

PHYSICAL SCIENCES

PATCH: A deep learning method to assess heterogeneity of artistic practice in historical paintings

Andrew Van Horn^{1,2,3*}, Lauryn Smith^{3,4}, Mahamad Salah Mahmoud⁵, Michael McMaster², Clara G. Pinchbeck³, Ina T. Martin^{2,6}, Andrew Lininger^{2,6}, Anthony Ingrisano⁷, Adam Lowe⁸, Carlos Bayod⁸, Elizabeth S. Bolman³, Kenneth Singer², Michael Hinczewski²

In the Renaissance and Early Modern period, paintings were largely produced by master painters who directed workshops of apprentices and others who often contributed to the piece. Discerning who created these masterworks and how they did so is a central question in technical art history and a nontrivial problem that machine learning can help solve by extending analysis to a microscopic scale. Analysis of workshop paintings presents a challenge, however, because information about the members of workshops and the processes by which artworks were created remains elusive. Thus, external examples are not available to train networks to recognize. Here, we present a novel machine learning approach we call pairwise assignment training for classifying heterogeneity (PATCH) that is capable of identifying individual artistic practice regimes with no external training data. We apply this method to two historical paintings by the Spanish Renaissance master, El Greco, and our findings regarding one of the works potentially challenge previous studies that assert that a considerable portion of the painting was completed by workshop members after El Greco's death. PATCH outperforms statistical and unsupervised machine learning methods in this complex pairwise comparison problem lacking "ground truth" data, making it potentially useful across similar cases in the social and natural sciences, including image segmentation in remote sensing, urban development and design, and anomaly detection manufacturing contexts, among others.

INTRODUCTION

There are many ways to create a painting. The process can be collaborative or solitary, and there is no end to the styles, tools, materials, and techniques that can be implemented. The evolution of what we call artistic practice—the technical, stylistic, material, and even physiological aspects of an artist's creative process—over time and its variation between artists remains a key element of technical art history scholarship. From the Late Middle Ages through the Early Modern period, painting was typically done in a bustling workshop (1). A master painter would establish a workshop and assemble a group of apprentices, journeymen, and specialist painters and even collaborate with associates to help fulfill commissions and create smaller works to support the enterprise between larger contracts.

Workshop practice adds a second layer of complexity to artistic practice. Master painters varied considerably in their managerial styles. Some performed much of the work themselves, while others dictated the composition but left many elements to the hands of workshop members [e.g., van Dyck (2)]. Still others involved apprentices in the conceptualization of a piece (3). Often, there was a hierarchy, with young apprentices typically assigned more mundane tasks such as the mixing of paints and the construction of canvases and more senior workshop members directed to paint portions of larger works in the master's style (4).

Because workshops varied so widely in their organization and management yet produced works with cohesive styles, deeper understanding

of variation in artistic practice can shed light on workshop practice and vice versa. Research into these key aspects of Renaissance and Early Modern art production is hindered by a lack of surviving textual sources regarding the structure, size, and membership of individual workshops as well as their day-to-day operation and the variation in artistic practices (4). While scholars have been able to reconstruct individual artists' workshops through a variety of methods, our collective understanding of the inner workings of workshops is still fragmentary (5). Artistic practice varies widely, encompassing materials, methods, styles, and individual contributions. Understanding who interacted with individual paintings and how they did so—discerning the artists at work in a painting, the interplay of their individual "hands," and the interaction of the artist's hand and eye with paint and brush (6)—is an important open question in art history (7).

Machine learning (ML) has recently emerged as a promising complement to traditional art historical analyses in applications including preservation, conservation, and the detection of forgeries (8–11), as well as in fighting illegal trafficking in antiquities (12). Much work involving style is focused on broad classification of works (13, 14), such as the use of a neural network to arrange works of the Western canon in the correct chronological order (15). Recently published work by some of the authors of the present study demonstrated the efficacy of ML for attribution of paintings by analyzing high-resolution topographic images of the surface texture of oil paintings by known authors (student painters) in a controlled experimental setting. In a supervised ML analysis, a convolutional neural network (CNN) was trained to sort microscopic surface textures among artists, assigning 1-cm² patches of paintings to the correct authors with ~95% accuracy (16). Thus, the method has shown that microscopic analysis can complement traditional art historical analyses of macroscopic features.

Unfortunately, supervised deep learning on topographic images cannot be used when there is no ground truth for the network to learn on. When the object of analysis is a painting produced by a Renaissance or Early Modern artist's workshop, the number of artistic practice

¹Department of Anthropology, Purdue University, West Lafayette, IN 47907, USA. ²Department of Physics, Case Western Reserve University, Cleveland, OH 44106, USA. ³Department of Art History and Art, Case Western Reserve University, Cleveland, OH 44106, USA. ⁴The Frick Pittsburgh Museums and Gardens, Pittsburgh, PA 15208, USA. ⁵Department of Computer and Data Sciences, Case Western Reserve University, Cleveland, OH 44106, USA. ⁶MORE Center, Case Western Reserve University, Cleveland, OH 44106, USA. ⁷Painting Department, Cleveland Institute of Art, Cleveland, OH 44106, USA. ⁸Factum Foundation, Madrid 28037, Spain.

*Corresponding author. Email: axv487@case.edu

regimes is generally unknown, known examples of the included artists' work may not exist, and knowledge of the painting practices used is incomplete. Unsupervised learning, wherein the machine analyzes individual objects (e.g., images or patches of an image) and creates classes based on the statistical properties of those objects, would appear more appropriate. However, a naive unsupervised analysis still depends on hyperparameters that can be tuned to yield larger or smaller numbers of classes. Whether performed by the trained art historian or by a neural network, the problem of identifying regions of shared artistic practice is, in a word, nontrivial.

Here, we demonstrate a novel technique, to our knowledge, for ML-based image attribution and use it to assign paintings to their respective artist(s). We then combine this method with network analysis to identify artistic practice regimes (combinations of artists and materials) and create a measure of the heterogeneity of artistic practice (HAP) within a given painting or set of paintings. We apply this method to two important historical paintings by El Greco, one considered to be entirely by the master himself, and one previously thought to feature the work of members of his workshop.

