

Two-by-two ordinal patterns in art paintings

Mateus M. Tarozo ^a, Arthur A. B. Pessa^a, Luciano Zunino ^{b,c}, Osvaldo A. Rosso ^d, Matjaž Perc ^{e,f,g,h,i} and Haroldo V. Ribeiro ^{a,*}

^aDepartamento de Física, Universidade Estadual de Maringá, Maringá, PR 87020-900, Brazil

^bCentro de Investigaciones Ópticas (CONICET La Plata - CIC - UNLP), Gonnet, La Plata 1897, Argentina

^cDepartamento de Ciencias Básicas, Facultad de Ingeniería, Universidad Nacional de La Plata (UNLP), La Plata 1900, Argentina

^dInstituto de Física, Universidade Federal de Alagoas, Maceió 57072-900, Brazil

^eFaculty of Natural Sciences and Mathematics, University of Maribor, Koroška cesta 160, Maribor 2000, Slovenia

^fCommunity Healthcare Center Dr. Adolf Drolc Maribor, Ulica talcev 9, Maribor 2000, Slovenia

^gDepartment of Physics, Kyung Hee University, 26 Kyungheedaero, Dongdaemun-gu, Seoul 02447, Republic of Korea

^hComplexity Science Hub, Metternichgasse 8, Vienna 1030, Austria

ⁱUniversity College, Korea University, 145 Anam-ro, Seongbuk-gu, Seoul 02841, Republic of Korea

*To whom correspondence should be addressed: Email: hvr@dfi.uem.br

Edited By Javier Borge-Holthoefer

Abstract

Quantitative analysis of visual arts has recently expanded to encompass a more extensive array of artworks due to the availability of large-scale digitized art collections. Consistent with formal analyses by art historians, many of these studies highlight the significance of encoding spatial structures within artworks to enhance our understanding of visual arts. However, defining universally applicable, interpretable, and sufficiently simple units that capture the essence of paintings and their artistic styles remains challenging. Here, we examine ordering patterns in pixel intensities within two-by-two partitions of images from nearly 140,000 paintings created over the past 1,000 years. These patterns, categorized into 11 types based on arguments of continuity and symmetry, are both universally applicable and detailed enough to correlate with low-level visual features of paintings. We uncover a universal distribution of these patterns, with consistent prevalence within groups, yet modulated across groups by a nontrivial interplay between pattern smoothness and the likelihood of identical pixel intensities. This finding provides a standardized metric for comparing paintings and styles, further establishing a scale to measure deviations from the average prevalence. Our research also shows that these simple patterns carry valuable information for identifying painting styles, though styles generally exhibit considerable variability in the prevalence of ordinal patterns. Moreover, shifts in the prevalence of these patterns reveal a trend in which artworks increasingly diverge from the average incidence over time; however, this evolution is neither smooth nor uniform, with substantial variability in pattern prevalence, particularly after the 1930s.

Keywords: spatial patterns, complexity, esthetic measure, art history

Significance Statement

Spatial patterns play a crucial role in both esthetic appreciation and in defining the characteristics of artworks and artistic movements. Despite advancements in image processing, encoding these patterns into simple quantitative representations remains challenging. Our study addresses this by analyzing the ordering of pixel intensities in small image partitions, allowing us to represent paintings by a set of 75 ordinal patterns grouped into 11 categories based on continuity and symmetry. These patterns capture key visual features—sharp edges, well-defined or textured elements, and vertical or horizontal structures—providing insights into the artistic style of paintings, highlighting their variability and the complex, nonhomogeneous nature of art evolution.

Introduction

The pioneering work of Birkhoff in 1933 (1), where he formulated an esthetic measure based on the ratio between order and complexity, is often viewed as one of the first contemporary attempts to define universal mathematical principles for evaluating artistic esthetics. His ideas, however, have origins extending back to ancient Greece and medieval philosophers (2). Despite this long historical path, the empirical characterization of artworks using

physical science approaches is relatively recent. The work by Taylor et al. (3) in the late 1990s on the fractal nature of Pollock's drip paintings was pivotal in this regard and spurred a multitude of further studies (4–13), which in turn contributed to the emergence of the field of quantitative study of visual arts (14–16). However, it was not until recent years that these quantitative efforts reached a larger scale and expanded beyond a limited number of artworks from specific artists or artistic styles.

Competing Interest: The authors declare no competing interests.

Received: June 14, 2024. **Accepted:** February 6, 2025

© The Author(s) 2025. Published by Oxford University Press on behalf of National Academy of Sciences. This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial License (<https://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited. For commercial re-use, please contact reprints@oup.com for reprints and translation rights for reprints. All other permissions can be obtained through our RightsLink service via the Permissions link on the article page on our site—for further information please contact journals.permissions@oup.com.

This expansion was primarily driven by the recent availability of extensive digitized art collections, which not only facilitated practical applications but also started to contribute to deepening our understanding of the cultural and social aspects of art and how they possibly evolved over the centuries (17).

Research by Manovich and coauthors (18–20) has pioneered the analysis of large-scale datasets of paintings and other forms of visual art to quantify artistic evolution through metrics such as average brightness and saturation. Changes in the use of color and contrast over time were also investigated by Kim et al. (21) and Lee et al. (22), who identified significant differences across historical periods and artistic styles. Art paintings were further explored using a physics-inspired approach based on permutation entropy and statistical complexity by Sigaki et al. (23), revealing a temporal evolution marked by transitions aligned with major historical periods. This methodology was subsequently applied by Valensise et al. (24) to assess the visual complexity of memes on social media platforms. Additionally, Lee et al. (25) dissected the compositional structure of landscape paintings, finding that from 1600 to 1850, artworks typically featured a primary horizontal partition with a secondary vertical division, a style that gradually fell into disuse, yielding a preference for dual horizontal partitions in 20th-century paintings. More recently, Karjus et al. (26) introduced a representation space termed compression ensembles, designed by calculating the normalized compression size of various transformations applied to original images, a method that proved effective in quantifying the complexity of a broad array of visual arts. In another recent study, Lee et al. (27) have used a large-scale dataset of contemporary paintings sold at auctions to demonstrate that visual features play only a marginal role in predicting artwork prices.

