Assessing the Expressivity of Iconclass to Embody Emotional Features in Classical Iconography

Alessandro Adamou[1], Davide Picca[2]

[1]Bibliotheca Hertziana - Max Planck Institute for Art History, Rome, Italy
[2]University of Lausanne, Lausanne, Switzerland

Abstract
As digital iconography takes hold in contemporary art history studies and gradually encompasses iconology and disciplines in hermeneutics, the question arises whether existing nomenclatures are suitable for representing different levels of interpretation of the image. One such nomenclature is IconClass, a hierarchical and alphanumeric classification system: we try to understand to what extent it can encode emotional variables that may aid iconological interpretation. In this study, we select a corpus of images from the photographic archive of the Bibliotheca Hertziana, categorized by IconClass terms from Classical mythology. We associated terms from the SenticNet vocabulary to each image, using the keywords that IconClass itself assigns to its categories. We then had the same images annotated by independent, non-expert raters, who were given no information on the scenes being depicted. A comparative study of these two ratings yielded that, although there is less disagreement on the sentiments most commonly associated to tragedy, the variety of emotional content in classical imagery remains largely unexpressed by what is arguably the most widely adopted classification system for visual art.

Keywords
art history, iconology, knowledge graphs, emotion mining, sentic computing, IconClass, Greek classics

1. Introduction

Iconography, as a pivotal branch of art history, delves into aspects of the subject matter of the artworks. It transcends the boundaries of mere style and structure, focusing instead on the rich tapestry of content and subjects that artworks embody [1, 2]. The themes and motifs detected by iconographical studies can, in turn, help us better understand their deeper meanings—what in art-historical parlance is sometimes referred to as *iconology*. The notion of iconology is debated in art history, with the very dychotomy with iconography being called into question, though the former is generally used to refer to the deeper interpretation of art [3].

The IconClass classification system[1] stands out as one of the key tools for digital iconography at the disposal of scholars and practitioners in the GLAM sector (Galleries, Libraries, Archives, and Museums) [4]. IconClass is a widely adopted taxonomy of visual art themes, figures, and concepts, which are mapped to an alphanumeric signature system of over 28,000 unique codes.

SWODCH'24: International Workshop on Semantic Web and Ontology Design for Cultural Heritage, October 30–31, 2024, Tours, France
alessandro.adamou@biblhertz.it (A. Adamou); davide.picca@unil.ch (D. Picca)
0000-0002-9272-908X (A. Adamou); 0000-0003-2014-0855 (D. Picca)

[1]IconClass, http://iconclass.org/

These cover a broad spectrum of themes and subjects across various cultures and historical periods. Let us consider, for example, the Greek mythology associated to the Trojan war, which is the focus of this study. In that case, the IconClass categories of interest begin with 94C to 94K, or 95. While the former are hierarchically organized along a phenomenological dimension, with subclasses covering specific books in Homer's poems and events therein, the latter is centered around heroes and heroines, with subclasses covering phases of their lives, or activities and relationships that shape their characters. For instance, `94F82` is the category that represents the death of Patroclus at the hands of Hector, whereas `95B(CIRCE)2` relates to the love affairs (in general) of Circe, the enchantress from the Odyssey. Different criteria govern other parts of the taxonomy: for example, `46C4` is the category of interstellar travel.

As IconClass has enjoyed wide adoption for over four decades, the question arises whether it is able, by nature or through its usage, to inform not only iconographical studies, but also iconological ones on possible interpretations of what is depicted. Particularly, with this paper we intend to explore to what extent IconClass can encode nuanced features related to emotions expressed through the iconography itself. We do so through the case study of the Homeric myth as it is expressed through visual arts ranging from classical to beyond the Italian Renaissance.

Our work borrows from the theories of *sentic computing* formulated by Cambria and Hussain [5], which attempt to bridge computational linguistics, semiotics, and affective computing. We use a controlled vocabulary of emotional coordinates from the SenticNet dataset and rate a corpus of 300 images from the photographic library of the Bibliotheca Hertziana according to these coordinates. We compare a purely linguistic rating, extracted from the IconClass categories and their associated keywords, to an independent rating by human annotators, and seek commonalities and discrepancies arising from their (dis-)agreement.