The method described here, pairwise assignment training for classifying heterogeneity (PATCH), uses supervised learning toward an unsupervised end. Rather than train a network to recognize individual classes, we test whether the network is capable of learning to distinguish between two objects and then use that information to build post hoc classes. The PATCH method relies on the inability of the network to correctly sort patches of paintings or regions of a

painting painted by the same artist. Given patches of two paintings by different painters—a different-artist pair—the network will learn to assign those patches to the correct painting with a high degree of accuracy, as demonstrated in our previous publication. However, if the network is given patches of two paintings by the same artist—a same-artist pair—it will not be able to learn to assign those patches to the correct painting; it will perform no better than a coin flip. We then create a network of same-artist pairs and use network analysis to construct classes corresponding to artistic practice regimes. In this way, we obviate the need for separate ground truth information for individual artists and artistic practice regimes, enabling successful applications in contexts such as workshop paintings where the ground truth is not known with certainty.

RESULTS

Development and validation of the PATCH algorithm

The PATCH method consists of two phases. The first phase, pairwise assignment training, identifies pairs of paintings or pairs of regions of a single painting that are the “same,” that is, created by the same artist under the same conditions (Fig. 1). The second phase, community finding, creates post hoc classes based on these pairings. The PATCH method was developed and then validated on a set of paintings by nine known student artists. Each artist individually painted three paintings, all using the same tools (a set of three brushes of different sizes and a palette knife), materials (paints and canvas paper),

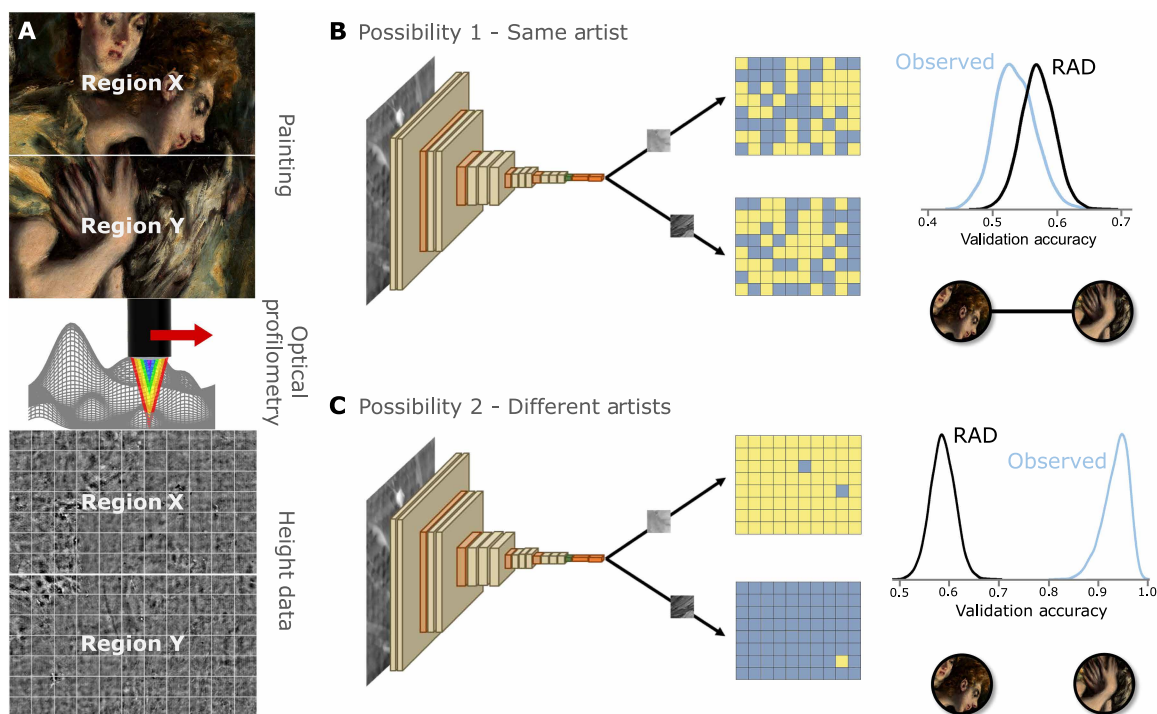


Fig. 1. Overview of the first phase of the PATCH method. (A) Paint height data are gathered using optical profilometry, and two regions are chosen for comparison. Each region corresponds to an entire painting in the case of the student experiment, or a portion of a historical workshop painting, for example, the ones from El Greco's *The Baptism* shown here. The height values are encoded in a 16-bit image file, which is subsequently divided into 1 cm-by-1 cm patches. (B and C) A CNN is repeatedly trained to sort patches of two images, recording its accuracy as a percentage of patches correctly sorted from the test set. This distribution of observed accuracies is compared to the random assignment distribution (RAD) of accuracies if the patches were sorted at random. If the observed accuracy is no better than the RAD accuracy, then the images were likely painted by the same artist, whereas if the observed accuracy is better than expected by chance, the images were likely painted by different artists. Image credit: painting images courtesy of Wikimedia Commons, public domain.

and subject (a photograph of lilies), as described previously (16). While the artists were instructed to realistically depict the subject, choices of style and technique were left up to the artists. One artist was only able to produce two paintings, and one painting was removed from the sample because of the optical profilometry data being corrupted, for a total of 25 paintings. Topographical information for each painting was recorded using high-resolution spectral confocal optical profilometry with a spatial resolution of 50 μm , and a height repeatability of 200 nm. The resulting data were processed to remove any large-scale warping of the canvas by subtracting a mean-filtered version of the height map (with filter radius of 0.5 cm) (16). Because in this dataset, the paintings were small (12 cm by 15 cm), each “region” chosen for the pairwise training consists of an entire painting. The region is then subdivided into 1-cm² square patches, which corresponds to the size where the earlier supervised ML analysis (16) was able to achieve maximum classification accuracy. This patch size is big enough to contain ample information about brushstrokes, while yielding sufficiently large datasets for training (in this case, 180 patches per painting).

Details of the network architecture are described in Materials and Methods. To measure the ability of the network to learn to distinguish between two paintings or regions of a painting, we ran 26 training folds. Each fold comprised the following steps: One-square centimeter patches from each painting/region were randomly selected, with replacement. Sampling with replacement creates a bootstrapping effect that drives down the success rate of same-artist pairs and drives up the success rate for different-artist pairs. To eliminate the influence of directional elements, the corners of each patch were trimmed to create an octagon, and the patch was randomly rotated to one of eight possible orientations. The network was then trained for 25 epochs with validation set size set to 30% of patches, and the maximum validation accuracy (percentage of validation patches correctly identified) was recorded. The distribution of the 26 maximum validation accuracies was then compared to the distribution expected if the network were assigning patches at random.