Previous research thus demonstrates that, although art is traditionally regarded as qualitative or metaphysical, quantitative analyses—often inspired by physics—are not only feasible but have already yielded significant insights into the nature of art and contributed to more practical applications. Notable examples include the use of wavelet transforms for dating paintings (28), authenticating unknown or disputed artworks (29, 30), removing cradle artifacts (31), as well as other computational techniques for crack detection, digital inpainting of cracks (32), and separating overlapping images in X-ray scans of paintings (33, 34). Another prominent research avenue is the automatic classification of painting styles, which initially relied primarily on shallow learning approaches (28, 35–37) but, inspired by the success of convolutional neural networks in image processing (38, 39), has shifted to deep learning methods (40–44). These studies, particularly those employing wavelet transforms and convolutional neural networks, underscore the importance of encoding the spatial structures within paintings and other visual artworks to deepen our understanding of art. Spatial patterns are indeed essential for esthetic appreciation (45, 46) and are frequently used to qualitatively describe and distinguish the key characteristics of artistic movements. This focus on spatial structures is notably evident in the works of Wölfflin, a seminal figure in art history's development of formal analysis (47), who, in 1915, introduced five contrasting conceptual pairs to distinguish between Renaissance and Baroque art (48). These celebrated pairs—linear vs. painterly, plane vs. recession, closed vs. open form, multiplicity vs. unity, and clearness vs. unclearness—primarily address spatial patterns and formal structures in paintings rather than color aspects (49). A somewhat similar emphasis on spatial patterns also appears in the more recent works by art historians such as Sypher (50) and Davis (51).

Nevertheless, defining individual elements as units within artworks remains a complex challenge, despite the substantial research conducted to date. These units must be simple enough to apply universally, yet sufficiently detailed to encapsulate the essence of the paintings and their artistic styles. To address this challenge, we analyze ordering patterns among pixel intensities within two-by-two partitions in painting images using a crowd-sourced dataset comprising almost 140,000 artworks spanning a timeline of nearly a thousand years. Our study identifies 75 unique ordinal patterns, which we categorize into 11 groups based on arguments of continuity and symmetry. These patterns are easily interpretable and their prevalence correlates with low-level visual characteristics of paintings, including the presence of sharp edges, well-defined or textured elements, and vertical or horizontal structures. We observe a universal pattern in the distribution of two-by-two ordinal patterns across paintings, characterized by an approximately constant prevalence within each group, yet modulated across groups by a nontrivial interplay between pattern smoothness and the likelihood of finding identical pixel intensities. The existence of a universal prevalence in the incidence of ordinal patterns thus allows us to define a standardized measure that not only compares paintings and styles but also sets up a metric informing how distant a painting or an artistic style is from the average prevalence. Our research shows that the prevalence of ordinal patterns carries information for identifying the artistic style of paintings, outperforming baseline classifiers and further highlighting the critical role of patterns associated with identical values within the two-by-two partitions. This analysis also shows that styles typically display large variability in the prevalence of ordinal patterns, reflecting both the limitations of representing artworks through these patterns and the heterogeneous quality of our crowd-sourced dataset, as well as the intrinsic ambiguity often present in attributing a unique style to artworks. Additionally, we quantify the temporal evolution of paintings in our dataset by examining shifts in the prevalence of ordinal patterns, revealing a general trend wherein artworks increasingly diverge from the average incidence. Notably, this evolution is neither smooth nor homogeneous, showing substantial variability in the prevalence of ordinal patterns, especially after the 1930s.

In what follows, we detail these results, beginning with the data presentation and the definition of our two-by-two patterns. This is accompanied by an interpretation of these patterns in relation to the low-level visual features of paintings. We then examine the prevalence and variability of these patterns across all paintings in our dataset as well as within specific artistic styles. We further explore the potential of quantifying the evolution of paintings in our dataset by investigating changes in the incidence and variability of ordinal patterns. Finally, we conclude by summarizing our findings, acknowledging limitations related to the heterogeneous quality of the dataset, historiographic biases, and the localized nature of our ordinal patterns. The Methods section provides detailed information about the implementation of our approach.

Results

The dataset used in our research is the same as the one introduced in Sigaki et al. (23), comprising 137,364 paintings encompassing 154 styles from 2,391 artists and spanning a timeline of nearly a thousand years. These images were collected from Wikiart, a crowd-sourced, encyclopedic compilation of digital reproductions of artworks, covering primarily Western paintings. Given the dataset's reliance on heterogeneous sources, there is a substantial

variation in the quality of digital reproductions and in the meta-data associated with the artworks. These factors naturally constrain the generalizability and interpretation of the findings we shall present and discuss. Nevertheless, in the absence of large-scale, well-curated alternatives, the Wikiart dataset stands out as perhaps the best source for digital reproductions of paintings and for exploring quantitative approaches to assessing artwork features.

