1 Above: Jan van Eyck, *Saint Francis Receiving the Stigmata*, 1435–1440, 29.3 x 33.4 cm, Galleria Sabauda, Turin. Below: attributed to Jan van Eyck, *Saint Francis Receiving the Stigmata*, 1440–1445, 12.4 × 14.6 cm, Philadelphia Museum of Art, John G. Johnson Collection, Philadelphia

Till-Holger Borchert, Alessandra Conti, Batu Arda Düzgün

and Carina Popovici

Integrating Art Historical Knowledge into AI-Based Authorship Detection in Jan van Eyck's *Saint Francis Receiving the Stigmata*

1. Introduction

The two fifteenth-century paintings depicting *Saint Francis Receiving the Stigmata*, both attributed to the Flemish Primitive Jan van Eyck (1390–1441), have challenged Van Eyck scholars in terms of their attribution. The works are presently held in the Galleria Sabauda in Turin and the Philadelphia Museum of Art, respectively illustrated above and below in Fig. 1. Both belong to a category of works that have sometimes been accepted as authentic paintings by Van Eyck, while other scholars have rejected the master's authorship[1]. Some scholars have favoured one version over the other, others have regarded both panels as workshop copies after a lost original by Jan van Eyck (Jones 2014, 30–47).

In addition, the documentary evidence – two versions allegedly painted by Jan van Eyck that were inherited by the daughters of Anselm Adornes – is convenient at first sight, but inconclusive: while it is convenient and mentions two paintings of the subject by the master, the source is not contemporary and not detailed enough to draw conclusions (Geirnaert 2000); (Dumolyn and Geirnaert 2024); (Coombs, Harrison and Guidicini 2026). Furthermore, the reception of the compositions or details thereof suggests that additional versions – believed to be by Van Eyck in the fifteenth century – may have circulated in Spain and Southern Italy in the 1450s–1460s (Natale 2001); (Fernando Benito Doménech and José Gómez Frechina 2001); (Borchert 2002); (Challéat 2012); (Kemperdick 2012); (Jones 2014) ; (Fransen 2017); (Espin forthcoming). A related example is preserved in the Museo del Prado, Madrid, formerly in the collection of Queen Isabella Farnese, as indicated by the white lily mark visible at the lower right of the painting's surface.[2] As reported by Pérez Preciado, this oil on panel (47 × 36 cm), likewise titled *Saint Francis Receiving the Stigmata*, has been attributed to the Master of Hoogstraten—a conventional designation introduced by Max J. Friedländer in the early twentieth century (Pérez Preciado 2024). In Northern Italy, evidence points to the circulation of such versions slightly later, in the 1470s and 1480s (Rosenauer 1993); (Rohlmann 1994); (Nuttall 2004).

Technical studies have since highlighted differences in technique and scale – the smaller Philadelphia version is painted on parchment mounted on wood – and pointed out that there are pentimenti in the underdrawing of the larger Turin version – usually understood as an argument for the artist's own creative process and intervention (Van Asperen de Boer, et al. 1997). But since in a medieval painters' workshop the underdrawing may not have necessarily been applied by the same painter who painted the surface, strictly speaking the existence of pentimenti does not present a conclusive argument for Van Eyck's authorship either. In short: traditional connoisseurship – even if combined with additional art historical methods – has not been able to yield a convincing answer regarding the attribution and specifically concerning the question of whether either painting was painted by Jan van Eyck, and if so, which one.

Traditional connoisseurship has been and continues to be an intuitive occupation; only in the nineteenth century did a physician from Bergamo, Giovanni Morelli, attempt to develop an objective methodology by comparing details like earlobes and noses to argue for attributions. His entire concept, however, was fundamentally embedded in the reality of academic artistry whereby an individual artist was the sole creator of a painting, possibly assisted by workers who prepared colors, canvases and eventually carried the painting around on the mountains (Anderson 2019). Although it worked very well and became a successful means for attribution – despite initially receiving much criticism as well – it remains a one-dimensional model.

Precisely for this reason, both versions of the *Saint Francis receiving the Stigmata* seemed to be a very interesting test case to put the algorithm of Art Recognition to the test. We predicted that the results would be controversial, and we also anticipated that the results would provoke comments and rejection by some critics.

Nonetheless, the results provided by the algorithm turned out to be much more surprising: at first, it classified both paintings as not by Van Eyck, with an unusually high degree of probability. This result raises questions that are relevant for both traditional connoisseurship and the use of algorithms for attribution purposes. Are both models too one-dimensional to account for the complexity of late medieval workshop practices? If so, which additional layers of information must be accounted for? One such layer is later interventions: overpainting and restorations. In the case of the two van Eyck paintings, subtracting the known restorations changed the algorithmic attribution of the Turin panel from negative to positive, while the attribution of the Philadelphia panel remained unchanged.

When it comes to medieval workshop practice, things are potentially complex: there were undoubtedly workshops where the master executed his paintings himself; there were masters who outsourced their compositions to other workshops to have them produce paintings in high demand for them; there were masters whose workshop was more like a factory where several collaborators were involved in all kinds of production stages (much as in the case of Warhol or Hirst, or Disney).

We do not know much in the way of facts about Van Eyck's workshop other than that he had a workshop and worked with several assistants. We do not know for sure what these assistants did, but they might have been involved in various stages of the production process. Instead of the horizontal working model – several assistants painting different paintings in a workshop – there could have been a vertical model of collaboration by which different production steps were assigned to different collaborators.

The workshop may also have been a dynamic environment, with collaborators changing over time. The key question, then, is how we account for such possibilities. Is authorship still to be located in the way earlobes are painted, or are other diagnostic features involved?

It is at this point that artificial intelligence can offer a new perspective, not as a replacement for art historical expertise, but as a tool that invites us to reconsider how attributional evidence is structured and interpreted. Its use in art historical analysis has introduced new methodologies within the broader field of Digital Humanities, as well as new research practices situated between art history and computational science. Such interdisciplinary approaches are already beginning to demonstrate how computational methods and traditional art history can complement one another. In a recent study of a painting attributed to Anthony van Dyck, for example, AI analysis was combined with connoisseurship and historical research to examine questions of attribution. (Büttner, De Feudis and Popovici 2024).

Central to this process is dataset curation: the construction and organization of image datasets for machine learning applications in cultural heritage. Unlike standard computer vision datasets, art historical datasets are embedded in interpretative and historiographical frameworks. Their selection, classification, and contextualization are therefore not neutral technical procedures, but scholarly decisions that shape what an algorithm can learn and how it interprets works of art.

In practical terms, these datasets provide the visual material from which the AI learns. By being trained on groups of securely attributed works, the algorithm identifies recurring visual patterns associated with a given artist, as well as stylistically related artists, workshop productions, modern imitations, and even AI generated art in the style of the artist. When an unknown or disputed painting is then submitted for analysis, the system compares its visual features with those learned from the dataset and produces a probability-based assessment. The result is not a verdict, but an additional form of evidence that must be interpreted in relation to technical, historical, and connoisseurial arguments.