If the network has failed to learn to assign objects to their correct classes, it should perform this task no better than if it were assigning them randomly. We are measuring the maximum validation accuracy from 25 training epochs in each of 26 folds. Thus, the distribution expected if the network fails to learn to distinguish between painters would be the distribution $p_{n,k}^{\max}(m)$, the probability that you will see a maximum of m heads (i.e., correct assignments) over k repetitions (k epochs) of an experiment where you flip a fair coin n times (the number of patches in the test set), which we call the random assignment distribution (RAD) (see Supplementary Text section S1 for an analytical derivation of this distribution).

Two decision criteria were selected to determine whether the distribution of maximum validation accuracies indicated that the network had failed to learn to distinguish between the input paintings: (i) if the observed mean had a z score less than 2 relative to the RAD mean and (ii) if the largest of the observed maximum validation accuracies (the right edge of the observed distribution) among the 26 folds was less than an empirically determined threshold based on finite sampling from RAD, plus 10% to account for the effects of bootstrapping (selection with replacement). Figure S3 shows a graphical summary of the decision process, including a flowchart. Using these decision criteria, the classification performance in determining whether the artist was the same or different was exceptional. We identified same-artist pairs ($n = 23$) and different-artist pairs ($n = 277$) with F_1

scores of 0.889 and 0.991, respectively (see table S2). The PATCH algorithm substantially outperformed a statistical method based on surface roughness and beat a variety of unsupervised and supervised clustering approaches (see details in Supplementary Text section S3) in F_1 for both same-artist and different-artist pairs. In addition, PATCH with height data outperforms PATCH on high-resolution photographs (see Supplementary Text section S9).

While pairwise assignment training performs exceptionally well in identifying artists with control for materials, situations where ground truth about the artistic practices used is not available require some means to create the post hoc classes to which regions of a painting will be assigned. To this end, we use network analysis. We construct a network where same-artist pairs (or, for historical paintings, same-practice pairs) of paintings or regions are connected by an edge and different-artist/practice pairs are unconnected.

With potentially spurious edges removed (see Supplementary Text section S7) and the remaining edges weighted equally, we then use a community-finding algorithm to identify groups of regions that are the most similar. If a set of regions has a large number of internal links within the set and fewer links to outside regions, it forms a “community” within the network (17–19), suggesting that the regions in a given community were likely painted under different circumstances (artists, materials, etc.) than those in other communities. To establish the degree of difference between communities, we can characterize the community structure of the network via measures such as modularity (Q)—the fraction of edges in the network that are internal to communities minus the mean fraction in a network with the same communities but where the edges are completely randomized (18). We implemented the Louvain community-finding algorithm (20) in Gephi (21). With the resolution set to 1.0, the algorithm returns the partition of the network with the maximum modularity. The Louvain algorithm returned a correct partition of the experimental dataset (nine disjoint communities for nine different artists) with $Q = 0.875$ (see fig. S5).

El Greco, *The Baptism*, and *Christ on the Cross*

El Greco, born Domenikos Theotokópoulos (1541 to 1614), is regarded as a pillar of the Renaissance in Spain and as an early progenitor of modernism (22). El Greco became a master of icon painting in Crete before journeying through Venice and Rome, eventually settling in Toledo, Spain (23, 24). It was there that the master married Byzantine and Venetian motifs (25) and developed the peculiar style for which he is most widely recognized, with its distorted human figures, “expressive hands” (26), and “exploitation” of “pure colors to their limits” (27). Two examples of that style are *The Baptism of Christ* (1624, Hospital Tavera, Toledo, Esp., henceforth, *The Baptism*), shown in Fig. 2A, and *Christ on the Cross with Landscape* (~1600 to 1610, Cleveland Museum of Art, Cleveland, OH, henceforth *Christ on the Cross*), shown in Fig. 2B.

These works were chosen for this analysis because of the difference in how art historians have characterized their authorship. *The Baptism* has long been thought to feature the work of El Greco and at least one other artist. Primary historical evidence indicates that El Greco began the painting under a contract with the Hospital Tavera but retained it at his death in 1614. It was delivered to the Hospital nearly a decade later (28), during which period art historians have proposed that the painting was finished by workshop members, particularly the master’s son, Jorge Manuel (23). Previous art historical studies have attempted to attribute regions of *The Baptism* to the master




Fig. 2. The two historical paintings analyzed using the PATCH algorithm. (A) El Greco's *The Baptism of Christ*, begun in approximately 1608 and suspected to have been finished by his workshop before delivery to the Hospital Tavera in 1624. (B) El Greco's *Christ on the Cross with Landscape*, ~1600 to 1610. Image credit: *The Baptism* and *Christ on the Cross* images courtesy of Wikimedia Commons and the Cleveland Museum of Art, respectively, both public domain.

and others (illustrated in fig. S4) largely through connoisseurship—visual analysis of artists' styles and artistic choices (29). El Greco himself is proposed to have painted the entire top, with the possible exception of the robe of the angel on the right (which was not scanned for this study because of its poor condition) (28, 30), as well as the angel in green (bottom left), with the exception of the wings. Jorge Manuel is proposed to have painted John the Baptist (bottom right) and the adjacent figure in red. Lopera (30) and Wethey (28) ascribe the image of Christ (bottom center) to Jorge Manuel and El Greco, respectively, suggesting that it could represent the work of both artists. Lopera believes that a third hand may have been involved in rendering the faces of the angels in the background on the bottom. The landscapes at the bottom were extended and “transformed into a river” by an unknown artist sometime between receipt of the painting and its installation in the epistle-side altarpiece ~1660 (30), which was not an unusual practice at the time. This likely required inpainting into the original canvas. The river was subsequently removed sometime after 1908, and the landscape at the very bottom right (an area we did not analyze because of its condition) was crudely repaired in 1936 (28, 30, 31). In contrast, *Christ on the Cross* has been entirely attributed to El Greco himself. It is generally accepted that El Greco created many works in their entirety without the aid of his workshop, including *Christ on the Cross*. Francis (32) states that most of El Greco's approximately 20 variations on the theme of the Crucifixion were “by his own hand,” and notes the “lightning dexterity of brushwork” in this particular example.