As detailed in the Data section, image files for each painting were converted to grayscale using the standard luminance transformation. Consequently, each painting is represented by a matrix A with n_y rows (image height) and n_x columns (image width), where each entry a_{ij} denotes the pixel intensity of the i th row and the j th column. Building on the ordinal framework inaugurated by Bandt and Pompe (52) and the recent classification of two-by-two ordinal patterns by Bandt and Wittfeld (53), we propose to characterize paintings by their distributions of ordinal patterns.

To describe our approach, let us consider a hypothetical three-by-three image

$$A = \begin{bmatrix} 4 & 2 & 9 \\ 5 & 3 & 1 \\ 2 & 6 & 3 \end{bmatrix}$$

and sample it using a two-by-two sliding partition that moves one pixel at a time, both horizontally and vertically, resulting in the four partitions:

$$A_0 = \begin{bmatrix} 4 & 2 \\ 5 & 3 \end{bmatrix}, \quad A_1 = \begin{bmatrix} 2 & 9 \\ 3 & 1 \end{bmatrix}, \quad A_2 = \begin{bmatrix} 5 & 3 \\ 2 & 6 \end{bmatrix}, \quad \text{and} \\ A_3 = \begin{bmatrix} 3 & 1 \\ 6 & 3 \end{bmatrix}.$$

Subsequently, we assign relative ranks to the pixel values within these partitions. Thus, representing the flattened elements of these partitions as (a_0, a_1, a_2, a_3) , the first partition $A_0 = (4, 2, 5, 3)$ is described by [2031] because $a_0 = 4$ is the second-largest element (rank 2), $a_1 = 2$ is the smallest element (rank 0), $a_2 = 5$ is the largest element (rank 3), and $a_3 = 3$ is the third-largest element (rank 1). Similarly, A_1 corresponds to [1320] and A_2 to [2103]. If a partition contains identical values, the same rank is attributed to them. Therefore, the partition A_3 , where the number 3 (rank 1) appears twice, corresponds to the pattern [1021]. We can identify up to 75 distinct ordinal patterns: 24 occur in partitions without identical values; 36 arise when two partition elements are equal (12 patterns for each of the three possible ranks of the identical values); 8 appear when three partition elements are identical (4 patterns for each of the two possible ranks of the identical values); 6 are found when there are two pairs of identical values within a partition; and 1 corresponds to a partition where all elements are identical.

In assessing all ordinal patterns of an image, we calculate their relative frequencies to define a probability distribution $P = \{p_i; i = 1, \dots, 75\}$, where p_i represents the relative frequency of the i th pattern. Unlike Bandt and Wittfeld's approach (53), our method explicitly incorporates the occurrence of identical pixel intensities rather than relying on their positions within partitions to resolve rank ties (see Ordinal patterns without ties section). Consequently, the total number of possible patterns expands from 24 to 75, and as we shall verify, these additional patterns provide essential information for a deeper understanding of paintings and artistic styles.

We classify the identified patterns into 11 groups, designated by capital letters A to K. This classification is based on the

number of distinct symbols each pattern exhibits and their smoothness degree, as depicted in Fig. 1. We observe that certain patterns correspond to rotations of other patterns. Thus, we further identify a subset of patterns as "primary patterns" because their rotations by quarter, half, and three-quarter turns generate all other patterns within their respective groups. Each group contains one, two, or four primary patterns, which are highlighted in Fig. 1 by dashed blue lines. For a given pattern represented as $[x_1 \ x_2 \ x_3 \ x_4]$, where each symbol x_i (for $i = 1, \dots, 4$) takes a value from the set $\{0, 1, 2, 3\}$, we quantify smoothness using three measures: sum of the absolute values of first-order differences among symbols ($\Delta = |x_1 - x_2| + |x_3 - x_4| + |x_1 - x_3| + |x_2 - x_4|$); maximum absolute value of these differences [$\delta^* = \max(|x_1 - x_2|, |x_3 - x_4|, |x_1 - x_3|, |x_2 - x_4|)$]; and absolute value of the first-order difference between symbols aligned along either a row or a column of identical symbols (δ^{\parallel}). The values of δ^{\parallel} are thus only determined for patterns that feature two identical elements aligned along a row or column. For instance, the pattern [0001] yields $\Delta = 2$, $\delta^* = 1$, and $\delta^{\parallel} = 1$, whereas the pattern [0110] results in $\Delta = 4$ and $\delta^* = 1$. We initially rank the smoothness of patterns based on Δ ; δ^* is subsequently used to resolve cases of equal Δ values, and δ^{\parallel} is used similarly when patterns display identical Δ and δ^* values.

The exclusive one-symbol pattern [0000] is the sole member of group A. Characterized by its uniformity, this pattern exhibits the maximum smoothness with $\Delta = 0$, indicative of identical pixel intensities within a partition. Two-symbol patterns occur in partitions that exhibit pairs or trios of identical values and are categorized into types D, H, and K. Among these, type D represents the most homogeneous configuration ($\Delta = 2$, $\delta^* = 1$, and $\delta^{\parallel} = 0$), typified by pairs of equal values aligned either in a column or a row with the primary pattern being [0011]. Type H, slightly less uniform, differs from type D by $\delta^{\parallel} = 1$ and corresponds to trios of identical values that form corners within partitions whose primary patterns are [0001] and [0111]. Type K, the least uniform within this group ($\Delta = 4$), consists of patterns where pairs of equal values are aligned across both principal and secondary diagonals, as illustrated by the primary pattern [0110]. Three-symbol patterns emerge when partitions contain only a single pair of identical values; they are distributed across four groups: F, G, I, and J. Type F, with [0112] as the primary pattern, is the smoothest within this category ($\Delta = 4$ and $\delta^* = 1$), displaying a configuration where the pair of identical values has the intermediate rank and is positioned diagonally, creating corner-like structures akin to type H. Type G is delineated by four primary types ([0012], [0021], [0122], and [0212]) and type I by two ([0121] and [0211]), both characterized by $\Delta = 4$ and $\delta^* = 2$ and corresponding to a configuration where the pair of identical values is located either along a column or a row. In type G, the identical values assume the extremal ranks, whereas in type I, they have the intermediate rank. This nuanced distinction results in a shift in δ^{\parallel} from 1 to 2, rendering type I less uniform compared to type G. Type J, identified by primary types [0120] and [0221], features the pair of identical values with the extremal ranks along the diagonals. The only difference between types F and J is the rank of their diagonal elements, but this variation alters the smoothness metric Δ from 4 to 6, categorizing type J as the least smooth among the three-symbol patterns.