While a contribution of this work lies in the proposed methodological approach, the critique that arises from its application does not yet seek to enlarge the landscape of data schemas for iconography. The ultimate outcome of our work, which will follow up on this paper, may very well introduce a new vocabulary, establish a novel ontology, or integrate emotional dimensions with the IconClass framework. What this paper aims for is to provide grounding for a case of complementing IconClass with emotional descriptors, to provide a more holistic approach to understanding artworks by accommodating both the intellectual and affective dimensions, thereby enhancing the analytical capabilities of art history and iconography, and offering a complementary lens through which artworks can be classified and understood [6, 7, 8, 9].

Using a theme that has been repeatedly represented through history, this work also offers insights into how similar emotions are depicted across cultures and epochs. It also seeks to explore the feasibility and benefits of an integration with IconClass, hypothesizing that a combined approach could significantly enrich the cataloging, analysis, and understanding of artworks. Through a balanced examination of IconClass alongside an emerging potential emotional classification system, this study aims to contribute to broader discussions about innovative and inclusive practices in the classification of art.

After an overview on related work in digital iconography, iconology and sentic computing, we delve into the materials and methods in Section 3. Section 4 then describes the comparative analysis and formulates conjectures over the results. Finally, Section 5 offers a concluding discussion touching upon potential strategies for acting upon these outcomes and our outlook for the future.

2. Related Work

Recent advancements in the detection and analysis of emotions within iconographic artworks illustrate a promising intersection of art history and cognitive science and technology. These interdisciplinary efforts aim to enhance our understanding of how emotions are conveyed and perceived in artistic representations, leading to more nuanced interpretations and accessible art historical knowledge.

One groundbreaking approach involves the adaptation of machine learning techniques, specifically convolutional neural networks (CNNs), to interpret emotional content in artworks. Research by González-Martín et al. [10] indicates that CNNs, traditionally trained on non-artistic images, can be effectively adapted to the artistic domain by addressing the cross-depiction problem. This involves algorithms such as QuickShift, which enhance the network's ability to generalize across different artistic styles, thereby improving accuracy in emotion detection. This methodology not only bridges the gap between digital image processing and art analysis, but also opens up possibilities for more robust cataloging and understanding of emotional expressions across diverse art forms.

Further enriching the domain are multimodal frameworks that integrate various types of data (e.g., visual, textual, and auditory) to analyze emotional content. Originally developed for conversational dynamics in videos, such frameworks can be adapted for art, considering how narratives within artworks contribute to emotional impact. This suggests a layered approach to emotion detection, where the interplay of different modalities can provide a deeper understanding of how artworks engage viewers emotionally. Hazarika et al. [11] present such a multimodal emotion detection framework, which could potentially be applied to the analysis of emotional content in iconic artworks .

Moreover, studies examining cognitive responses to iconic versus realistic depictions reveal that iconic representations, such as those in cartoons or stylized graphics, communicate emotional information more effectively. Kendall et al. [12] highlight the differences in neural processing that occur when viewers encounter iconic versus realistic images, suggesting that iconicity enhances emotional communication through visual art, potentially due to the simplified and exaggerated features that better capture and convey emotional states.

The influence of demographic factors on emotional responses is also critical. Research by Ko and Yu [13] investigates gender differences in responses to iconic designs, using facial expression recognition software to analyze how different genders perceive and react emotionally to the same visual stimuli. Such studies highlight the need for considering a variety of viewer backgrounds when analyzing emotional responses to art, providing insights into how personal experiences and cultural contexts might influence emotional interpretation .

Additionally, the universality of emotional responses to sensory inputs extends beyond visual arts to other domains such as poetry, where Auracher et al. [14] explore how sound iconicity in poetry can evoke specific emotions. This research suggests parallels in visual art, where certain visual forms or styles may universally trigger emotional responses, enhancing our understanding of the cross-sensory dimensions of emotional perception in art.

These diverse approaches not only underscore the complexity of emotion detection in art but also highlight the potential for developing more sophisticated tools and methodologies, which can cater to the multifaceted nature of art perception and appreciation.

3. Dataset and Methodology

The choice of classical iconography as a case study has a number of peculiarities. For one thing, as stated earlier, it draws on a multitude of diverse representations of the Greek myth throughout centuries and cultures, paving the way to iconological considerations on how the scene should be interpreted in the context in which it was depicted. Secondly, its representation in IconClass is quite rigid, opting for an almost exquisitely event-centric lens, which means that the IconClass categories themselves are expected to scarcely touch upon emotional cues. While there are categories that embody some of them explicitly, such as `94F83` (Achilles' grief over Patroclus), others like `94C11` (the Judgment of Paris), require them to be drawn from context.