Application to El Greco's *The Baptism*

Two large sections of *The Baptism* were scanned in situ using the Factum Foundation's Lucida scanner. The Lucida scanner is a non-contact laser triangulation scanner with 100- μm lateral resolution

and is capable of acquiring scans in situ over a large area by stitching together 48-cm square tiles. The scanned regions used in this study are indicated in Fig. 3. The height data were processed in the same way as for the student paintings described above, including correction for possible canvas warping.

Regions of the topographic image of *The Baptism* were selected by hand to minimize variation in subject matter (e.g., faces and robes) within each region and to constrain region size to between 180 and 540 cm^2 . Hand selection allows art historical knowledge to inform region selection. For example, using a grid to select regions could divide areas such as the face of Christ which are known to have often been painted by masters themselves. The lower bound was selected because the method had been developed and validated using 180-patch student paintings, 180 patches being 6 patches more than the theoretical lower bound determined by sample size. The upper bound was selected by balancing computational resources with our ability to ensure adequate representation of patches when regions of significantly different size were compared (see Supplementary Text section S8 for more details on the bounds). We avoided selecting areas with significant cracking or damage, including the robes of the large angels to the left and right of God, and the red swatch held by the “green angel” in the bottom half of the painting. When comparing regions of different sizes during the PATCH analysis, the sample size was set equal to the number of patches in the smaller region. To test the robustness of our conclusions to the details of region selection, we also performed PATCH on both El Greco paintings using a different set of regions (Supplementary Text section S10), which led to qualitatively similar conclusions to the ones described below.

Pairwise comparisons of all 55 regions (a total of 1485 pairwise tests) of *The Baptism* yielded a network with 356 edges, which

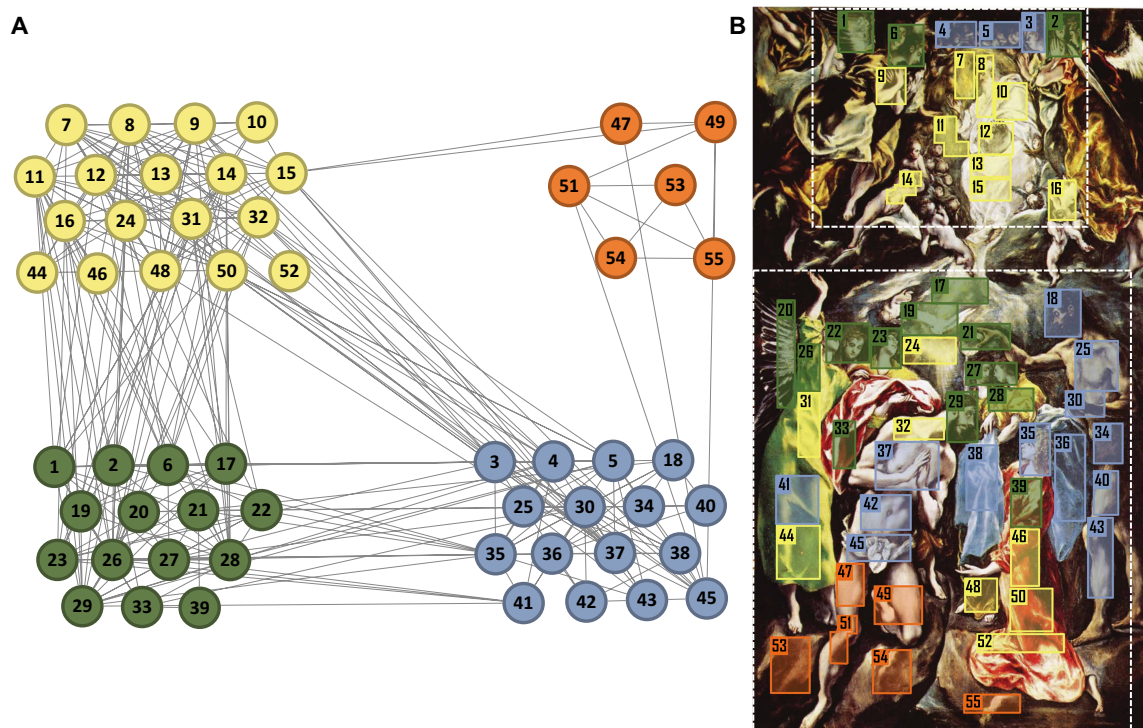


Fig. 3. Results of PATCH analysis on El Greco's *The Baptism of Christ*. (A) Network diagram showing the maximum modularity (Q) partition. Four communities are present. (B) The communities mapped onto the corresponding regions of the painting that were analyzed. Scanned areas of the painting are outlined by dashed white lines. Image credit: Painting image is courtesy of Wikimedia Commons, public domain.

was reduced to 314 edges after pruning. The Louvain community-finding algorithm (20), with resolution set to 1.0, returned a partition with $Q = 0.341$ that featured four communities (Fig. 3A). A value of $Q \geq 0.3$ is often considered evidence of structure in real-world networks (18).

The four communities are mapped onto the painting in Fig. 3B. The first community overlaps much of the top portion of the painting, including the image of God the Father. The second overlaps the image of John the Baptist and Christ's torso and hands. The third covers much of the center of the painting, the face of Christ and the faces of many of the background angels. The final community is confined to Jesus' legs and the rocks at the bottom of the painting.

Application to El Greco's *Christ on the Cross*

As in the case of *The Baptism*, regions of the topographic image of *Christ on the Cross* were selected by hand to minimize variation in subject matter and to constrain region size. Areas with apparent cracking or damage, such as the background and trees on the lower left and the cross beam to the immediate left of Christ's face, were not included.

Pairwise comparison of all 24 regions (276 pairwise tests) yielded a network with 123 edges, which was reduced to 108 edges after pruning. The Louvain community-finding algorithm (20) at resolution 1.0 generated a partition with $Q = 0.231$ that featured two communities. Two is the smallest possible number of communities for a non-fully connected graph, as Q compares inter- and intracommunity connections. Communities are mapped onto the painting in Fig. 4. Constituent regions of both communities are not randomly distributed, with one community overlapping much of Christ's torso and face, while the

other community overlaps Christ's legs, the lower portion of the cross and regions of the background.