Patterns with four symbols emerge when there are no identical values within a partition and comprise types B, C, and E. These are equivalent to the patterns named types I, II, and III by Bandt and Wittfeld (53) (see Ordinal patterns without ties section). Type B patterns, with [0123] and [0213] as primary patterns, occur when the partition values align in a 3D plane, indicating simultaneous

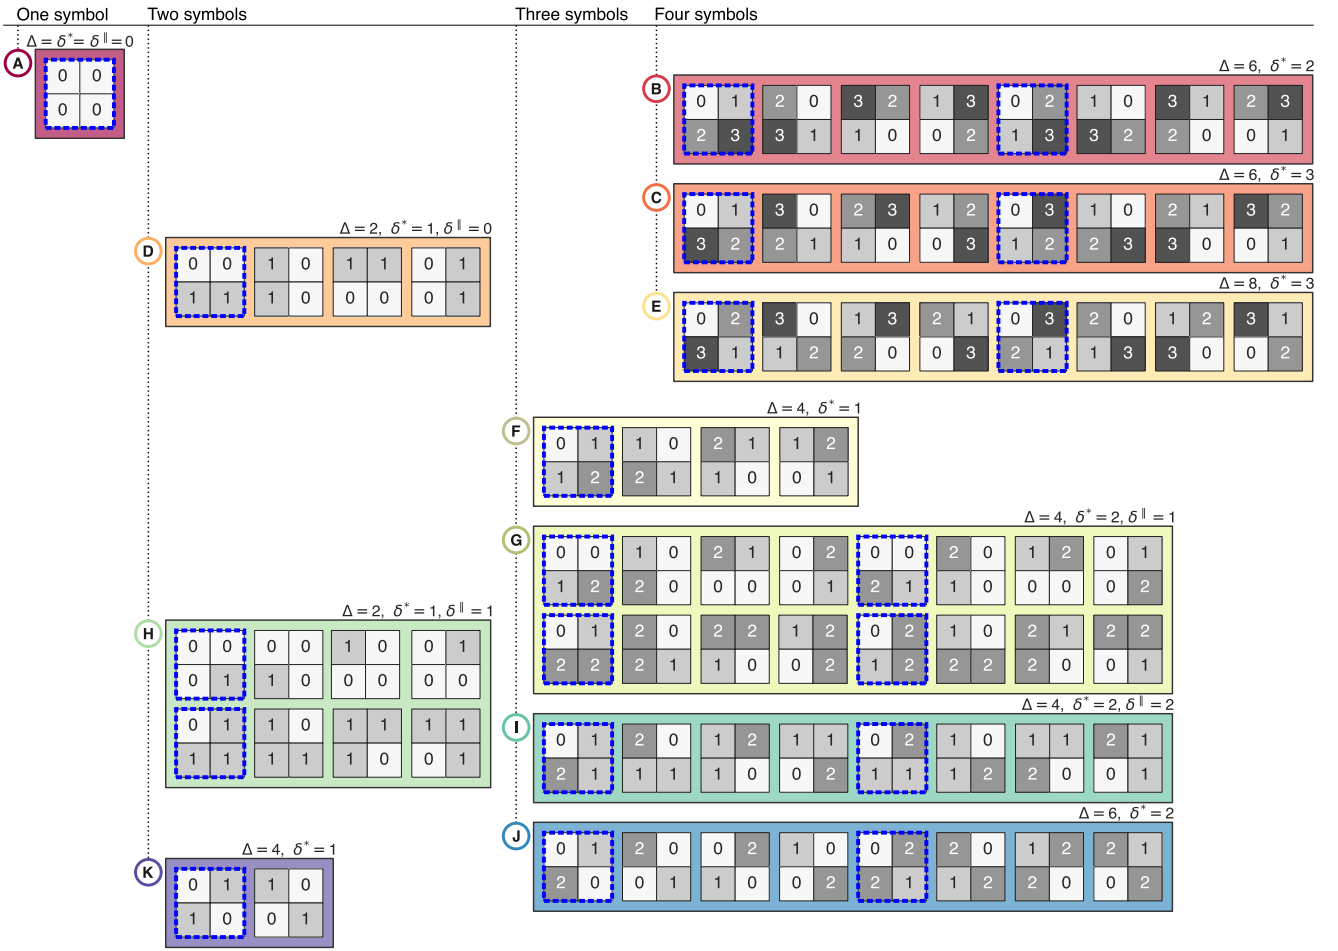


Fig. 1. Overview of the 75 possible two-by-two ordinal patterns categorized into 11 groups based on their number of unique symbols and smoothness. Patterns are designated by capital letters from A to K and distinguished by different background colors. Smoothness is quantified by three metrics: sum of the absolute values of first-order differences among symbols (Δ); maximum absolute value of these differences (δ^*); and absolute value of the first-order difference between symbols aligned with either a row or column of identical symbols ($\delta^||$). The pattern [0000], the sole type A pattern, exhibits maximal smoothness with $\Delta = \delta^* = \delta^|| = 0$ and corresponds to identical values in the associated partition. Groups D, H, and K, ranked by smoothness, contain patterns with two distinct symbols. Type D consists of four patterns with $\Delta = 2$, $\delta^* = 1$, and $\delta^|| = 0$. Type H includes eight patterns with $\Delta = 2$, $\delta^* = 1$, and $\delta^|| = 1$. Type K comprises two patterns with $\Delta = 4$ and $\delta^* = 1$. Patterns with three distinct symbols are categorized into four groups: F, G, I, and J, in order of smoothness. The four type F patterns are characterized by $\Delta = 4$ and $\delta^* = 1$. Type G, the largest group, includes 16 patterns with $\Delta = 4$, $\delta^* = 2$, and $\delta^|| = 1$. Type I comprises eight patterns, distinguished from type G by $\delta^|| = 2$. Type J includes eight patterns with $\Delta = 6$ and $\delta^* = 2$. Finally, patterns with four distinct symbols are organized into groups B, C, and E, each containing eight patterns. Type B is the smoothest ($\Delta = 6$ and $\delta^* = 2$), followed by type C ($\Delta = 6$ and $\delta^* = 3$), and type E ($\Delta = 8$ and $\delta^* = 3$). Primary patterns within each group are highlighted with dashed blue lines. Rotations of primary patterns by quarter, half, and three-quarter turns generate all other patterns within each group.