3.1. The Bibliotheca Hertziana Photographic Library: A Pivotal Resource in Art Historical Research

The Photographic Library of the Bibliotheca Hertziana, or Fotothek,[2] is an archive that holds a significant position within the domain of art historical research. Established in Rome, the archive was designed to support scholarly endeavors in the study of art history, particularly those concerning Italian and Mediterranean art. This repository offers an extensive collection of photographic reproductions that document a wide array of artworks, from ancient sculptures and frescoes to Renaissance paintings and architectural monuments. The Fotothek was initiated as a core component of the Hertziana's research infrastructure, reflecting the vision of its founder, Henriette Hertz. The aim was to create a comprehensive visual resource that would aid scholars in accessing and studying Italian art, regardless of their geographical location.

Over the decades, the collection has grown to over 1,500,000 images depicting monuments and artworks, which serve as critical tools for researchers. The strength of the Fotothek lies in its meticulous organization and the breadth of its collections: it includes not only photographs of well-known masterpieces but also of lesser-known works, providing a broader view of the artistic landscape. In recent years, the Fotothek has embraced digital technology to increase accessibility to its collections. A significant portion of the archive has been digitized, allowing online access via the IIIF protocols[3] to scholars worldwide. Each digitized photograph is accompanied by metadata following the MIDAS schema [15], including information about the artwork's creator, date, and style. The artwork subject matter is mainly aligned with IconClass classification codes. It should be noted that the assignment of IconClass codes is often carried out by scholars in art history as part of dedicated digitization campaigns, under the supervision of resident staff.

3.1.1. Materials

Knowing that we could avail ourselves of a cohort of twenty annotators, an iconographic corpus was created with the intent of striking a balance between workload assigned to annotators and dataset size and variety.

Items in the Fotothek are grouped by work, or "object", meaning that paintings belonging to the same pictorial cycle, sides to a coin, or details of the same historical building, all be-

[2] Fotothek, https://foto.biblhertz.it/.
[3] Internet Image Interoperability Framework, https://iiif.io/.

long together as parts of that object. Each object or part has its IIIF manifest, which groups together related images, such as photos of the same sculpture from different angles, or color and black/white photos of one painting. Metadata tend to propagate top-down from the object. Both the general object and its specific parts can accommodate IconClass category codes.

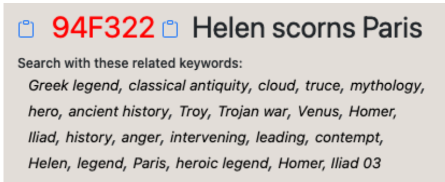

Figure 1: Left: Giuseppe Cammarano, "Paris and Helen, with Hymen", drawing, 1782 or 1792. Corpus Gernsheim, licensed CC BY-NC 4.0 (https://foto.biblhertz.it/document/obj/08123696). Right: detail of IconClass 94F322, with the associated set of keywords.

Through a low-level query[4] we retrieved the identifiers of all the objects that were tagged with at least one IconClass category belonging in the 94C to 94K range, or one in the 95 group, which is about individual characters. From the latter, we excluded characters not associated to Trojan War mythology, including but not limited to[5] the "Iliad" and "Odyssey". Because the image corpus was to be submitted to third parties for annotation, it had to comprise only images that were licensed free of charge for public use ('*freigegeben*') by the rights owner. A harvester was written for fetching the actual images and ran from outside the network of the Max Planck Society, so that the retrieval of non-licensed images would fail. We filtered the resulting dataset to limit the number of different photos of the same work. Introducing a few negative examples, i.e. images with other IconClasses in the 94-95 range, yielded a corpus of 300 images.

IconClass associates to each category code a bag of keywords available in multiple languages, which represent topics relevant to that category. For instance, Figure 1 shows a 18th-century Italian drawing of a scene from the troubled encounter between Helen and Paris. In the Fotothek, the image is annotated with IconClass 94F322 ("Helen scorns Paris"), which would lead to believe that the scene illustrates Helen's resistance to Venus' induction to fall in love with Paris. That category is decorated with keywords like 'Trojan war', 'Helen', 'truce', but also 'anger' and 'contempt'. Concerning the last two, we observe that, while 'contempt' is directly assigned to this category, 'anger' is inherited from 94, which is general for the tenth year of the Trojan war. We expect this keyword inheritance mechanism to deeply affect the outcome of our study.

[4]The Fotothek metadata sit upon the XML database eXist-db (http://exist-db.org/), hence XQuery was used.
[5]A typical outlier prominently figured in art history is the Laocoön, who is found in Virgil but not in Homer.