DISCUSSION

To our knowledge, the PATCH learning method has not previously been demonstrated or applied both to paintings from a controlled experiment and to historical paintings by a well-known artist with a workshop. Among AI/ML-based methods designed to aid in the analysis of historical paintings, PATCH is notable in that it requires no external, ground truth data. Networks trained on known examples can accomplish tasks associated with workshop practice and authorship with exceptional accuracy. For example, Ugail and colleagues (33) recently used transfer learning to authenticate paintings by the Renaissance master Raphael (Raffello Sanzio). By combining edge detection with typical feature extraction in a residual neural network, they were able to recognize works by Raphael with 98% accuracy. However, known examples of artists' work may not be available, as exemplified by the workshop context. Our method shows exceptional promise as a tool for ML analysis of complex historical works where known samples of work by the artists involved do not exist. It may have extensive application in other image analysis tasks where there is little to no ground truth information available.

Further, the second phase of PATCH creates what we consider a measure of the HAP within a given dataset. Communities found within a dataset represent different artists, materials, or both. The degree to which those communities are interconnected provides information about the possibility of shared authorship, materials, or

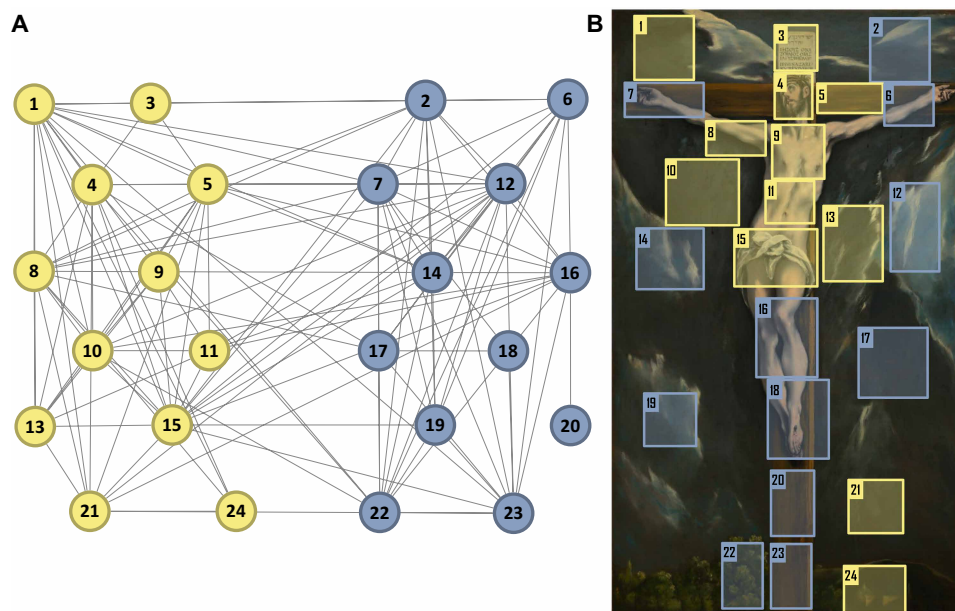


Fig. 4. Results of PATCH analysis on El Greco's *Christ on the Cross with Landscape*. (A) Network diagram showing the maximum modularity (Q) partition. Two communities are identified, although this is the minimum number of communities discoverable using a maximum-modularity partition. (B) The communities mapped onto the corresponding regions of the painting that were analyzed. Image credit: Painting image is courtesy of the Cleveland Museum of Art, public domain.

other practices. A fully connected network with no distinct communities has $Q = 0$. As the number of communities increases and the interconnection between communities decreases, modularity increases toward 1 (a completely disconnected network has undefined modularity but would effectively represent the maximum heterogeneity). The networks from our three studies provide empirical examples of lower and higher modularity networks. *Christ on the Cross* and *The Baptism* have Q values that cluster around 0.3, considered the threshold for evidence of structure, indicating low heterogeneity. The network for the experimental student paintings has much greater heterogeneity at $Q = 0.875$ (see fig. S5). The maximum modularity depends on the number of edges in the network (34) but also on the number of communities in the network. A network with two fully connected and completely disjoint communities has $Q = 0.5$, for example. Hence, network modularity can function as a measure of absolute observed heterogeneity (Fig. 5).

HAP informs interpretation of our analyses of *Christ on the Cross* and *The Baptism*. For the former, our analysis identifies two communities of regions, but with Q below the threshold generally accepted to indicate structure (0.231 versus 0.3). The regions in each community do cluster in space, however, which is worth considering. One community overlaps Christ's head and torso and surrounding areas. Conservation files generously provided by the Cleveland Museum of Art indicate that the figure of Christ (areas overlapped by the community in yellow in Fig. 4) contains inpainting from previous conservation, which could create a distinguishable signal, although damage to the canvas is another, more likely culprit. Further mapping of damage and past conservation efforts will improve our understanding of this work. However, the results comport with the notion of *Christ on the Cross* as the work of a single artist with some variation introduced by artistic practice or possibly early conservation practice.




Fig. 5. Heterogeneity scale. Networks and modularity (Q) values for *Christ on the Cross*, *The Baptism*, and the experimental student paintings along the axis of Q from least to most heterogeneous. Image credit: Painting image fragments are courtesy of Wikimedia Commons and the Cleveland Museum of Art, all public domain.

At first glance, the communities identified in our analysis of *The Baptism* appear to roughly correspond to the proposed attributions by previous authors discussed above. The community in yellow in Fig. 3 overlaps much of the top of the painting, previously attributed to El Greco, while the community in blue overlaps John the Baptist and the figure in red, assigned to Jorge Manuel. We also find a combination of communities in the face and torso of Christ, which was attributed to both artists. A single, small, and particularly unique community (in orange in Fig. 3) overlaps the landscape at the bottom of the painting and Christ's legs. The bottom of the painting is known to have been altered by another artist long after delivery and part of that alteration later removed (see Supplementary Text section S12 for historical details). That PATCH so clearly indicates the difference between these regions and the rest of the painting illustrates the method's utility.