For this reason, the creation of datasets must be understood as both a computational and a curatorial practice, requiring collaboration between computer scientists, art historians, and domain experts. In the present study, this methodological awareness is essential, since the algorithmic results depend not only on the architecture of the model, but also on the art historical assumptions embedded in the dataset itself.

A further methodological layer concerns the material condition of the paintings themselves. Restorations, overpainting, and later interventions leave visual traces that may affect AI authenticity analysis, since the algorithm processes the painting as visual data. For this reason, the present study also examines how conservation history influences the computational results. By incorporating restoration documentation and applying digital masking to restored areas, the analysis seeks to reduce the impact of non-original material and to reconcile conservation history with the methodological requirements of computer-based art historical research.

2. Art-Historical Context and Methodological Framework

This section outlines our methodological framework and examines how art historical questions intersect with artificial intelligence throughout the process, particularly in the construction of authentic and contrast datasets. By addressing both the art historical and computational dimensions, the section explains how the datasets were created and how the AI model was developed, trained, and evaluated in relation to the Van Eyck corpus.

2.1 Computational Approaches to Art Authentication

The use of computational image analysis in questions of art authentication has developed significantly over the last decade, particularly through the application of convolutional neural networks, transfer learning, and, more recently, transformer-based architectures. Early work in this field demonstrated that machine-learning systems could be trained to recognize artist-specific visual patterns across digitized paintings, including at different image scales, where local details may prove especially informative for attribution (Van Noord and Postma 2017). More recent studies have extended these approaches to problems closer to authentication, including the analysis of surface topography, brushstroke-related features, and the distinction between authentic and non-authentic image sets (Ji, et al. 2021). In this context, deep learning has been framed not as a replacement for connoisseurship, provenance research, or technical examination, but as an additional analytical tool that may support art-historical judgment (Ugail, et al. 2023).

The broader literature also shows that computational approaches to painting classification now range from general surveys of fine-art image classification to more targeted studies of forgery detection and authentication. Recent survey work provides a useful overview of image-processing, machine-learning, and deep-learning methods used for fine-art classification (Rathi, et al. 2025). Other studies illustrate the growing interest in adapting neural networks to the specific problem of distinguishing genuine from forged works (Sundravadivelu, et al. 2024). Taken together, these studies indicate that AI-based analysis is best understood as a comparative and probabilistic instrument: it can identify visual regularities within a defined corpus, but its results remain dependent on the quality of the dataset, the relevance of the comparison material, and interpretation within a broader art-historical and material-technical framework.

2.2 Training Principles and Algorithmic Learning

Art Recognition employs a computer-based approach to image analysis, utilizing Vision Transformers (ViTs) for the authentication of artistic works. This methodology has been extensively investigated by Schaerf et al. (Schaerf, Postma and Popovici 2024). In contrast to traditional rule-based image analysis, Vision Transformers do not rely on manually selected stylistic markers. Instead, the system learns visual relationships directly from training data through repeated exposure to examples.

The training process is based on the comparison between two categories of images grouped in a positive dataset and a negative, or contrast, dataset. The positive dataset consists of works that are securely attributed to the artist under investigation, while the contrast dataset contains works by related artists, workshop productions, imitators, copies, and synthetic images. The purpose of this dual structure is not simply to teach the algorithm what a Jan van Eyck painting looks like, but also what it does not look like. Through this comparative process, the model gradually identifies recurring visual structures and characteristics that distinguish one artistic corpus from another.

The training phase therefore constitutes a process of statistical learning rather than stylistic interpretation in the traditional art historical sense. During training, the algorithm analyses thousands of image fragments extracted from the datasets and adjusts its internal parameters in order to improve its ability to differentiate between authentic and non-authentic examples. Once the training process is completed, the model can evaluate previously unseen works and generate a probability-based attribution assessment. In the case of early modern artworks, however, the quality and structure of the dataset become particularly important. Since paintings are embedded in complex histories of restoration and conservation, workshop collaboration, copying, and material alteration, the construction of the training corpus directly affects the interpretability of the results. The following subsections examine the construction of the authentic and contrast datasets employed in the Van Eyck training model.

2.3 Constructing the Authentic Dataset

The initial phase entailed the identification of works attributed to Jan van Eyck that are regarded by scholars and experts as authentic. In this instance, the construction of the authentic set was facilitated by the employment of exclusive material and writings by Professor Till-Holger Borchert. The authentic dataset is the most relevant segment of the pre-training phase in terms of structure. In this context, the approach to each training program is characterized by its complexity and variability. The specific nature of these variations is contingent on the particular painting or drawing that is the subject of analysis, as well as on the artist and workshop responsible for its creation and the historical context in which they operated.

Jan van Eyck is regarded as one of the most enigmatic exponents of the Flemish Primitive painting tradition. One of the masterpieces of international renown attributed to him is the *Ghent Altarpiece*, completed in 1432. However, the question remains as to whether the scenes contained within the panels of this artwork can be incorporated into the *Authentic set*. The inclusion or exclusion of this work is motivated by several factors.

Firstly, in 1934, the panels representing *John the Baptist* and the *Just Judges* were stolen. The latter was replaced by a copy executed by Jef van der Veken (1872–1964), while the other panel was later recovered and restored. Secondly, the collaborative nature of the altarpiece's production must be considered, as the work was executed by Jan van Eyck, his brother Hubert van Eyck described as *"[...] maior quo nemo repertus [...]"* (greater than whom no one was found), and their workshop.

Early Netherlandish painting was often produced in workshop environments in which multiple assistants contributed to a single work. This raises important questions regarding authorship and stylistic attribution, particularly in the context of art authentication in both traditional and artificial intelligence–based approaches. Moreover, very little is known about the artistic production of the probably older master Hubert van Eyck, to the extent that he was considered by some scholars a legendary figure. His recognition was largely prompted by the now famous quatrain found on the frame of the altarpiece during the nineteenth-century cleaning: "*Pictor Hubertus Eyck, maior quo nemo repertus, Incepit pondus; quod Johannes arte secundus Frater perfecit, Judoci Vijd prece fretus. Versu sexta Mai vos collocat acta tueri.*" (Borchert forthcoming).

In response to this complexity, the present case study involved the creation of two distinct authentic datasets, hereafter referred to as the Expanded Authentic Dataset and the Core Authentic Dataset. The objective of this dual dataset strategy was to address uncertainties in the attribution history while maintaining a rigorous methodological framework for training and evaluating machine learning models. The Expanded Authentic Dataset was conceived as a comprehensive body of material reflecting the broader Van Eyck artistic environment, including works attributed to Jan and Hubert van Eyck, as well as works associated with their immediate workshop context. The Core Authentic Dataset, by contrast, was constructed according to a more restrictive principle and was limited to painted works unanimously attributed to Jan van Eyck in current art historical scholarship.