3.2. Experiment on Emotion Tagging in Classical Iconography

The primary goal of this experiment is to capture a range of emotional interpretations that viewers associate with specific scenes from classical mythology. It was designed to explore how different viewers perceive and respond to classical themes and to the artist's nuanced interpretation of them. It also aims to assess the potential for integrating emotional responses into traditional iconographic cataloging systems, in the event that the latter are found to be lacking on that front. In the longer run, we intend to devise a way to add a layer of emotional data to the existing IconClass codes used in the archive, potentially enriching the academic and educational value of the Fotothek's collections.

Having built the iconographic corpus described in Section 3.1.1, we proceed to generate two comparable "ratings" of it. The categories and keywords of IconClass itself represent the source of the first rating: we intend to obtain a dataset that represents how IconClass assesses the emotions associated to a scene depicted in art, based only on what event we know is being portrayed, and on what emotions are associated by it by art historians and classicists, who presumably contributed the keyword cloud. On the other hand, we gather the independent responses of non-experts whose bias is minimized, and who represent likely consumers of visual art, based on what emotions they perceive as elicited from an observation of the artwork itself. These two ratings are then encoded as matrices whereupon it is possible to conduct various types of analytics, most importantly, detecting patterns of agreement between them.

3.2.1. Extracting emotional terminology from IconClass

In the RDF description of an IconClass code, categories are related to one another using SKOS, while the associated keywords that help convey aspects of the category semantics are represented in the Dublin Core schema as `dc:subject` predicates over plain literals, not aligned with any vocabulary, dictionary, or authority file. Another example is given in Figure 2.[6]

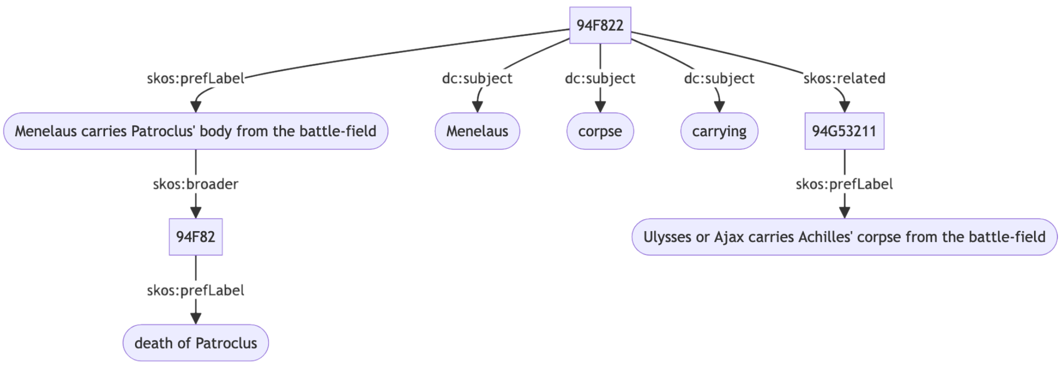

Figure 2: RDF representation of category `94F822` (https://iconclass.org/94F822.rdf).

It is also implied from the human-readable description of the IconClass code, that it also inherits the keywords of the one(s) that is subsumes: therefore, emotion-laden terms appearing high up in the notation hierarchy are likely to sway the sentiment associated to specific episodes.

[6]See e.g. https://iconclass.org/94F822.rdf for https://iconclass.org/94F822

As a controlled vocabulary of emotions, and by extension the rating dimensions of this experiment, we adopt the SenticNet vocabulary. SenticNet was built from 100,000 concepts automatically extracted using a blend of symbolic and sub-symbolic AI techniques. Each concept includes a multiword expression, weights for four affective dimensions (pleasantness, attention, sensitivity, aptitude), primary and secondary mood labels, a polarity score, and semantically related concepts. Particularly, it features a vocabulary of 24 terms, such as "anger", "delight" or "responsiveness", that denote emotions [16]. For instance, the aforementioned "truce" is associated to the SenticNet primary emotion "contentment" and secondary emotion "serenity".

The strategy for extracting the IconClass emotion rating for our image corpus was to build a knowledge graph. First, we extract subject-related metadata from the IIIF manifests of the images in the Fotothek and represent them as RDF. Then, we crawl iconclass.org for the RDF data of the matching IconClass codes, traversing the SKOS hierarchy and storing the results.