Overall, however, the PATCH findings suggest a different story underlying the existing art historical interpretation. The partition of *The Baptism* has $Q = 0.341$, which is only slightly above the threshold for evidence of structure. Nearly a third of all edges in the network are interconnections between three of the four communities

we have identified. The yellow and green groups have 49 intercommunity edges, blue and green 30, and blue and yellow 22. If the blue, green, and yellow communities represented three distinct artists working with the same or similar materials, we would expect the communities to be more insular and the modularity of the network to be substantially higher, similar to what we observed in our controlled student painting dataset. Yet, the modularity of *The Baptism* is much closer to that of *Christ on the Cross* ($Q = 0.231$), which appears to feature one artist. As shown in the analysis of student paintings, the CNN is particularly proficient at recognizing different-artist pairs and is more likely to incorrectly label a same-artist pair as different than the converse. Thus, the intercommunity edges in *The Baptism* network that survived our initial pruning are not liable to be spurious.

This suggests that some unifying factor connects these three communities. One possibility is a single artist working with different brushes. Renaissance artists used brushes made from the bristles or hair of several different mammal species, variation in the characteristics (e.g., coarseness) of which could affect the deposition of paint (35–37) (see Supplementary Text section S5 for more information). It is also possible that more than one artist could be working with the same set of brushes. However, in the controlled experiment, the artists were each supplied with the same set of three brushes of different sizes, suggesting that variation in artist would overwhelm the unifying signal from brushes. The evidence could also indicate the work of one artist with changing style or technique across subjects or over time. It is not inconceivable that El Greco, a master icon painter, might change technique in rendering the “invisible” (i.e., God and the angels) and the “visible” (i.e., Christ and John the Baptist). Biannuci and colleagues (38) present evidence that El Greco suffered a series of ischemic events, one in the 1590s and a second in 1608 (the year the contract for the altarpieces for the Hospital Tavera was initiated), “resulting in progressive disabilities with fluctuating course characterized by temporary improvements and worsening before his death.” The regions that were previously identified as the work of different individual painters could potentially represent variation in the master’s individual style (however this may have arisen) over the course of his final years. Investigation of changes in style and technique and their effect on PATCH analysis will be an important direction for future research.

PATCH has the capability to make a substantial contribution to research as a complement to existing art historical methods. The accuracy of the method has been demonstrated in the validation on student paintings. Our analyses of *Christ on the Cross* and *The Baptism* show the true potential of PATCH by contributing important new information to art historical scholarship on the creation of objects in El Greco’s workshop in Toledo. In the analysis of *The Baptism*, the fact that the orange group, which overlaps a region known to have been altered after the delivery of the painting, is so disconnected from the remainder of the communities (see Supplementary Text section S4 for further analyses) speaks to the method’s capability of recognizing a unique contribution or substantial departure from the technique used in the rest of the painting. That the other three communities roughly correspond to art historical attributions suggests that PATCH has a kind of semantic fidelity (39)—our method recognizes areas that a trained art historian would classify as different. However, the connections between those areas revealed by PATCH also call those attributions into question. Computer vision allows us to view the surface of a painting at a different scale. Previous analyses

indicate that features as small as the diameter of a single paint brush bristle (length scales of 0.2 to 0.4 mm) may be integral to the network’s identification of an artist (16). By analyzing spatial correlations at this microscopic scale, the network may be revealing aspects of the physics of paint application and of the physiology of hand movements. PATCH adds microscale features to the macro (brushstrokes) and metascale (historical and material) data, allowing for a full-scale analysis and the discovery of heretofore unseen evidence of the processes by which paintings were created.

Future research with PATCH will focus on teasing apart the influences of materials and artists on the HAP. Expansion of PATCH applications to include interpainting comparisons and other art historical applications will be important as well. However, uses of PATCH are not limited to those contexts: It provides a general-purpose approach to assessing regional similarity in images or data that can be encoded in image-like arrays.

Hence, PATCH is a promising method for applications where supervised learning is impossible (because the ground truth is unknown) and alternative unsupervised approaches are ineffective, especially for a posteriori cluster number. Potential applications span the physical and social sciences and include image segmentation tasks in applications such as medical imaging (40), agricultural remote sensing (41), urban development and design (42), and microstructural analysis, as well as anomaly detection in a range of manufacturing contexts (43, 44), among others.

MATERIALS AND METHODS

The neural network architecture chosen for the PATCH algorithm was VGG-16 (45) implemented in TensorFlow and pretrained on the ImageNet dataset (46). Given the computational complexity of the training—the combinatorics of many region-to-region comparisons, repeated over many folds, described in the Results—the relatively light-weight but accurate VGG-16 architecture was ideal for the PATCH approach. However, the idea of the PATCH algorithm is quite general, and other network architectures could be used in future applications. The sections in the Supplementary Text provide additional details of methodology, which we here summarize: Section S1 gives an analytical derivation of the RAD and describes the mean and maximum of the distribution that are used as criteria for the same versus different test in the PATCH algorithm. Section S2 discusses how the neural network model hyperparameters (i.e., learning rate, number of dense layers, etc.) were optimized to yield the best performance on the student painting dataset. Section S3 compares PATCH to alternative methods of predicting same- and different-artist pairs, ranging from simple approaches (like categorizing regions by surface roughness) to more sophisticated unsupervised and supervised classification techniques. Section S4 analyzes in more depth the networks produced by the PATCH algorithm, focusing in particular on the degree distributions within and between the identified communities. Section S5 discusses material aspects of brush construction during the period of El Greco’s paintings, and Section S6 shows a number of additional figures (including a graphical summary of the algorithm decision criteria, previous art historical attributions for different regions in *The Baptism*, and the PATCH-produced network diagram for the student painting dataset). Section S7 describes the method by which a small number of spurious edges were identified and pruned from the networks. Section S8 goes over the rationale for the patch numbers used in the analysis and

shows how the number of folds in the algorithm depends on the patch numbers in each pair of compared regions. Section S9 summarizes the results for the PATCH algorithm applied to photo data (rather than height data) for the student painting dataset, and section S10 illustrates how different region selection choices for the El Greco paintings yielded the same qualitative results as those discussed in the main text. Section S11 covers some further details of the limitations and challenges in applying the algorithm to historical paintings, and finally section S12 gives a brief overview of the known historical facts relating to provenance and conservation of the two El Greco paintings analyzed in this study.

