, the instances of “person” predicted by our model has an overlap of more than 50% with those manually annotated by a human expert in 94% of the test samples. Some examples of automatic detection are shown in Figure 7.

§4.6 The “person” detector was also tested on sample images from the Ashmolean Museum of the University of Oxford, the De Liagre Böhl collection of the University of Leiden, and the Bibliothèque Nationale d’France. These images differ in style, preservation, and imaging conditions from those of the British Museum and therefore present a more challenging test for the detector. The results of automatic detections on these samples are shown in Figure 8. These results show that the trained model can successfully identify

the “person” motif across collections. However, a larger and more systematic evaluation is needed to fully assess the model’s performance on external datasets.


§5. Software Tools and Online Demos

§5.1 All the software tools, manual annotations and person detection model files have been released publicly and are available to download from <https://www.robots.ox.ac.uk/~vgg/research/cylinder-seal/> which also includes online demos for Annotation Explore and the “person” motif detector. The software tools used for this research are available as open source project in the following [software repository](#). We have released the trained model, manual annotations and software tools under [Apache 2.0](#) open source license to support future research on these seals images which are publicly available from [CDLI](#). The person detection model is derived from [EfficientDet-D2 model weights](#) that are also available under Apache 2.0 license.

§6. Future Use Cases

§6.1 Although the dataset analysed in this study derives primarily from Old Babylonian material, the annotation and spatial normalisation framework described here is not limited to this corpus. The methodology is particularly suited to glyptic traditions characterised by recurrent compositional structures and stable motif hierarchies.


§6.2 A natural extension of the present work would be the presentation scenes of the Ur III period (c. 2112–2004 BCE), which directly precede and inform Old Babylonian glyptic. Ur III scenes exhibit a highly standardised compositional schema, typically depicting a worship-



(a) P473147



(b) P475251



(c) P474987



(d) P475112


Figure 7: Automatic detections of the "person" class in unwrappings of cylinder seals are shown as green rectangular boxes. The duplicate regions contained in the unwrappings are caused by the image acquisition process which ensures full capture of the cylindrical seal by rotating it twice during the capture process. (a) is in the Ashmolean Museum, Oxford, UK; (b) ; (c); and (d) are all in the British Museum, London, UK.

Subset name	Number of images
Training	276
Validation	48
Test	140


Table 2: Dataset subsets and number of images



(a) P473152



(b) P530730



(c) P530764

Figure 8: Results obtained by applying our "person" detector on cylinder seal impressions taken from collections held by other museums in the UK and the Netherlands. (a) is in the Ashmolean Museum, Oxford, UK; (b) and; (c) are in the De Liagre Böhl Collection, Netherlands Institute for the Near East, Leiden, The Netherlands.

per led by an interceding goddess before a seated deity or deified king (Elhewaily 2017, 136–139). Because Old Babylonian presentation scenes are derived from these prototypes (Elhewaily 2017, 144) applying the same spatial annotation framework to Ur III material would allow direct diachronic comparison. Such analysis could quantify the continuity and modification of compositional conventions across the transition from Ur III to Old Babylonian glyptic.

§6.3 A second use case concerns Old Akkadian presentation and mythological scenes. During the Old Akkadian period (ca. 2340–2200 BC), deities are clearly distinguished by horned headdresses and attributes and incorporated into structured compositions (Elhewaily 2017, 134–135). Although less rigid than Ur III examples, these scenes display recurrent spatial relationships between principal and secondary figures. The annotation framework described in §2 would enable the aggregation of these relationships at scale, supporting comparative analysis between narrative and ceremonial scene types.

§6.4 The methodology is also applicable to Early Dynastic banquet scenes and their subsequent

transformation into presentation scenes. Diachronic spatial aggregation would allow measurable comparison between horizontally organised banquet arrangements and the more vertically structured authority scenes that become dominant in later periods (Elhewaily 2017, 145–146). In this way, compositional change may be examined quantitatively rather than described in a narrative way.

§6.5 In addition to figural motifs, the framework can be extended to geometric and emblematic elements. Celestial symbols, framed inscriptions, drill-hole clusters, and repetitive border devices frequently occupy predictable positions within the visual field. Because such motifs often exhibit clearer geometric regularity than anthropomorphic figures, they may be particularly suitable for automated detection models. Mapping the density and placement of these elements across corpora could contribute to the identification of workshop practices, regional variation, and chronological trends.