SenticNet extraction for each image was performed through querying this knowledge graph. Through a SPARQL query, we traverse the hierarchies of all the categories that an image is tagged with at the Fotothek, collecting all its `dc:subject` keywords in the process. We then look up each keyword on the SenticNet dataset and, for every match, take both its primary and secondary emotion. Therefore, for every image i we end up with a vector $F_i = \{n_{ij}\}$, where n_{ij} is an integer that counts the occurrences of emotion j on the IconClass categories with which i is annotated in the Fotothek. The matrix $\mathcal{F} = \{F_i\}$ of size $n \times m$, where n is the number of images and m are the 24 SenticNet emotions, is thus the dataset of the emotional ratings of the image corpus according to IconClass itself.

3.2.2. User study: image annotation with SenticNet

To obtain a counter-rating of our iconographic corpus by its viewers, our starting hypothesis is that even who is not an expert in the classical subject matter is able to detect nuances in the emotions evoked from the depicted scene. Whether these nuances are the result of the artist's own interpretation of the scene, or of shortcomings in the IconClass scheme, is an iconological research question in and of itself, which the results should help us explore.

A second rating, to be compared against the one emerging from IconClass, was obtained through a user study.[7] Each participant was asked to annotate a subset of the image corpus with emotional tags, so that exactly four users would have annotated the same image. The annotation process consisted of selecting one or more regions of each image, with each region being tagged with exactly one emotion from SenticNet.[8] One region could only be tagged with multiple emotions by replicating the region itself. This task was performed independently, and the participants were encouraged to consider both the emotional tone of the scene and the emotional responses they believed the artwork was intended to evoke in an audience. They were not informed on what artwork they were looking at, what scene it depicted, or what IconClass terms the Hertziana Fotothek had it annotated with. They were also given complete freedom to choose the shape and size of the regions—whether it highlighted a face, an entire

[7] The user base consisted of a cohort of 20 undergraduate students enrolled in the Computer Science for the Humanities course at the University of Lausanne.

[8] For the sake of future studies, an additional option for "emotionlessness" was given to the users in case they still wanted to highlight a region of interest without assigning it a SenticNet term.

body, another body part or, possibly, even an inanimate object. These measures were all taken in the interest of minimizing annotator bias, as well as allowing them to express the emotions conveyed in one depicted scene at a fine granularity, should they choose to do so.

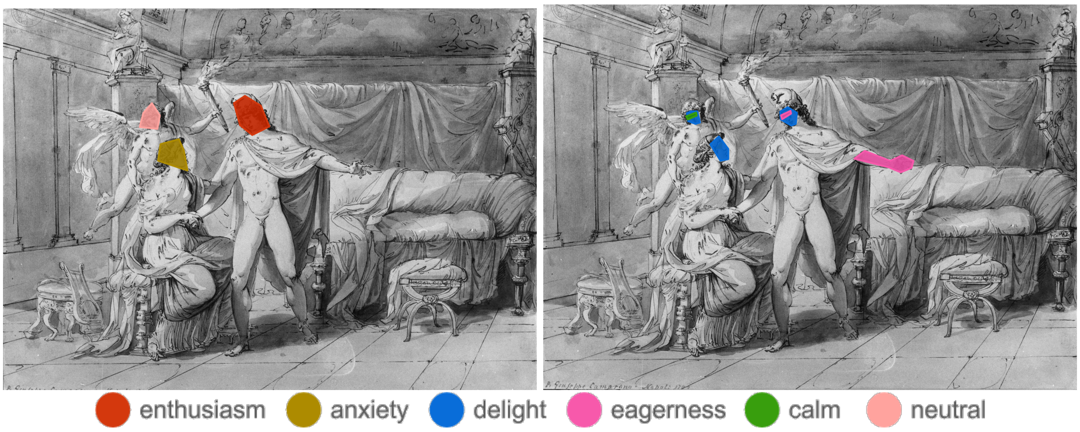

Figure 3: Two different annotations of the drawing "Paris and Helen, with Hymen".

Figure 3 shows two different ways in which the drawing from Figure 1 was annotated. The participants interpreted Helen's facial expression very differently, noting 'anxiety' and 'delight' respectively—none of which matches the repulsion expected from the IconClass category name, "Helen scorns Paris". They also used different but agreeable tags for the expressions of Paris and Hymen, yet one of them reinforced Paris' 'eagerness' by also annotating his tense arm with it.

The platform employed for recording annotations is *LabelStudio*[9], a data labelling tool that is available both on the cloud and as an open source local package. It was chosen due to its flexibility in allowing multi-user rating, its automatic calculation of simple rater metrics, and its ability to also annotate text, which will prove useful in future developments of this study.