The Expanded Authentic Dataset therefore sought to balance methodological rigor with art historical completeness. It includes the altarpiece traditionally assigned to Jan van Eyck and Hubert van Eyck, as well as works generally associated with their immediate workshop context. However, certain objects were deliberately excluded in order to preserve the methodological consistency of the dataset. Notably, the panel executed by Jef Van der Veken was excluded by default due to its modern authorship and therefore incompatible material and historical context.

At the same time, the dataset intentionally retained several works whose art historical interpretation presents particular methodological interest. Among these are the diptych commonly titled *Crucifixion of Christ and the Last Judgment*, dated between 1435 and 1440 and attributed to Jan van Eyck and his workshop, and the work known as *Saint Barbara* (1437) (Borchert forthcoming). Although the diptych has been discussed in scholarship in relation to workshop participation and questions concerning the degree of direct authorship by Jan van Eyck, it remains a significant component of the broader corpus associated with the Van Eyck circle. Including this work in the dataset allows the training material to reflect the historical reality of collaborative workshop practices that characterized early Netherlandish painting.

Similarly, *Saint Barbara* was included in the first authentic training dataset despite its distinctive visual and technical characteristics. The work has no contradictory attribution history and is traditionally attributed to Jan van Eyck. Nevertheless, it differs from most paintings in the corpus because it appears closer to a highly finished drawing than to a completed painted panel. Its unusual technique and unfinished state make it an atypical object within the Van Eyck oeuvre. Precisely for this reason, its inclusion provides an opportunity to test how machine learning models respond to variations in medium and execution within a single artistic corpus.

The construction of the first authentic dataset therefore sought to balance methodological rigor with art historical completeness with a total of forty-eight images. By excluding clearly incompatible objects, such as modern reconstructions, while retaining historically relevant works with distinctive technical features, the dataset attempts to mirror the complexity of the Van Eyck corpus as it is understood in contemporary scholarship.

The *Core Authentic Dataset* was constructed through a more targeted methodological approach, designed to reduce attributional ambiguity and ensure a higher degree of consistency within the training material. Unlike the broader dataset described above, this corpus was limited exclusively to painted works that are unanimously attributed to Jan van Eyck in current art historical scholarship. The purpose of this restriction was to create a dataset characterized by a stronger attributional consensus, thereby providing a more controlled foundation for computational analysis. Following this methodology, the number of images that were used for training was restricted to twenty.

Despite this careful selection process, the preparation of the dataset still presented several methodological challenges. Even among works whose authorship is widely accepted, the material history of individual paintings can significantly affect the visual information available for computational analysis. One notable example is the panel *Portrait of a Man (Self Portrait?)* dated 1433 and currently held in the National Gallery in London. The painting has a well-documented provenance history and is known to have formed part of the celebrated collection of the English magistrate and collector Thomas Howard, Earl of Arundel (1585–1646). However, technical and art historical investigations have revealed that the painting has undergone alterations affecting the background of the composition. As Borchert has suggested, the background was possibly executed in a blue tonal range and at a later stage, this area appears to have been repainted in black (Borchert forthcoming). In order to address this issue, during the training, a masking procedure was applied specifically to the altered background area of the portrait where the most intrusive alteration occurred.

2.4 Constructing the Contrast Dataset

The contrast dataset was constructed to provide the model with a structured body of non-authentic comparative material against which the authentic Van Eyck corpus could be tested. It includes three principal categories: works by related artists, understood as stylistic and technical proxies for Van Eyck's artistic environment; imitations that closely reproduce the master's visual language; and synthetic images, or cultural deepfakes, generated from authentic works and workshop-related material.

Related artists is a definition that in this particular context indicates *proxies*, described as works by followers of the master subject of the investigation or by collaborators, both of which are closely linked by similarities in style and technique. This part of the dataset enables the model to learn not only what belongs to the authentic corpus, but also which visually similar works should not be classified as authentic (Schaerf, Postma and Popovici 2024).

In the case of the Jan van Eyck training, artworks by multiple masters were sourced for a total of eight artists. The ratio of artworks to be selected is contingent on the quantity of authentic artworks on which the model has been trained. Consequently, adjustments have been made to the training process, with the model being trained on the expanded authentic dataset and the core authentic dataset. The following painters are to be considered: Robert Campin (Valenciennes, c. 1370 – Tournai, c. 1444), Petrus Christus (Baarle-Hertog, ? – Bruges, c. 1475), Niccolò Antonio Colantonio (Naples, c. 1420 – ?), Lluis Dalmau (Valencia, c. 1400 – 1461), Jacques Daret (Tournai, c. 1400 – 1470), Gerard David (Oudewater, c. 1460 – Bruges, 1523), Juan Rexach (Barcelona, c. 1411 – Valencia, ?) and Rogier van der Weyden (Tournai, c. 1399 – Brussels, 1464).

2 Synthetic reproduction of Jan Van Eyck *Annunciation*, authentic artwork in the National Gallery of Art, Washington

Imitations can be defined as works of art that reproduce, in a highly faithful and often nearly identical manner, the stylistic and compositional characteristics of a specific master. These works are typically produced by artists who deliberately model their work on a known prototype, aiming to replicate its visual appearance, technique, and stylistic language. Unlike other categories of artworks, where temporal proximity to the artist's lifetime may carry interpretative significance, the key criterion for inclusion here is stylistic similarity.

A well-known example that illustrates this category is the twentieth-century reconstruction of the panel *The Just Judges*, originally part of the *Ghent Altarpiece*. The original panel, attributed to Jan van Eyck and his workshop, was stolen in 1934 and has never been recovered. In 1939, the restorer Jef Van der Veken produced a replacement panel intended to visually reconstruct the missing work within the altarpiece. This reconstruction was executed as a highly faithful copy based on historical photographs, surviving preparatory material, and stylistic knowledge of early Netherlandish painting. Although created in the twentieth century, the panel was deliberately designed to imitate the visual language of the fifteenth-century original as closely as possible.

Synthetic images are *cultural deepfakes* of original artworks which can also be described as *doppelgangers* of authentic works of art.[3] In this case, oil paintings by Jan van Eyck were used as models for the generation of the replicas, as shown in Fig. 2. The synthetic images were generated using an image-to-image workflow, rather than producing an image entirely from text. This process was further guided by a custom lightweight model adaptation that allows the system to reproduce visual characteristics learned from an image corpus (Podell, et al. 2023); (Hu, et al. 2022).

2.5 Model Training and Testing

Before training, it is important to clarify what the model is understood to learn in this context. The AI model does not assess authorship through a single visible motif or isolated stylistic marker. Rather, it learns recurring visual patterns across the corpus.

The process begins with a pre-trained model, already capable of recognizing general visual structures, which is then further trained to identify characteristics associated with a specific artist. These characteristics can be understood at several levels. At a lower level, the model may respond to contours, brushstrokes, color relationships, and surface textures. At a middle level, it may learn shapes, forms, and spatial relationships that recur across the selected works. At a higher level, it may register broader compositional structures and aspects of the artist's overall manner. In this sense, model training is not directed toward identifying one fixed stylistic feature, but toward recognizing patterns of visual correspondence across the corpus.