increases or decreases along both columns and rows. Type C, with primary patterns [0132] and [0312], displays $\Delta = 6$ and $\delta^* = 3$ and is less smooth than type B ($\Delta = 6$ and $\delta^* = 2$). These patterns arise when pixel intensities within partitions simultaneously increase or decrease along one direction but exhibit distinct trends along the other. For instance, in [0132], intensities in the first row increase from left to right, whereas those in the second decrease, with both columns increasing from top to bottom. Type E is the least smooth among four-symbol types and comprises the primary patterns [0231] and [0321]. In these patterns, pixel intensities within partitions display distinct trends along both rows and columns. For example, in [0231], the first row increases from left to right, while the second decreases and the first column increases from top to bottom while the second decreases.

We thus calculate the probability distribution of ordinal patterns $P = \{p_i; i = 1, \dots, 75\}$ across all paintings in our dataset. In addition, we use our classification of ordinal patterns to arrange the elements of P according to the type they belong from A to K

and, within each type, we organize the probabilities p_i as they appear in Fig. 1. Consequently, p_1 corresponds to the probability of finding the pattern [0000], p_2 is associated with [0123], p_3 refers to [0231], and so forth, up to p_{75} that is related to [1001]. By doing so, we are mapping each image into a 75D vector, where each dimension represents the relative frequency of a specific two-by-two ordinal pattern. Although 75 dimensions certainly correspond to a high-dimensional space, it is significantly reduced compared to the original image resolution, which in our dataset typically exceeds one million pixels. Furthermore, these ordinal patterns all have straightforward interpretations and can be further categorized into 11 types.

Figure 2 presents a visualization of the probability distribution P for all paintings using the uniform manifold approximation and projection (UMAP) dimensionality reduction algorithm (54) (see Low-dimensional projection of the ordinal pattern probabilities section). The UMAP projects the 75D space of P onto a plane, striving to preserve both local and global structures. In this projection,