The result of the annotation process is the matrix $\mathcal{A}$ which, like $\mathcal{F}$, is of size $n \times m$, so that it could be used for the subsequent analysis of the emotional responses. These were further compared to the traditional thematic classifications provided by IconClass codes.

4. Results and Discussion

The great freedom that was granted to human annotators in the selection of regions, aimed at facilitating the variety of perceptions, affects the number of times an annotated value repeats itself. Therefore, great agreement between the four annotators per image is not to be expected. Because our goal is to assess the agreement between artwork viewers and IconClass itself, we consider matrix $\mathcal{F}$ as representative of one rater, and matrix $\mathcal{A}$ as another, single rater. We therefore proceed to calculate various forms of agreement between these two.

As a reliability measure, we adopt Krippendorff's Alpha. This is a very flexible mechanism, as it is able to accommodate categorical rating scales—as is the case of SenticNet's emotions—

[9]LabelStudio, https://labelstud.io/.

Table 1
α-agreement on emotions, rated considering zeroes both as relevant values and as missing values.

Emotion	Alpha (zeroes as 0)	Alpha (zeroes as missing)
acceptance	-0.02	-0.13
anger	-0.12	**1**
annoyance	-0.13	—
anxiety	-0.15	-0.17
calm	-0.2	0
delight	-0.01	-0.22
dislike	-0.12	-0.24
eagerness	-0.15	0
enthusiasm	-0.02	—
grief	0.03	-0.08
joy	-0.08	—
melancholy	-0.11	—
pleasantness	-0.04	0
responsiveness	-0.17	—
sadness	-0.13	0
serenity	-0.1	-0.07
terror	-0.02	0.25

and missing data. Unlike e.g. Fleiss' Kappa, however, the Alpha has the desirable feature of supporting cases where the total ratings do not amount for the same number for each item [17], which is again our case due to allowing users to annotate multiple regions with the same emotion and allowing one emotion to emerge from multiple IconClass keywords.

Krippendorff's Alpha, or α-agreement, is normalized between -1 and 1. A value greater than zero denotes inter-rater agreement; one close to zero means that the rating is as unreliable as in the case of randomness; one lower than zero denotes disagreement.

The first step was to consider the SenticNet emotions as the items being rated, thus calculating the Alpha between the *inverted* $\mathcal{F}$ and $\mathcal{A}$, to detect whether certain categories of emotion gather more agreement than others. Two calculations were made: one that considers a value of zero as an effective zero, and one that considers it a missing rating. This was done because considering zero-values causes the agreements to approach randomness, but can help make subtleties emerge, which would not stand out by considering them as missing values. The results can be seen in Table 1, for only those emotions that displayed a degree of polarity in agreement.

What is striking yet to be expected out of knowledge of the domain, is that the annotators agreed with IconClass on emotions typically associated to the tragedy of the Greek *epos*, as indicated by the perfect agreement on *anger* and the 25% agreement on *terror*. Sentiments that show disagreement in the 13-24 percent range appear either to be more subtle on the negativity front (e.g. *anxiety*, *dislike*), or to embody positiveness (e.g. *delight* or *acceptance*). This disagreement is largely in favor of the human annotators, who used these terms more often due to perceiving nuances of positive feelings in the depicted scenes, which are not considered characteristic of the corresponding episodes or characters in the Trojan war.

To confirm or disprove this assumption, we can look at the IconClass codes themselves. Recall

Table 2
The most agreed or disagreed IconClass codes, rated considering zeroes as relevant values.

IconClass	Alpha	Name
94C1131	-0.32	the Judgment of Paris (without Mercury)
94D132	-0.29	the sacrifice of Iphigenia
94F31	-0.27	Paris and Menelaus duelling
94G533	0.17	Thetis mourning Achilles
94H151	-0.27	the Trojan horse
94H243	-0.27	Polyxena is sacrificed by Neoptolemus on Achilles' tomb
94H2452	0.20	Hecuba finds her dead son Polydorus on the sea-shore
94I134	0.21	Polyphemus is blinded with a pointed stake
94I6	-0.28	the end of the suitors
95A(ACHILLES)	-0.28	(story of) Achilles
95A(DIOMEDES)4	0.35	Diomedes - aggressive, unfriendly activities and relationships
95A(LAOMEDON)31	-0.29	Neptune builds the walls of Troy as Apollo tends Laomedon's flock
95A(ORESTES)312	0.18	Orestes kills his mother Clytaemnestra, and Aegisthus
95A(PARIS)	-0.36	(story of) Paris (Alexander)
95A(ULYSSES)	-0.35	(story of) Ulysses
95B(CASSANDRA)6	0.24	(story of) Cassandra - suffering, misfortune
95B(CIRCE)	-0.33	Circe
95B(LAODAMIA)21	-0.27	Laodamia and Protesilaus
95B(LAODAMIA)6	-0.27	suffering, misfortune of Laodamia