Following data collection and dataset construction, the model was trained to distinguish between the authentic Van Eyck corpus and the contrast dataset. During this phase, the algorithm analyzed both entire images and image fragments extracted from the selected works, learning visual patterns associated with each category. The aim of the training was not to identify a single stylistic feature, but to model recurring visual relationships across the corpus.

To evaluate the trained models, 10% of each dataset was held out from training and used as a testing set. After training, the Core Dataset models achieved an average accuracy of 89.61% on the testing set, while the Expanded Dataset models achieved an average accuracy of 88.95%. Since these accuracies were calculated on different datasets, they should not be directly compared. Considered individually, however, both model groups achieved high performance metrics.

At the same time, these results must be understood in relation to the visual and material characteristics of the images on which the models were trained. In the case of Jan van Eyck's oeuvre, it may be argued that craquelure visible on the painted surfaces of the works analyzed here exerts a potentially distortive effect on visual interpretation.[4] The present AI-based analysis operates exclusively on digital images of the artworks and therefore constitutes a surface-level examination. As such, it does not directly account for the material support of the painting, except in cases where this becomes visible through areas of paint loss. Nevertheless, the nature of the support remains indirectly significant, as different supports are known to produce distinct cracking patterns that may alter the appearance of the pictorial surface.

To mitigate such interference, the models employed in this study have been trained using strategies designed to minimize the influence of incidental surface features, including craquelure. These measures aim to prevent the model from developing spurious correlations or heuristic shortcuts based on support-induced patterns, thereby ensuring that the analysis remains focused on relevant pictorial and stylistic features.

3. Case Study: Saint Francis Receiving the Stigmata in Turin

This section applies the methodological framework outlined above to the case study of *Saint Francis Receiving the Stigmata,* measuring 29.3 x 33.4 cm, in Turin. Particular attention is given to the painting's conservation state, the use of masking as a preprocessing strategy, and the comparative outcomes generated by the Expanded and Core Authentic Datasets. The section concludes with an examination of heatmap visualisations in order to assess which regions of the painting contributed most significantly to the predictions.

3.1 Restoration and Conservation History and Methodological Implications

The Turin panel has undergone several restoration campaigns that have significantly affected its current visual appearance. These interventions, documented in restoration reports and archival materials, have altered portions of the pictorial surface and consequently represent an important factor to consider in any computational image analysis of the painting. As discussed by Carlenrica Spantigati, three major restoration campaigns were carried out during the twentieth century, each introducing different levels of intervention and modification (Spantigati 1997). From the perspective of computer vision analysis, the restoration and conservation history of the object becomes a crucial methodological variable. The presence of repainted areas, retouching, filling and conservation treatments can introduce visual features that are not part of the original pictorial layer, potentially influencing the behavior of AI models trained to identify stylistic patterns. For this reason, the present study integrates restoration and conservation history into the dataset preparation process. In order to reduce the impact of non-original material on computational analysis, masking techniques are employed to exclude areas that have undergone pictorial changes. By isolating portions of the painting that are likely to preserve the original pictorial execution, this targeted strategy seeks to improve the reliability of stylistic feature extraction while acknowledging the complex material history of the object.

As Spantigati explains, the first major restoration campaign took place around 1952. The intervention is extensively documented in the restoration archive of the *Soprintendenza per i Beni Artistici e Storici* in Piedmont, which preserves a substantial corpus of photographic documentation related to the conservation work. These archival photographs provide an important record of the condition of the panel prior to and during the intervention, allowing scholars to reconstruct the extent of the treatments that were performed. According to the documentation, the restoration addressed a number of structural and pictorial issues affecting the painting. Among the areas affected by retouching were St. Francis's hair, parts of the central area of the composition, and sections of the green landscape located between the two friars. The habit (saio) of Brother Leo was also subject to repainting, indicating that the intervention extended to multiple figurative and landscape elements within the composition (Spantigati 1997).

A second restoration campaign was conducted in 1970. Compared to the earlier intervention, this campaign focused primarily on the cleaning of the pictorial surface. The treatment involved the removal of accumulated surface deposits and aged varnish layers in order to restore the chromatic clarity of the painting. Importantly, however, the integrations introduced during the previous restoration were not removed. Instead, they were preserved as part of the existing condition of the artwork (Spantigati 1997).

The third restoration campaign, undertaken in 1982, constituted a comparatively radical intervention. The treatment included a systematic cleaning aimed at reassessing and, where necessary, reversing earlier restoration measures. In this context, non-original overpaint, detectable primarily under ultraviolet illumination, was selectively removed to recover the legibility of the original pictorial layers. Where losses in the paint surface were identified, these areas were reintegrated through new retouching designed to restore visual continuity while remaining distinguishable upon close examination. Finally, a layer of varnish was applied to the panel in order to protect the surface and stabilize its appearance (Spantigati 1997).

From the perspective of computational analysis, these successive restoration campaigns have produced a complex stratigraphy of pictorial interventions that directly affects the visual data available for algorithmic processing. Areas that have undergone repainting or retouching may display characteristics that differ from the original paint layers. Consequently, when constructing image datasets for machine learning analysis, it becomes necessary when possible to account for these alterations in order to avoid misleading stylistic signals. The following section explains how this issue was addressed in practice through the use of digital masking as a preprocessing strategy.

3.2 Digital Masking as Image Preprocessing

Masking is employed as a prediction technique to exclude regions of a painting that may compromise the reliability of AI-based analysis, particularly areas affected by restoration, surface damage, or overpainting. As illustrated in Fig. 2, the excluded regions are represented in white, clearly distinguishing them from the areas retained for study. These areas were identified on the basis of conservation documentation, particularly the restoration history outlined first by Spantigati and later documented within the VERONA project (Spantigati 1997); (Royal Institute for Cultural Heritage 2018).

3 Masked image of Jean van Eyck *Saint Francis Receiving the Stigmata*, 1435–1440, 29.3 x 33.4 cm, Galleria Sabauda, Turin.

In the remainder of this study, the term *"unmasked"* will refer to the painting in its current visible state, including areas affected by restoration or later intervention, while the term *"masked"* will refer to the digitally adjusted version in which these areas were excluded from the analysis.

3.3 Comparative Results: Expanded and Core Dataset Models

Both the unmasked and masked versions of the Turin panel were subjected to a twofold AI authenticity analysis. Each version was analysed using models trained on the Expanded and Core Authentic Datasets, comprising forty-nine and twenty images, respectively. This approach was designed to assess the extent to which restored areas influence the extraction of stylistic features by the AI model, while also evaluating how dataset selection affects the results. The results are presented in Table 1 and discussed below.