§6.6 The availability of software tools and data described in this research under open source license is critical for such extensions and reuse.

Author Contributions (CRediT)

- Lara Bampfield [CRediT roles: Conceptualisation, Data curation, Formal analysis, Investigation, Methodology, Visualisation, Writing - original draft, Writing - review and editing]
- Jacob L. Dahl [CRediT roles: Conceptualisation, Funding acquisition, Methodology, Project administration, Supervision, Writing - review and editing]
- Abhishek Dutta [CRediT roles: Conceptualisation, Data curation, Formal analysis, Investigation, Methodology, Resources, Software, Validation, Visualisation, Writing - original draft, Writing - review and editing]
- Andrew Zisserman [CRediT roles: Funding acquisition, Methodology, Project administration, Resources, Software, Supervision, Validation, Writing - review and editing]

BIBLIOGRAPHY

- Bain, Max et al. (2021). "Automated audiovisual behavior recognition in wild primates". *Science Advances* 7.46, eabi4883. DOI: [10.1126/sciadv.abi4883](https://doi.org/10.1126/sciadv.abi4883). eprint: <https://www.science.org/doi/pdf/10.1126/sciadv.abi4883>. URL: <https://www.science.org/doi/abs/10.1126/sciadv.abi4883>.
- Dahl, Jacob L., Bertrand Lafont, and Nordine Ouraghi (2019). "Nouvelles recherches sur la collection des sceaux-cylindres orientaux de la Bibliothèque nationale de France". *Syria: Archéologie, Art Et Histoire*, 309–334.
- Dutta, Abhishek, Giles Bergel, and Andrew Zisserman (2021). "Visual Analysis of Chapbooks Printed in Scotland". In: *Proceedings of the 6th International Workshop on Historical Document Imaging and Processing*. HIP '21. Lausanne, Switzerland: Association for Computing Machinery, 67–72. ISBN: 9781450386906. DOI: [10.1145/3476887.3476893](https://doi.org/10.1145/3476887.3476893). URL: <https://doi.org/10.1145/3476887.3476893>.
- Dutta, Abhishek, Natalia Pérez-Campanero, et al. (2025). "Detect+Track: robust and flexible software tools for improved tracking and behavioural analysis of fish". *Royal Society Open Science* 12.7, 242086. DOI: [10.1098/rsos.242086](https://doi.org/10.1098/rsos.242086). eprint: <https://royalsocietypublishing.org/doi/pdf/10.1098/rsos.242086>. URL: <https://royalsocietypublishing.org/doi/abs/10.1098/rsos.242086>.
- Dutta, Abhishek and Andrew Zisserman (2019). "The VIA Annotation Software for Images, Audio and Video". In: *Proceedings of the 27th ACM International Conference on Multimedia*. MM '19. Nice, France: Association for Computing Machinery, 2276–2279. ISBN: 9781450368896. DOI: [10.1145/3343031.3350535](https://doi.org/10.1145/3343031.3350535). URL: <https://doi.org/10.1145/3343031.3350535>.
- Elhewaily, S. (2017). "The Intercession Scenes in Ancient Mesopotamian Cylinder Seals Till the End of the Old Babylonian Period". *Egyptian Journal of Archaeological and Restoration Studies (EJARS)* 7.2, 133–147. DOI: [10.21608/EJARS.2017.6838](https://doi.org/10.21608/EJARS.2017.6838).
- Englund, Robert K. (2014). "Seals and Sealing in CDLI files". *Cuneiform Digital Library Notes* 2014.4. [Online; accessed 2025-12-01]. ISSN: 1546-6566. URL: <https://cdli.earth/articles/cdln/2014-4>.
- Osley, Julian, Jane Savidge, and Gerry White (1991). "Art & Architecture Thesaurus. New York; Oxford: Oxford University Press for the Getty Art History Information Program, 1990." *Art Libraries Journal* 16.2, 29–33.
- Tan, Mingxing, Ruoming Pang, and Quoc V Le (2020). "Efficientdet: Scalable and efficient object detection". In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10781–10790.