that codes starting with 94 denote events, and those starting with 95 denote characters and their personal lives. Therefore, if agreement is primarily found over a universally tragic event, or a character primarily known for their dire fate, then we are closer to proving that IconClass has a gap in representing classical iconography in their affective variety and subtleties.

We aggregate the ratings in $\mathcal{F}$ and $\mathcal{A}$ to obtain two matrices of size 173×24, where 173 is the number of IconClass categories in the range being considered, with which our corpus is annotated at the Fotothek. We then calculate the α-agreement between these.

Table 2 shows the categories for which the highest agreement or disagreement was found. Indeed, the human annotators are shown to mostly agree with IconClass over episodes—such as the mourning of a loved one or the retrieval of their dead body—whose tragic nature is always represented beyond a doubt. By contrast, the highest disagreement in the event classes is found where multiple engagements are at play, such as in the Judgment of Paris, or the human sacrifices of Iphigenia and Polyxena. In the latter case, the terror and anxiety of the sacrificial victims are overshadowed, in the scene depictions, by the sense of satisfaction of the other attendees, in stark contrast with the focus on the victim that is established by IconClass.

A similar argument applies to character classes: the lack of agreement over complex characters such as Achilles or Paris, appearing through their most general codes, owes to the intricacy of their stories, whereas characters like Orestes and Cassandra, whose artistic focus tends to be on their most grievous vicissitudes, enjoy a more faithful emotional representation in IconClass.

The materials and tools employed for this analysis are available on GitHub[10].

[10] https://github.com/unil-ish/Hertziana_IconClass_Public

5. Conclusion

This experiment highlighted the potential for incorporating emotional tagging into the classification and analysis of artworks. The ideal way of doing so is being investigated, however, the likelihood of the highlighted limitations being partly a consequence of using a rigid taxonomy does not rule in favor of an 'enhanced IconClass'. While one branch of its taxonomy, starting with code 56, specifically covers emotions, it is largely underutilized in the Fotothek, possibly owing to the scene being central in classical art, whereas modern abstract paintings are more likely to be annotated with those categories. Still, to efficiently interoperate with IconClass, a more expressive ontological structure should be thought of with usability in mind.

Our next step is to integrate an actual art-historical perspective. From conversations with scholars in art history, it emerged that a frequent criticism on the use of IconClass is that, even in the rare event that multiple codes are employed for an image, they fail to convey the deeper meaning of a visual work. We have so far concentrated on a restricted set of codes, yet are aware that, if we were to factor in motifs from other co-occurring codes, this would enable richer if more complex forms of emotion mining. Collaborations with art historians are being sought in order to devise a strategy for integrating emotional data with IconClass annotations, as well as a further user study being planned, this time involving classicists and art historians.

In the direction of studies in digital hermeneutics, it comes natural, in this context, to relate iconographical and iconological studies to the interpretation of text, which connects to our earlier work on the Iliad [6, 9] and on the profiling of literary characters [18]. Because IconClass codes in Classics are organized by episode, we are confident that a *trait d'union* between linguistic and iconographical studies lies in detecting the events shared by the poems and artworks: this branch of our study is currently underway.

Representing and integrating emotional data in art history catalogs would enable iconological studies on their deeper interpretations. These studies can, in turn, lead to interactive archives that engage users not only intellectually but also emotionally, cross-referencing works not only by thematic or subject matter criteria but also by the affective responses they invoke.

References

[1] W. Godzic, Iconographic-iconological method in film research, Artibus et Historiae 2 (1981) 151–157. doi:10.2307/1483108.

[2] N. O. P. Vago, F. Milani, P. Fraternali, R. da Silva Torres, Comparing CAM algorithms for the identification of salient image features in iconography artwork analysis, Journal of Imaging 7 (2021). doi:10.3390/jimaging7070106.

[3] D. Džalto, Iconography and iconology, in: Oxford Research Encyclopedia of Religion, Oxford University Press, 2023. doi:10.1093/acrefore/9780199340378.013.914.