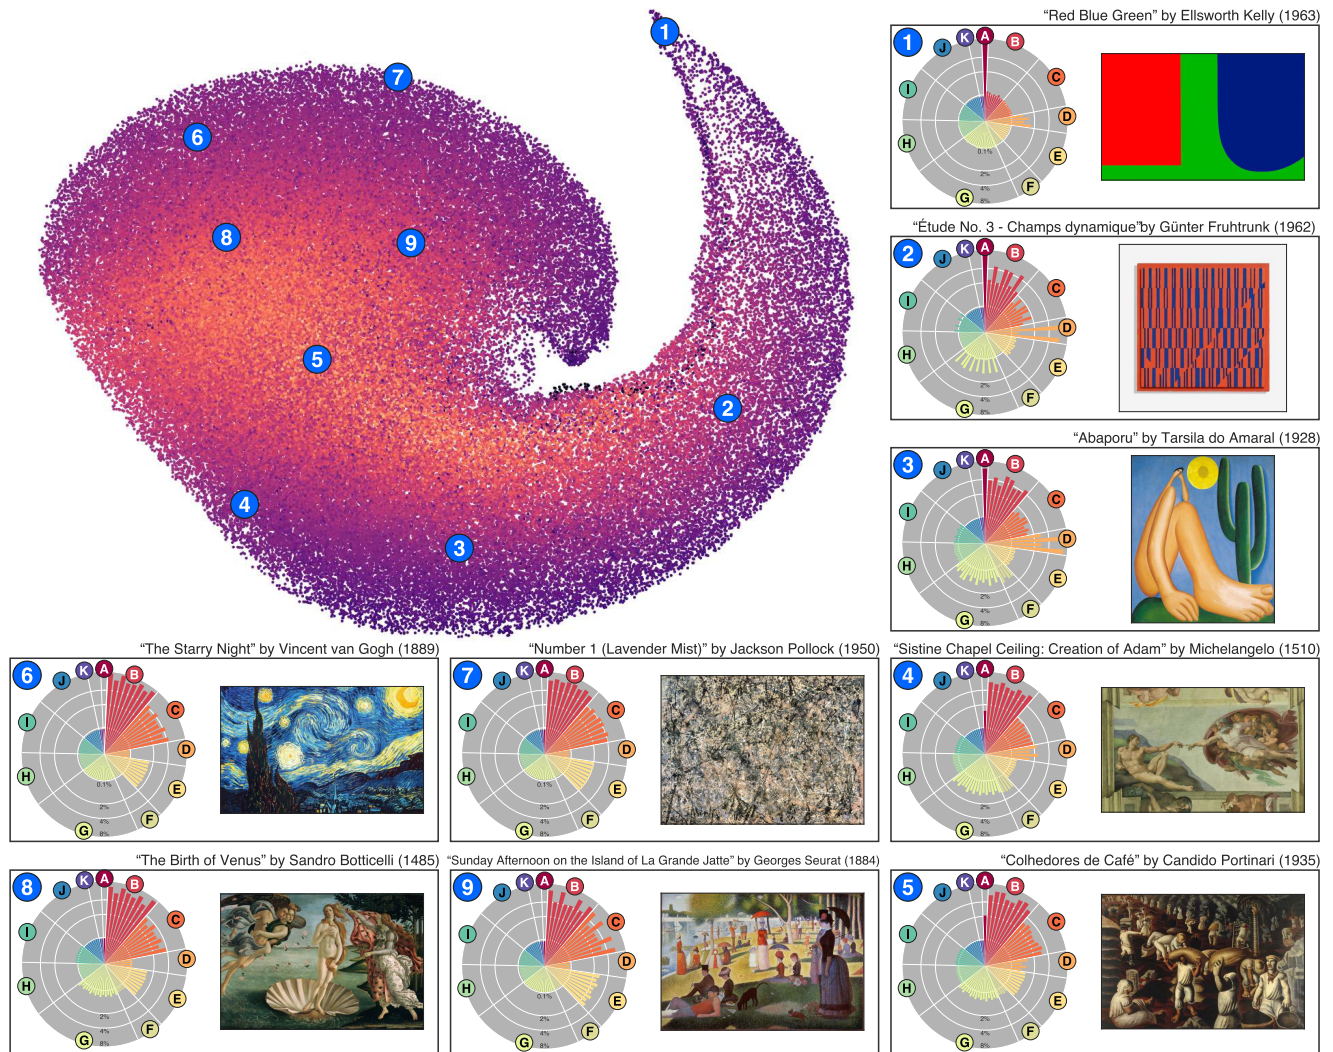


Fig. 2. Representing paintings using two-by-two ordinal patterns. All paintings in our dataset are represented as 75D vectors, with each dimension corresponding to the relative frequency of a specific two-by-two ordinal pattern. The main panel displays a 2D projection obtained from UMAP, where each dot corresponds to a painting. The numbers within blue circles indicate the positions of nine selected artworks. Surrounding panels illustrate the distribution of ordinal patterns, grouped by type, using circular bars on a logarithmic scale, accompanied by thumbnail images of the corresponding paintings. Paintings characterized by simple, well-defined design elements with sharp edges, such as the hard-edge painting by Ellsworth Kelly, typically exhibit a high frequency of the type A pattern. In contrast, paintings lacking sharp edges and well-defined elements, such as Jackson Pollock's drip paintings, display a negligible presence of this pattern. The ordinal patterns also reveal subtle nuances in artworks, such as the predominance of edges in specific directions, exemplified by the high frequency of two type D patterns in the Concretism work by Günter Fruhtrunk, or significant discontinuities among neighboring pixels, as observed in the pointillism artwork by Georges Seurat through the high frequency of two type E patterns.

paintings or clusters of paintings that are proximal typically exhibit similar frequencies of ordinal patterns, while those that are distant display more dissimilar distributions. Notably, while UMAP preserves meaningful local structures (55), it does not preserve high-dimensional distances (56); thus, insights derived from these projections remain primarily exploratory (55). We annotate the positions of nine paintings, depicting their ordinal pattern distributions alongside images of the artworks. These distributions provide insights into various visual characteristics of the paintings. For instance, Ellsworth Kelly's "Red Blue Green" (1963) predominantly features type A and D patterns (with type A being much more frequent), reflecting the three sharply delineated colored regions of this abstract painting of the style Hard Edge Painting. We find a somewhat similar pattern in Günter Fruhtrunk's "Étude No. 3—Champs dynamique" (1962) but with a significant presence of patterns [1010] and [0011], indicative of the painting's dominant vertical structures. Although more

subtly, Tarsila do Amaral's "Abaporu" (1928) also reveals these vertical structures through a high occurrence of the same patterns. Conversely, artworks such as Jackson Pollock's "Number 1 (Lavender Mist)" (1950), Vincent van Gogh's "The Starry Night" (1889), and Georges Seurat's "Sunday Afternoon on the Island of La Grande Jatte" (1884) practically do not exhibit the type A pattern. These ordinal patterns also capture more subtle nuances of artworks. For instance, Seurat's work is considered a notable example of the pointillist style of painting, in which the painter uses small brushstrokes to produce dots of locally distinct colors that, in turn, yield significant discontinuities among neighboring pixels. These discontinuities produce a relatively high frequency of two type E patterns. Furthermore, we observe a higher frequency of four type B patterns ([0123], [3210], [1032], and [2301]) compared to the other four ([2031], [1302], [0213], and [3120]) in Seurat's work. The patterns with high prevalence have smaller first-order differences horizontally than vertically, indicating

that pixel intensities change more abruptly along the vertical direction and agreeing with the fact that Seurat started his artwork by applying a layer of small horizontal brushstrokes. Similar trends are observed in type C patterns, where patterns with greater first-order vertical differences along the vertical direction ([0132], [2310], [1023], and [3201]) are more prevalent.