Supplementary Materials

This PDF file includes:

Supplementary Text
Tables S1 to S3
Figs. S1 to S10
References

REFERENCES

1. S. Alpers, *The Vexations of Art: Velázquez and Others* (Yale Univ. Press, 2005).
2. C. Brown, *Van Dyck* (Cornell Univ. Press, 1983).
3. B. Talvacchia, *Raphael* (Phaidon Press, 2007).
4. C. Bambach, *Drawing and Painting in the Italian Renaissance Workshop: Theory and Practice, 1300–1600* (Cambridge Univ. Press, 1999).
5. J. Brooks, D. Allen, X. F. Salomon, *Andrea del Sarto: The Renaissance Workshop in Action* (J. Paul Getty Museum, 2015).
6. S. van Hoogstraten, *Inleyding tot de hooge schoole der schilderkonst: anders de zichtbaere werelt* (Tot Rotterdam, 1678). [Introduction to the Academy of Painting, or the Visible World].
7. E. Van de Wetering, The search for the master's hand: An anachronism? (A summary). *Künstlerische Austausch/Artistic Exchange: Akten des XXVIII, Internationalen Kongresses für Kunstgeschichte* (1993), pp. 627–630.
8. J. M. Hughes, D. J. Graham, D. N. Rockmore, Quantification of artistic style through sparse coding analysis in the drawings of Pieter Bruegel the Elder. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 1279–1283 (2010).
9. T. Friedman, D. J. Lurie, A. Shalom, Authentication of Rembrandt's self-portraits through the use of facial aging analysis. *Isr. Med. Assoc. J.* **14**, 591–594 (2012).
10. S. Kogou, G. Shahtahmassebi, A. Lucian, H. Liang, B. Shui, W. Zhang, B. Su, S. van Schaik, From remote sensing and machine learning to the history of the Silk Road: Large scale material identification on wall paintings. *Sci. Rep.* **10**, 19312 (2020).
11. Y. Assael, T. Sommerschild, B. Shillingford, M. Bordbar, J. Pavlopoulos, M. Chatzipanagiotou, I. Androustopoulos, J. Prag, N. De Freitas, Restoring and attributing ancient texts using deep neural networks. *Nature* **603**, 280–283 (2022).
12. T. Winterbottom, A. Leone, N. Al Moubayed, A deep learning approach to fight illicit trafficking of antiquities using artefact instance classification. *Sci. Rep.* **12**, 13468 (2022).
13. C. Sandoval, E. Pirogova, M. Lech, Two-stage deep learning approach to the classification of fine-art paintings. *IEEE Access* **7**, 41770–41781 (2019).
14. M. Bigerelle, R. Guibert, A. Mironova, F. Robache, R. Deltonbe, L. Nys, C. A. Brown, Fractal and statistical characterization of brushstroke on paintings. *Surf. Topogr. Metrol. Prop.* **11**, 015019 (2023).
15. A. Elgammal, B. Liu, D. Kim, M. Elhoseiny, M. Mazzone, The shape of art history in the eyes of the machine, in *Proceedings of the AAAI Conference on Artificial Intelligence* (2018), vol. 32:1.
16. F. Ji, M. S. McMaster, S. Schwab, G. Singh, L. N. Smith, S. Adhikari, M. O'Dwyer, F. Sayed, A. Ingrisano, D. Yoder, E. S. Bolman, I. T. Martin, M. Hinczewski, K. D. Singer, Discerning the painter's hand: Machine learning on surface topography. *Herit. Sci.* **9**, 152 (2021).
17. M. Girvan, M. E. Newman, Community structure in social and biological networks. *Proc. Natl. Acad. Sci.* **99**, 7821–7826 (2002).
18. M. E. Newman, M. Girvan, Finding and evaluating community structure in networks. *Phys. Rev. E* **69**, 026113 (2004).
19. S. Fortunato, D. Hric, Community detection in networks: A user guide. *Phys. Rep.* **659**, 1–44 (2016).
20. V. D. Blondel, J.-L. Guillaume, R. Lambiotte, E. Lefebvre, Fast unfolding of communities in large networks. *J. Stat. Mech.* **2008**, P10008 (2008).
21. M. Bastian, S. Heymann, M. Jacomy, "Gephi: An open source software for exploring and manipulating networks," in *Proceedings of the International AAAI Conference on Web and Social Media*, (2009), vol. 3:1, pp. 361–362.
22. N. Hadjinicolaou, "He is, indeed, A prophet of the moderns," in *El Greco of Toledo: Painter of the Visible and the Invisible*, F. Marias, Ed. (Ediciones El Viso, 2014).
23. F. Marias, *El Greco: Life and Work—A New History* (Thames & Hudson, 2013).
24. F. Marias, "The Greek, between invention and history," in *El Greco of Toledo: Painter of the Visible and the Invisible*, F. Marias, Ed. (Ediciones El Viso, 2014).
25. R. G. Mann, Delno C. West award winner: Tradition and originality in El Greco's work: His synthesis of byzantine and renaissance conceptions of art. *Quidditas* **23**, 7 (2002).
26. F. J. Sanchez Canton, *Spanish Drawings From the 10th to the 19th Century* (Shorewood Publishers, 1964).
27. M. L. Krumrine, Color in El Greco. *Arte Lombarda* (Vita e Pensiero - Pubblicazioni dell'Università Cattolica del Sacro Cuore, 1993), pp. 42–48.
28. H. E. Wethey, *El Greco and His School* (Princeton Univ. Press, 1962).
29. C. B. Scallan, *Rembrandt, Reputation, and the Practice of Connoisseurship* (Amsterdam Univ. Press, 2004).
30. J. A. Lopera, *El Greco: Estudio y Catalogo* (Fundacion De Apoyo a La Historia Del Arte Hispanico, 2005) [El Greco: Studio and Catalog].
31. M. B. Cossio, *El Greco* (Madrid, V. Suárez, 1908).
32. H. Francis, A crucifixion by El Greco. *Bull. Cleveland Museum Art* **40**, (1953); <https://www.jstor.org/stable/25141846>.
33. H. Ugail, D. G. Stork, H. Edwards, S. C. Seward, C. Brooke, Deep transfer learning for visual analysis and attribution of paintings by Raphael. *Herit. Sci.* **11**, 268 (2023).
34. S. Fortunato, M. Barthelemy, Resolution limit in community detection. *Proc. Natl. Acad. Sci. U.S.A.* **104**, 36–41 (2007).
35. R. J. Gettens, G. L. Stout, *Painting Materials: A Short Encyclopedia* (Courier Corporation, 2012).
36. A. Palomino-y Velasco, The pictorial museum and optical scale, in *Artists' Techniques in Golden Age Spain: Six Treatises in Translation*, Z. Veliz, Ed. (Cambridge Univ. Press, 1986), pp. 141–191.
37. C. Cennini, C. J. Herringham, *The Book of the Art of Cennino Cennini: A Contemporary Practical Treatise on Quattrocento Painting* (Routledge, 2018).
38. R. Bianucci, F. M. Franco, O. Appenzeller, Historical evidence supports El Greco's depiction of a neurological condition in his attributed self-portrait. *J. Neurol. Sci.* **372**, 316–317 (2017).
39. Y. Takagi, S. Nishimoto, "High-resolution image reconstruction with latent diffusion models from human brain activity," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 14453–14463.
40. K. Raza, N. K. Singh, A tour of unsupervised deep learning for medical image analysis. *Curr. Med. Imaging Rev.* **17**, 1059–1077 (2021).
41. S. Khanal, K. Kc, J. P. Fulton, S. Shearer, E. Ozkan, Remote sensing in agriculture—Accomplishments, limitations, and opportunities. *Remote Sens.* **12**, 3783 (2020).
42. J. Francis, J. Bright, S. Esnaashari, Y. Hashem, D. Morgan, V. J. Straub, Unsupervised feature extraction of aerial images for clustering and understanding hazardous road segments. *Sci. Rep.* **13**, 10922 (2023).
43. A. Kharitonov, A. Nahhas, M. Pohl, K. Turowski, Comparative analysis of machine learning models for anomaly detection in manufacturing. *Procedia Comput. Sci.* **200**, 1288–1297 (2022).
44. J. Liu, G. Xie, J. Wang, S. Li, C. Wang, F. Zheng, Y. Jin, Deep industrial image anomaly detection: A survey. *Mach. Intell. Res.* **21**, 104–135 (2024).
45. K. Simonyan, A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *3rd International Conference on Learning Representations (ICLR 2015)* (2015), pp. 1–14.
46. J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, ImageNet: A large-scale hierarchical image database, in *2009 IEEE Conference on Computer Vision and Pattern Recognition* (2009), pp. 248–255.
47. J. Bergstra, B. Komer, C. Eliasmith, D. Yamins, D. D. Cox, Hyperopt: A Python library for model selection and hyperparameter optimization. *Comput. Sci. Discov.* **8**, 014008 (2015).
48. L. McInnes, J. Healy, N. Saul, L. Großberger, UMAP: Uniform manifold approximation and projection. *J. Open Source Softw.* **3**, 861 (2018).
49. L. McInnes, J. Healy, S. Astels, hdbscan: Hierarchical density based clustering. *J. Open Source Softw.* **2**, 205 (2017).
50. A. Fuga, *Artists' Techniques and Materials* (Getty Publications, 2006).
51. T. M. P. Hardy, J. I. Hardy, *Animal Fibers Used in Brushes*, Tech. Rep. "802," United States Department of Agriculture, (Washington, DC, 1949).
52. J. J. del Agila, *El TOP: La Represión de la Libertad* (Planeta, 2001) [The TOP: the Repression of Freedom].