	Probability of authenticity on masked image	Probability of authenticity on unmasked image
Expanded Authentic Dataset	73.8%	17.9%
Core Authentic Dataset	57.6%	47.7%

Table 1 Comparative results of AI authenticity analyses on masked and unmasked images of Saint Francis Receiving the Stigmata (Turin). "Expanded Authentic Dataset" refers to the larger set of authentic and non-authentic, whereas "Core Authentic Dataset" denotes the smaller set of artworks. "Unmasked" indicates the artwork in its current visible state, while "masked" refers to the image with areas affected by restoration excluded from the analysis.

The model trained on the Expanded Dataset produced markedly divergent results when applied to the masked and unmasked versions of the painting. In the condition where masking was applied, the model achieved a positive classification score of 73.8%. This result suggests that, when the analysis is restricted to regions more likely to preserve the original pictorial execution, the algorithm is able to detect stylistic features consistent with the Van Eyck corpus with a high degree of confidence. By contrast, when the unmasked image was analysed, thus including all restored and altered areas, the positive classification rate dropped sharply to 17.9%. This substantial decline indicates that the presence of later interventions introduces visual noise that significantly interferes with the model's ability to extract reliable stylistic patterns. Repainted or retouched areas, which do not correspond to the original painted surface, in some instances appear to distort the feature space on which the algorithm relies.

We now turn to the second training iteration, conducted using the Core Dataset. This methodological choice can be considered unorthodox, as conventional machine learning practice suggests that model confidence generally increases with the size of the training dataset. Nevertheless, the restricted dataset was deliberately designed to prioritize attributional certainty, reflecting the art historical principle that content quality can, in some cases, outweigh sheer quantity in computational analyses of cultural heritage objects.

Similar to the previous experimental setup, the training was performed on both the masked and unmasked versions of the Turin panel. The outcomes of this smaller-scale analysis proved surprising, albeit for reasons distinct from those observed in the larger dataset. The masked version of the panel yielded a positive authenticity classification of 57.6%; in contrast, the unmasked version produced a comparatively close result of 47.7%. This narrowing of the gap between masked and unmasked conditions represents a noteworthy divergence from the patterns observed in the full forty-nine-image dataset. While masking in the larger dataset substantially increased classification accuracy, in the restricted twenty-image set the effect was attenuated, suggesting that the limited number of training examples constrains the model's ability to discriminate between original and restored features.

Additionally, this outcome illustrates the delicate balance between dataset size and data integrity: although smaller datasets may ensure higher attributional certainty, they may also reduce the algorithm's capacity to leverage masking interventions effectively. It is pivotal to note that, despite yielding less confident predictions, the core dataset produces results that are consistent with both those of the expanded dataset and prevailing art historical scholarship.

4. Case Study: Saint Francis Receiving the Stigmata in Philadelphia

This section examines the Philadelphia version of *Saint Francis Receiving the Stigmata*. Unless otherwise stated, the methodological procedures outlined in the previous section also apply to the present analysis.

4.1 Material and Conservation Context

As noted by Rishel, the panel depicting *Saint Francis Receiving the Stigmata,* now in Philadelphia, was acquired in London in 1894 by the American lawyer and art collector John Graver Johnson (1840–1917) and entered the museum's collection in 1932. The artwork arrived in London via Portugal, most likely Lisbon, and it was recorded in the collection of Lord Heytesbury (1779–1860) in 1843 as by the hand of Albrecht Dürer (1471–1528), an attribution that was soon thereafter abandoned. The composition closely corresponds to that of the Turin version, although the Philadelphia panel is significantly smaller, measuring 12.4 × 14.6 cm. Of particular note is its composite structure: the paint layer is executed on parchment adhered to a wooden support assembled from five joined elements. The painting underwent restoration in New York in 1906 by Roger Fry (1866–1934), while in the Johnson collection. (Rishel 1997)

It has been documented by Caroline Elam how, during this period Fry was also engaged in selling works of art to Johnson, some of which had reportedly been declined by the Metropolitan Museum of Art, particularly after 1907. Fry instructed his assistant, Bryson Burroughs, in methods of cleaning and restoration—practices that would later attract considerable criticism. In the case of Saint Francis Receiving the Stigmata, however, the intervention was contemporaneously regarded as a "successful cleaning," as reported in The Burlington Magazine (Elam 2019, 105). Prior to restoration, Saint Francis Receiving the Stigmata exhibited significant alterations to its original format, including an added extension that enlarged the sky above the saint and extensive overpainting along the margins. The removal of these later additions revealed a burgundy-painted border framing the composition, along with traces of gilding. In this regard, Rishel and Elam cite an observation made by Fry twenty years later in a letter, in which he associated the burgundy and gold borders with practices derived from manuscript illumination (Rishel 1997, 11n8); (Elam 2019, 105). Overall, the Philadelphia version appears to have undergone comparatively less invasive intervention to its original paint surface.

Similarly to the Turin painting, this work has also been connected to Anselmo Adorno's document recording donations to his daughters, described as follows:

> "*Item, I gave to each of my daughters, to be theirs, to wit, Marguerite, Carthusian, and Louise, Sint-Truiden, a picture wherein Saint Francis in portraiture from the hand of Master Jan van Eyck, and make the condition that in the shutters of the same little pictures be made my likeness and that of my wife, as well as can be made.*"[5]

As Rishel pointed out, as reported by James Weale, it must be emphasized that Adorno could not have acquired the painting directly from Van Eyck, as he was barely fifteen years old at the time of the master's death in 1441 (Rishel 1997). This discrepancy introduces further ambiguity into the art historical traceability and geographical provenance of the works. Moreover, it may also be argued that the two paintings in question do not fully correspond to the genre of portraiture as described in the aforementioned note, thereby further complicating the relationship between the primary source and the artworks.

4.2 Comparative AI Authenticity Analyses

As in the Turin case study, masking was employed as a preprocessing strategy in order to exclude regions affected by restoration and later intervention. The masked areas were identified on the basis of conservation documentation and visible alterations associated with earlier restoration campaigns. Fig 4 shows the masked version of the Philadelphia panel used for the computational analysis.

The computational analysis of the Philadelphia panel produced markedly different results from those observed in the Turin case study. In both the Expanded and Core Dataset conditions, the model generated low authenticity probabilities for both the masked and unmasked versions of the painting, consistently classifying the work as non-authentic. Unlike the Turin panel, where masking substantially altered the classification outcome, the Philadelphia results remained comparatively stable regardless of whether restored regions were excluded from the analysis. The results are reported in Table 2.

4 Masked image. Attributed to Jean van Eyck *Saint Francis Receiving the Stigmata*, 1440–1445, 12.4 × 14.6 cm, Philadelphia Museum of Art, John G. Johnson Collection, Philadelphia

When trained on the Expanded Dataset, the model produced authenticity probabilities of 24.47% for the masked image and 25.43% for the unmasked image. Similarly, the Core Dataset yielded probabilities of 32.69% and 32.58%, respectively. The minimal divergence between masked and unmasked conditions suggests that restoration layers exert a significantly weaker influence on the model's decision-making process than in the Turin panel.