It is important to acknowledge, however, that the frequency of ordinal patterns may be influenced by the resolution of the digital representation of the artwork, particularly for large paintings, where only high-resolution scans can adequately capture finer details. Seurat's "Grande Jatte," executed on a monumental canvas ~3 m wide by 2 m high, exemplifies this limitation, especially as we obtained a high-resolution representation of this artwork at 30,000×19,970 pixels (Fig. S1A). At this resolution, the small brushstrokes characteristic of this pointillist painting may emerge as well-defined regions, potentially increasing the prevalence of pattern types associated with identical pixel intensities, such as the type A pattern. To examine this effect, we analyze downscaled versions of the high-resolution image, starting from 200×300 pixels (equivalent to sampling one pixel every 100 pixels along both the horizontal and vertical axes) up to the original resolution, incrementally increasing pixel sampling by one unit at each step and estimating the ordinal pattern distribution at each resolution. We observe that the prevalence of the type A pattern ([0000]) rises from 0 to 0.24% as the pixel density increases, with similar behavior for pattern types D and H, and others associated with identical pixel intensities (Fig. S1B and C). Despite the primary features of the ordinal distribution remaining relatively stable across resolutions (Fig. S2), this illustrative case underscores the importance of caution when associating features of ordinal distributions with aesthetic patterns in artworks.

Considering the previous limitations, the examples highlighted in Fig. 2 illustrate the insights that can be gleaned from analyzing distributions of ordinal patterns. Integrating this analysis with domain-specific knowledge from the art field and high-quality digitalization of artworks offers a promising method for the individual examination of paintings. From a broader perspective, it raises the question of whether a universal prevalence of ordinal patterns exists in paintings. This query was also explored by Bandt and Wittfeld (53) using their three types of ordinal patterns. In their study using a dataset of 25 natural textures, they observed that type I patterns are more prevalent than type II, which in turn surpasses type III in prevalence. This finding indicates that the prevalence of patterns correlates with their degree of smoothness, with smoother patterns appearing more frequently in natural textures. Despite the limited scale and variability of their dataset, it is plausible to assume that the smoothness of our 11 pattern types influences their occurrence. Indeed, that is the primary rationale for using smoothness to categorize our types. However, it is reasonable to suppose that the number of identical values within a partition also impacts the prevalence of corresponding patterns. To address this question, we calculate the average prevalence of the 75 ordinal patterns across all images in our dataset. Figure 3A reveals a clear hierarchical organization in the prevalence of ordinal patterns. Their average probabilities monotonically decrease across alphabetically arranged types and patterns within each group exhibit nearly identical prevalence. The smoothest pattern, type A, is also the most frequent, whereas the less frequent patterns belong to group K. Within groups sharing the same number of symbols (identical values), the average probability monotonically decreases as the roughness of their patterns increases. For instance, among patterns with two symbols, type D is more prevalent than type H, which is more prevalent

than type K, following the same sequence in smoothness. This behavior is consistent for patterns with three and four symbols, ranked by prevalence and smoothness as (F, G, I, J) and (B, C, E), respectively. Conversely, no clear relationship between smoothness and prevalence exists across types with different numbers of symbols. Therefore, the prevalence of pattern types reflects a non-trivial interplay between their smoothness and the reduced likelihood of encountering identical values in two-by-two partitions.

We further explore the variability in the prevalence of ordinal patterns. Figure 3B and C depict the SD and the coefficient of variation (the ratio of SD to mean) of the probability of each pattern across all images. The ranking by SD largely corresponds to that by average prevalence, except for type D patterns, which exhibit the second-largest SD. In contrast, the coefficient of variation presents a remarkably different ranking, with low-prevalence patterns (such as J and K) showing significantly higher relative dispersion than high-prevalence patterns (such as B and C). This finding suggests that even low-prevalence patterns may carry crucial information for the analysis of paintings, as visually confirmed in Fig. 2. It also underscores the importance of caution when comparing raw prevalence values among different paintings and artistic styles, since minor changes can correspond to substantial variations relative to the SD for several patterns.

To account for these distinct scales of variations, we define a standardized measure of prevalence, the z-score, as

$$z_i = \frac{p_i - \mathbb{E}[p_i]}{\mathbb{S}[p_i]},$$

where p_i is the probability of finding the i th pattern in a painting, and $\mathbb{E}[p_i]$ and $\mathbb{S}[p_i]$, represent, respectively, the average and the SD of its prevalence estimated across all images. The value of z_i quantifies how many SD units the prevalence of the i th pattern in a specific painting is above ($z_i > 0$) or below ($z_i < 0$) its overall incidence across all images. The values of z_i for all 75 patterns are comparable in scale and further allow us to measure how distant the distribution of ordinal patterns of a painting is from the average prevalence. Specifically, we define this distance to the average distribution as the absolute sum of the z-score probabilities $s = \sum_{i=1}^{75} |z_i|$. Indeed, we have already used the values of s to color-code the UMAP projection in Fig. 2, with lighter shades representing small distances and darker shades indicating large distances. Thus, in this representation, the proximity of paintings to the average pattern decreases as they move further from the center of the spiral-like shape. For instance, Candido Portinari's "Colhedores de Café" (1935) is near the center of the spiral with $s \approx 16$, while Ellsworth Kelly's "Red Blue Green" (1963) is positioned at one end of the spiral with $s \approx 113$.