Acknowledgments: We would like to thank Fundación Casa Ducal de Medinaceli for providing data for *The Baptism* and J. M. A. Solis for feedback on our initial draft, which helped us reevaluate our conclusions and improve the manuscript. We would also like to thank the Cleveland Museum of Art for use of the data from *Christ on the Cross*, and the Lapis Senior Conservator of Paintings and Head of Paintings Conservation, D. Yoder, for help and feedback. Computational support was provided by the CWRU High Performance

Computing cluster and a workstation donated through the Dell Seed Program. **Funding:** The study was supported by the Experimental Humanities initiative of the Case Western Reserve University (CWRU) College of Arts and Sciences and the Jack, Joseph and Morton Mandel Foundation through grants SPC603071 and SPC603976. **Author contributions:** A.V.H. conceived the PATCH method. L.S. and M.S.M. conceived the initial project. A.V.H., M.H., K.S., and M.S.M. developed the method. A.I. oversaw creation of paintings. A. Lo., C.B., and I.T.M. performed scans. A.V.H. and M.M. wrote the PATCH code. A.V.H., A. Li., M.M., and M.H. performed experiments and analyses. A.V.H., M.H., K.S., E.S.B., C.G.P., A. Li., and L.S. wrote the manuscript. **Competing interests:** The authors declare that they have no competing interests. **Data, code, and materials availability:** The dataset used to develop and validate the PATCH method is available on OSF at osf.io/a5xnh/. Height data from *The Baptism* and *Christ on the Cross* are the property of the institutions that own those paintings: the

Fundación Casa Ducal de Medinaceli and the Cleveland Museum of Art, respectively. The code used to run the PATCH algorithm is available on GitHub: <https://github.com/hincz-lab/PATCH>. Archived, permanent versions of the data and code are available at Zenodo at the links <https://doi.org/10.5281/zenodo.17861983> (data) and <https://doi.org/10.5281/zenodo.17861942> (code). All data and code needed to evaluate and reproduce the results in the paper are present in the paper and/or the Supplementary Materials. This study did not generate new materials.

Submitted 24 June 2025

Accepted 18 March 2026

Published 17 April 2026

10.1126/sciadv.aea0485

PATCH: A deep learning method to assess heterogeneity of artistic practice in historical paintings

Andrew Van Horn, Lauryn Smith, Mahamad Salah Mahmoud, Michael McMaster, Clara G. Pinchbeck, Ina T. Martin, Andrew Lininger, Anthony Ingrisano, Adam Lowe, Carlos Bayod, Elizabeth S. Bolman, Kenneth Singer, and Michael Hinczewski

Sci. Adv. **12** (16), eaea0485. DOI: 10.1126/sciadv.aea0485

View the article online

<https://www.science.org/doi/10.1126/sciadv.aea0485>

Permissions

<https://www.science.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of service](#)

Science Advances (ISSN 2375-2548) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science Advances* is a registered trademark of AAAS.

Copyright © 2026 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. Distributed under a Creative Commons Attribution NonCommercial License 4.0 (CC BY-NC).