	Probability of authenticity on masked image	Probability of authenticity on unmasked image
Expanded Authentic Dataset	24.47%	25.43%
Core Authentic Dataset	32.69%	32.58%

Table 2: Comparative results of AI authenticity analyses on masked and unmasked images of *Saint Francis Receiving the Stigmata* (Philadelphia). Dataset terminology and masking definitions follow Table 1.

A clarification is necessary regarding the probability values reported in media coverage of this project. Earlier results discussed in the press indicated that the Philadelphia panel was classified as 91% negative and the Turin panel as 86% negative with respect to Van Eyck's authorship. (Alberge 2026) These figures were obtained using a different analytical configuration from the one presented in the present study. The media-reported analysis was conducted exclusively on the post-restoration visible versions of the paintings and employed only the Expanded Dataset. Furthermore, the models presented here were trained using a more conservative optimization criterion. In machine learning, this criterion determines how strongly the model is rewarded for correct classifications and penalized for incorrect ones during training. The revised approach was designed to reduce overconfidence and to favour greater caution when assigning authenticity probabilities. For these reasons, although the media-reported figures correspond most closely to the Expanded Dataset results presented here, they should not be directly compared with them; as a result of the more conservative optimization criterion, the probabilities reported in the present study are generally lower than those reported previously in the media.

5. AI Explainability and Heatmaps

Heatmaps are instrumental in elucidating the regions of the artwork that are pivotal to the AI's decision-making process and resultant outcomes. The color scale, ranging from blue to bright red, serves as a metric for assessing the relative influence of each area in the analysis (Fig. 5, Fig. 6). The regions identified by red are of the utmost importance, suggesting that the AI system allocated a higher proportion of its attention to these areas during the decision-making process. The area's relevance undergoes a transition from orange to yellow, indicating a decrease in its influence. Meanwhile, turquoise and blue areas exhibit less contribution. Fundamentally, the most influential areas for the final decision are colored red, while the least influential areas are colored blue.

These heatmaps are generated with an algorithm called ScoreCAM (Wang, et al. 2020). The method works by masking small areas of the painting and observing how each masked region affects the model's prediction. If masking a particular region produces a substantial change in the predicted probability, that region is considered important for the model's decision.

It is important to note that the heatmaps only indicate which regions contribute to the final prediction the most, and not whether the final classification is authentic or not authentic. For an authentic prediction, red areas indicate regions that align with the artist's style, while for a non-authentic classification, the red areas indicate the areas that misalign with the artist's style the most. Therefore, the heatmaps must always be interpreted while considering the prediction result of the model on the image. It is also important to acknowledge that a direct comparison between different heatmaps is not feasible from an empirical standpoint. Due to the individual assignment of colors, an area colored deep red in

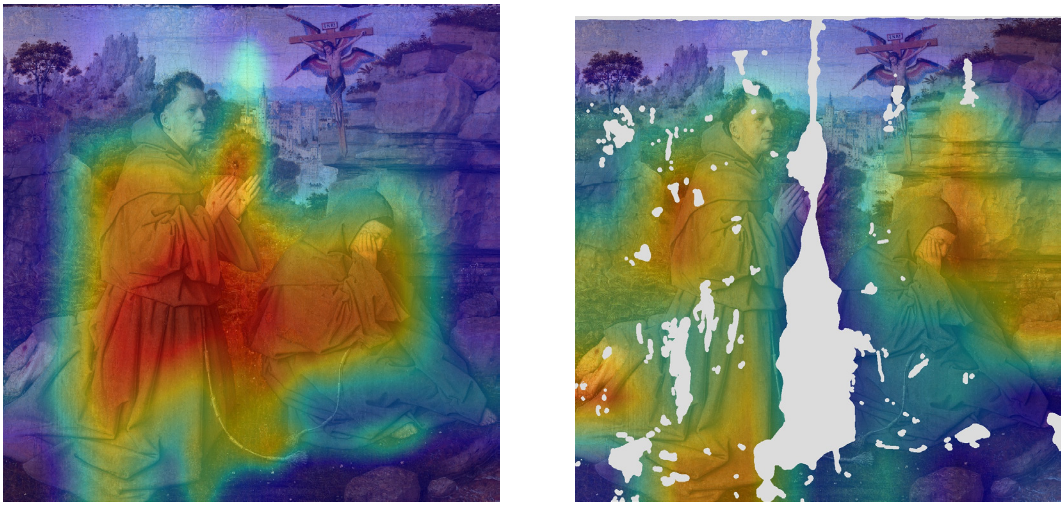

5 Heatmaps of unmasked (left) and masked (right) versions of Jan van Eyck *Saint Francis Receiving the Stigmata*, 1435–1440, 29.3 x 33.4 cm, Galleria Sabauda, Turin, following training on the Expanded Dataset. The model produced an authenticity probability of 17.9% for the unmasked image and 73.8% for the masked image. Red areas indicate regions that contributed most strongly to the model's prediction; in the unmasked image, this contribution supports a non-authentic classification, whereas in the masked image it supports an authentic classification. See text for further discussion.

one heatmap might have a lower importance score compared to a yellow area in a different heatmap.

In this section, the heatmaps are presented for the models trained on the Expanded Dataset. This dataset is particularly relevant for the present discussion because it produced a pronounced contrast between the masked and unmasked analyses in the case of the Turin panel. The heatmaps therefore provide insight into which visual features contributed to this shift in classification and how the exclusion of restored regions altered the basis of the prediction.

Fig. 5 shows the heatmaps for the unmasked and masked versions of the Turin panel. In the unmasked image, shown on the left, the red zones correspond to the central portion of the painting, particularly the vegetation between the two figures, Francis on the left and Leo on the right, as well as their respective *saio*. This is noteworthy, since scholars have identified the central portion of the Turin painting as especially susceptible to alteration. In the masked image, shown on the right, the heatmap displays a more complex distribution, with multiple centres of interest across the composition.

In the Turin case, the heatmaps reveal a clear difference between the unmasked and masked analyses. In the unmasked image, the model concentrates strongly on the central portion of the composition and on the kneeling friar, precisely in areas known to have undergone substantial repainting. This suggests that later interventions introduced visual features that significantly affected the model's decision-making process. Once these areas were masked, the distribution of attention changed, indicating that the overpainted regions had been exerting a measurable influence on the authenticity assessment.

The Philadelphia case presents a different pattern. As Fig. 6 shows, the heatmaps for the masked and unmasked images highlight largely similar areas. Moreover, the Philadelphia heatmaps closely resemble those observed in the unmasked Turin analysis. In all of these cases, which resulted in negative authenticity assessments, the model concentrates primarily on the drapery-related features.