[4] J. Becker, H. V. D. Waal, L. D. Couprie, R. Fuchs, E. Tholén, Iconclass: an iconographic classification system, Art Libraries Journal 8 (1983) 32 – 49. doi:10.1017/S0307472200003436.

[5] E. Cambria, A. Hussain, Sentic computing, Cogn. Comput. 7 (2015) 183–185. doi:10.1007/S12559-015-9325-0.

[6] J. Pavlopoulos, A. Xenos, D. Picca, Sentiment analysis of Homeric text: The 1st Book of Iliad, in: Proceedings of the Thirteenth Language Resources and Evaluation Conference, 2022, pp. 7071–7077. URL: https://aclanthology.org/2022.lrec-1.765/.

[7] D. Picca, C. Richard, Unveiling emotional landscapes in Plautus and Terentius comedies: A computational approach for qualitative analysis, in: Proceedings of the Ancient Language Processing Workshop, 2023, pp. 88–95. URL: https://aclanthology.org/2023.alp-1.10/.

[8] D. Picca, Fluid Dynamics-Inspired Emotional Analysis in Shakespearean Tragedies: A Novel Computational Linguistics Methodology, in: Proceedings of the 2nd Workshop on Mathematical Natural Language Processing@ LREC-COLING 2024, 2024, pp. 11–18. URL: https://aclanthology.org/2024.mathnlp-1.2/.

[9] D. Picca, J. Pavlopoulos, The Iliad: a French dataset for emotion recognition in classical literature, in: Proceedings of the fourteenth language resources and evaluation conference (LREC-COLING 2024), 2024.

[10] C. González-Martín, M. Carrasco, T. G. W. Wielandt, Detection of emotions in artworks using a convolutional neural network trained on non-artistic images: a methodology to reduce the cross-depiction problem, Empirical Studies of the Arts 42 (2024) 38–64. doi:10.1177/02762374231163481.

[11] D. Hazarika, S. Poria, R. Mihalcea, E. Cambria, R. Zimmermann, ICON: Interactive Conversational Memory Network for Multimodal Emotion Detection, in: E. Riloff, D. Chiang, J. Hockenmaier, J. Tsujii (Eds.), Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Brussels, Belgium, 2018, pp. 2594–2604. doi:10.18653/v1/D18-1280.

[12] L. N. Kendall, Q. Raffaelli, A. Kingstone, R. M. Todd, Iconic faces are not real faces: enhanced emotion detection and altered neural processing as faces become more iconic, Cognitive Research: Principles and Implications 1 (2016) 19. doi:10.1186/s41235-016-0021-8.

[13] C.-H. Ko, C.-Y. Yu, Gender Differences in Emotional Responses to Iconic Design, in: 2016 5th IIAI International Congress on Advanced Applied Informatics (IIAI-AAI), 2016, pp. 785–790. URL: https://ieeexplore.ieee.org/document/7557718. doi:10.1109/IIAI-AAI.2016.35.

[14] J. Auracher, S. Albers, Y. Zhai, G. Gareeva, T. Stavniychuk, P Is for Happiness, N Is for Sadness: Universals in Sound Iconicity to Detect Emotions in Poetry, Discourse Processes 48 (2010) 1–25. doi:10.1080/01638531003674894.

[15] J. Bove, L. Heusinger, A. Kailus, Marburger Informations-, Dokumentations- und Administrations-System (MIDAS), Digitales Informationssystem für Kunst- und Sozialgeschichte (DISKUS), 2001. URL: http://archiv.ub.uni-heidelberg.de/artdok/3770/1/Bove_Heusinger_Kailus_MIDAS_Handbuch_2001.pdf.

[16] Y. Susanto, A. G. Livingstone, B. C. Ng, E. Cambria, The hourglass model revisited, IEEE Intelligent Systems 35 (2020) 96–102. doi:10.1109/MIS.2020.2992799.

[17] K. Krippendorff, Content Analysis: An Introduction to Its Methodology, SAGE Publications, Inc., Thousand Oaks, California, 2013, pp. 221–250. doi:10.4135/9781071878781.

[18] M. Egloff, D. Picca, A. Adamou, Extraction of character profiles from the Gutenberg Archive, in: E. Garoufallou, F. Fallucchi, E. W. De Luca (Eds.), Metadata and Semantic Research - 13th International Conference, MTSR 2019, volume 1057 of *Communications in Computer and Information Science*, Springer, 2019, pp. 367–372. doi:10.1007/978-3-030-36599-8_32.