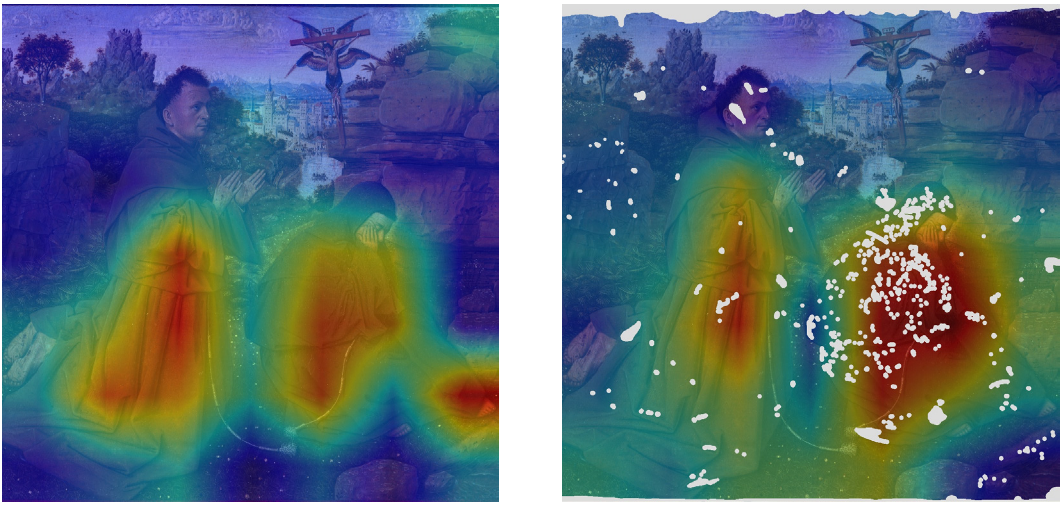

6 Heatmaps of unmasked (left) and masked (right) versions of attributed to Jan van Eyck Saint Francis Receiving the Stigmata, 1440–1445, 12.4 × 14.6 cm, Philadelphia Museum of Art, following training on the Expanded Dataset. The authenticity probabilities were 25.43% for the unmasked image and 24.47% for the masked image.

6. Conclusions

The present study examined the two fifteenth-century panels depicting *Saint Francis Receiving the Stigmata* traditionally associated with Jan van Eyck through a combined methodology integrating art historical research, conservation history, and Artificial Intelligence. Rather than treating AI as an autonomous attributional authority, the study approached computational analysis as an additional interpretative instrument capable of interacting with established methods of connoisseurship and technical examination.

The results demonstrate that restoration history and dataset construction play a decisive role in the interpretation of AI authenticity analyses. In the case of the Turin panel, the exclusion of heavily restored areas substantially altered the computational outcome, transforming a strongly negative authenticity assessment into a positive one. This finding suggests that later interventions and overpainting may significantly distort the stylistic signals detected by machine learning models. The Philadelphia panel, by contrast, produced consistently low authenticity probabilities under all testing conditions, with only minimal variation between analyses conducted on the full visible image and on the version excluding restored regions. The comparison between the two panels therefore highlights the extent to which conservation history can influence computational attributional results.

The study also demonstrates the methodological importance of dataset curation. The distinction between the Expanded and Core Authentic Datasets revealed how differing approaches to attributional certainty, workshop material, and corpus selection affect the behaviour of the model. These findings underscore that AI systems do not operate independently of art historical interpretation; rather, they are fundamentally shaped by the scholarly assumptions embedded within the training data itself.

In this sense, artificial intelligence should be understood as contributing to a broader interdisciplinary methodology situated between connoisseurship, restoration and conservation science, and the Digital Humanities.

In the case of the two *Saint Francis Receiving the Stigmata* panels, adding an extra layer to the investigation – namely the consideration of restoration history through digital masking – turned out to yield a more nuanced result. It may be possible that, in the future, we will encounter works of art whose identification ultimately requires one or more additional layers to do justice to the complex circumstances of their production. As it stands today, it is clear that the AI can be a useful corrective to purely connoisseurial debates, and that connoisseurial debates can also improve the algorithm. We will have to keep an open mind towards interdisciplinary collaboration to maximize the potential cross-fertilization that seems to be possible in the future.

References

Alberge, Dalya. 2026. "AI Analysis Casts Doubt on Van Eyck Paintings in Italian and US Museums." *The Guardian* , 7 February: https://www.theguardian.com/artanddesign/2026/feb/07/ai-analysis-van-eyck-paintings-turin-philadelphia.

Anderson, Jaynie. 2019. *The Life of Giovanni Morelli in Risorgimento Italy.* Milano.

Büttner, Nils, Alita De Feudis, and Carina Popovici. 2024. "The Interplay of Art Historical Connoisseurship and Artificial Intelligence in Authenticating a Painting Attributed to Anthony van Dyck." *Kunstgeschichte.*

Borchert, Till-Holger. forthcoming. "Katalog der Gemälde."

—. 2002. *Le siècle de Van Eyck: Le monde méditerranéen et les primitifs flamands, 1430–1530.* Amsterdam: Ludion.

Challéat, Claire. 2012. *Dalle Fiandre a Napoli: Committenza artistica, politica, diplomazia al tempo di Alfonso il Magnanimo e Filippo il Buono.* Rome: L'Erma di Bretschneider.

Coombs, Bryony, Jill Harrison, and Giovanna Guidicini. 2026. "Anselm Adornes: Travel, Trade, Cultural Exchange, and Intellectual Networks in Scotland, Bruges, and Jerusalem." Edited by Bryony Coombs, and Jill Harrison, eds. Giovanna Guidicini. *Late Medieval and Early Modern Studies* (Brepols) 31.

Cornelis, B., T. Ružić, E. Gezels, A. Dooms, A. Pižurica, L. Platiša, J. Cornelis, M. Martens, M. De Mey, and I. Daubechies. 2013. "Crack detection and inpainting for virtual restoration of paintings: The case of the Ghent Altarpiece." *Signal Processing, Volume 93, Issue 3* 605–619.

Dumolyn, J., and N. (eds) Geirnaert. 2024. *A celestial Jerusalem, The Adornes Estate and the Jerusalem Chapel.* Veurne.

Elam, Caroline. 2019. *Roger Fry and Italian Art.* Cambridge: Cambridge University Press.

Espin, Elsa. forthcoming. *Van Eyck and Valencia.*

Fernando Benito Doménech and José Gómez Frechina, eds., ed. 2001. *La clave flamenca en los primitivos valencianos: Museo de Bellas Artes de Valencia, del 30 de mayo al 2 de septiembre de 2001.* Valencia: Generalitat Valenciana.

Fransen, Barten. 2017. "Van Eyck in Valencia." Edited by Bart Fransen, Valentine Henderiks, Cyriel Stroo, and Dominique Vanwijnsberghe Christina Currie. *Van Eyck Studies: Papers Presented at the Eighteenth Symposium for the Study of Underdrawing and Technology in Painting, Brussels, 19–21 September 2012.* Löwen: Peeters. 468–478.

Geirnaert, Noel. 2000. "Anselm Adornes and His Daughters: Owners of Two Paintings of Saint Francis by Jan van Eyck?" In *Investigating Jan van Eyck*, by Delphine Cool, Susan Jones and Susan Foister, edited by S. Jones, S. Foister D. Cool, 163–168. London: Archetype Publications.

Hu, Edward J., Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. "LoRA: Low-Rank Adaptation of Large Language Models." *International Conference on Learning Representations (ICLR).*

Ji, F., M. S. McMaster, S. Schwab, G. Singh, and L. N. Smith. 2021. "Discerning the painter's hand: machine learning on surface topography." *Herit Sci 9* 152.

Jones, Susan Frances. 2014. "Jan van Eyck and Spain." *Boletín del Museo del Prado 32* (Museo del Prado) (50): 30–49.

Kemperdick, Stephan. 2012. "Copies after Jan van Eyck and his Predecessors." In *The Road to Van Eyck*, by Stephan Kemperdick and Lammertse, 80–81. Rotterdam: Museum Boijmans Van Beuningen.

Kenderdine, Sarah, and Lily Hibberd. 2025. "Lexicon." In *Deep Fakes. A Critical Lexicon for Digital Museology*, 40–57. London: Routledge.

Museo del Prado. 2015. *Museo del Prado.* 28 04. Accessed 6 5, 2026. https://www.museodelprado.es/en/the-collection/art-work/saint-francis-of-assisi-receiving-the-stigmata/13b53b14-2396-471a-8b0f-817359bb77ee.

Natale, Mauro, ed. 2001. *El Renacimiento mediterráneo: Viajes de artistas e itinerarios de obras entre Italia, Francia y España en el siglo XV: Madrid, Museo Thyssen-Bornemisza, del 31 de enero al 6 de mayo de 2001; Valencia, Museu de Belles Arts de València, del 18 de mayo al 2 de septiembre de 2001.* Madrid: Fundación Colección Thyssen-Bornemisza.

Nuttall, Paula. 2004. *From Flanders to Florence. The Impact of Netherlandish Painting, 1400–1500.* New Haven: Yale University Press.

Ostmeyer, Johan, Ludovica Schaerf, Pavel Buividovich, Tessa Charles, Eric Postma, and Carina Popovici. 2024. "Synthetic images aid the recognition of human-made art forgeries." *PLOS ONE 19* 1–15.

Pérez Preciado, Juan José. 2024. *Fifteenth-Century Netherlandish Painting at the Museo Nacional del Prado: Catalogue Raisonné.* Madrid: Museo del Prado.

Podell, Dustin, Zion Englis, Kyle Lace, A. Blattman, Tim Dockhor, Jonas Mulle, Joe Penn, and Robin Rombach. 2023. "SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis." *ArXiv abs/2307.01952.*

Rathi, Vishwas, Abhilasha Sharma, Aditya Venkata Nithin, Amit Kumar Singh, and Brij B. Gupta. 2025. "A survey of computational techniques for fine art painting classification." *Image and Vision Computing Volume 161.*

Rishel, Joseph J. 1997. "The Philadelphia and Turin Paintings: The Literature and Controversy over Attribution." In *Jan van Eyck: Two Paintings of Saint Francis Receiving the Stigmata*, by J. R. J. Van Asperen de Boer, Bé Kenneth, Marigene H. Butler, Peter Klein, Katherine Crawford Luber, Joseph J. Rishel, Maurits Smeyers, James Snyder and Carlenrica Spantigati, edited by J.R.J Van Asperen de Boer, 3–12. Philadelphia: Philadelphia Museum of Art.

Rohlmann, Michael. 1994. *Auftragskunst und Sammlerbild. Altniederländische Tafelmalerei im Florenz des Quattrocento.* Alfter: Verlag und Datenbank für Geisteswissenschaften.

Rosenauer, Artur. 1993. *Art of the Gothic Era: Architecture, Sculpture, Painting.* London: Laurence King Publishing.

Royal Institute for Cultural Heritage. 2018. *Closer to Van Eyck: Rediscovering the Ghent Altarpiece.* Accessed 6 5, 2026. https://closertovaneyck.kikirpa.be/ghentaltarpiece/.

Schaerf, Ludovica, Eric Postma, and Carina Popovici. 2024. "Art authentication with vision transformers." *Neural Computing & Applications 36* 11849–11858.

Spantigati, Carlenrica. 1997. "Le stigmate di San Francesco della Galleria Sabauda di Torino." In *Jean van Eyck. Opere a confronto: Jan van Eyck (1390–1441)*, by Carlenrica Spantigati and Joseph J. Rishel, 27–43. Turin: Allemandi Editore.

Sundravadivelu, K., S. Karthika, R. R. Al-Fatlawy, J. R. Monish, and M. V. R. Sundari. 2024. "Rescaled Hinge Loss Function based Convolutional Neural Network for Identification of Art Forgery by Image Analysis." *First International Conference on Software, Systems and Information Technology (SSITCON).* Tumkur, India: IEEE. 1–5.

Ugail, Hassan, David G. Stork, Howell Edwards, Steven C. Seward, and Christopher Brooke. 2023. "Deep transfer learning for visual analysis and attribution of paintings by Raphael." *Heritage Science volume 11* 268.

Van Asperen de Boer, J. R. J., Bé Kenneth, Marigene H. Butler, Peter Klein, Katherine Crawford Luber, Joseph J. Rishel, Maurits Smeyers, James Snyder, and Carlenrica Spantigati. 1997. *Jan van Eyck: Two Paintings of Saint Francis Receiving the Stigmata.* Philadelphia: Philadelphia Museum of Art.

Van Noord, Nanne, and Eric Postma. 2017. "Learning scale-variant and scale-invariant features for deep image classification." *Pattern Recognition, Volume 61* 583–592.

Wang, Haofan, Zifan Wang, Mengnan Du, Fan Yang, Zijian Zhang, Sirui Ding, Piotr Mardziel, and Xia Hu. 2020. "Score-CAM: Score-Weighted Visual Explanations for Convolutional Neural Networks." *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW 2020)* (WA:IEEE) 24–25.

Endnotes

1 Positive attributions by Rishel/Spantigati 1997 and Crawford Luber 1998; negative attributions by Dhanens 1980; Smeyers 1997; Reynolds 2000.

2 The work is generally dated to circa 1510, thus postdating the Eyckian prototypes, which are conventionally situated between 1435 and 1445. (Museo del Prado 2015)

3 The term *cultural deepfakes* was coined in 2021 in relation to the exhibition *Deep Fakes* (Kenderdine and Hibberd 2025). For a discussion of the application of synthetic images to improve model precision in AI-based art authentication, see (Ostmeyer, et al. 2024).

4 In this regard, see the article *Crack Detection and Inpainting for Virtual Restoration of Paintings: The Case of the Ghent Altarpiece*, which addresses related issues in the context of the Ghent Altarpiece and virtual restoration in digitalized paintings (Cornelis, et al. 2013).

5 Translation from the original testament of Anselme Adornes (10 February 1470), as published by Joseph J. Rishel (Rishel 1997, 11n10).