As the values of s encapsulate information from all ordinal patterns, they provide an opportunity to globally assess whether the heterogeneous quality of digital reproductions in our dataset introduces any clear bias in our analysis. To examine this aspect, we estimate the association between the distance to the average pattern s and the total number of pixels for each image in our dataset. We also explore the relationship between s and the pixel density for ~20% of the images for which the physical dimensions of the artworks were available on Wikiart. The results indicate no evident association between s and either image size or pixel density (Fig. S3), as the correlations between s and these image properties are very weak. These findings suggest that ordinal patterns are not significantly biased by the heterogeneous quality of the

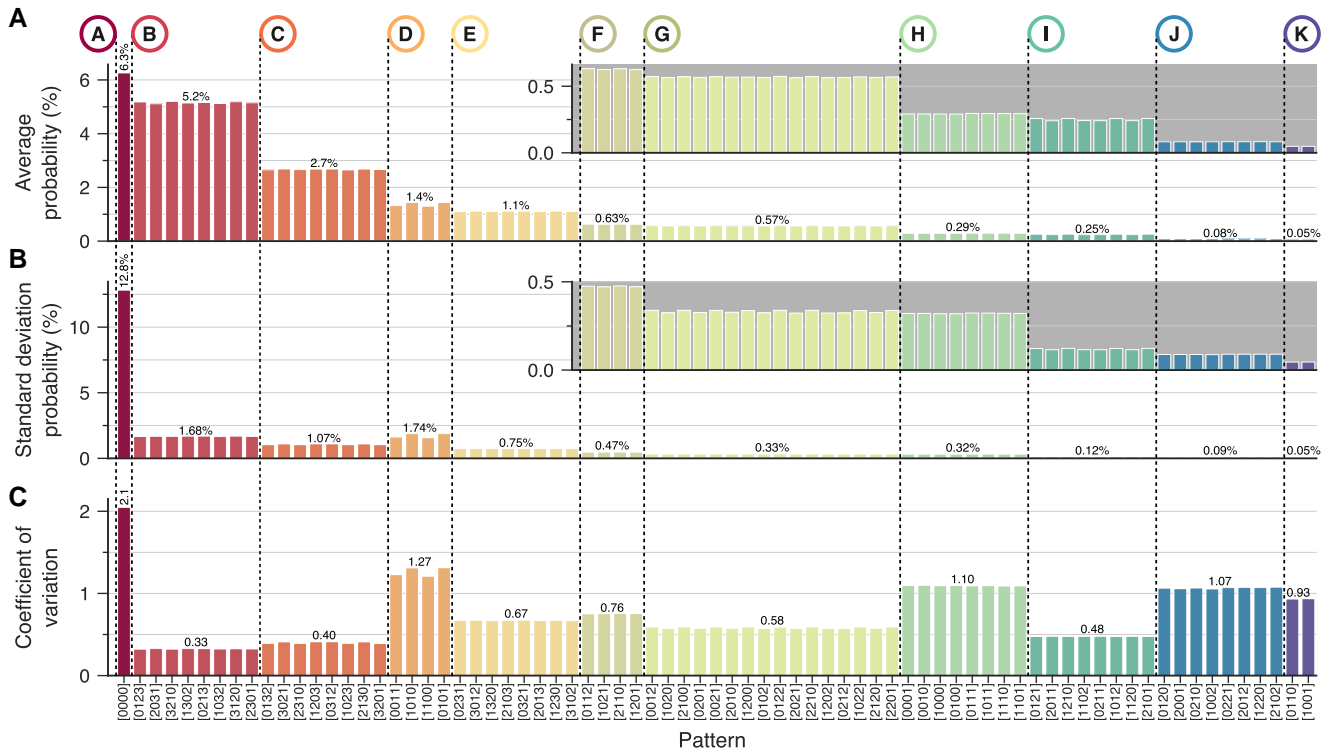


Fig. 3. Universal patterns in the occurrence of two-by-two ordinal patterns. A) Average probability of finding each of the 75 ordinal patterns across all images in our dataset. These patterns are categorized into 11 groups (A to K) and represented by different colors. Within each group, the probability of finding patterns remains approximately constant but decreases monotonically across alphabetically arranged groups. Within groups sharing the same number of distinct symbols, average probability diminishes as the roughness of the patterns increases. For example, type D patterns exhibit the greatest smoothness among types with two different symbols, followed by types H and K, with average probabilities of 1.4, 0.29, and 0.05%, respectively. Across all types, the sequence of occurrence reflects a nontrivial interplay between pattern smoothness and the diminishing likelihood of repeating symbols. For instance, type F and K patterns, comparable in levels of smoothness, differ significantly in frequency; type K is less common due to the lower probability of encountering just two distinct symbols, compared to three in type F patterns. B) SD and C) coefficient of variation (SD divided by the mean) for each ordinal pattern across all images in our dataset. The ordering of the SD largely mirrors that of the average probabilities, except for type D patterns, which present the second-highest SD. Conversely, the coefficient of variation across pattern types displays a distinct behavior, with types A, D, H, J, and K showing the highest coefficients of variation, respectively.

images in our dataset. Nevertheless, caution remains essential when examining specific nuances of individual artworks, as illustrated in our prior discussion of Seurat's "Grande Jatte."

Using these z-score probabilities, we further investigate whether different artistic styles exhibit distinct incidences of ordinal patterns. We select 92 styles, each represented by at least 100 paintings and calculate the average z_i after grouping the images by style. Additionally, we evaluate the average value of s to rank the styles according to their deviation from the average pattern. Figure 4 depicts these averages in a matrix plot where styles are organized in descending order of s , with the latter represented by a horizontal bar plot. Furthermore, to quantify the relative dispersion in the incidence of ordinal patterns within styles, we estimate the coefficients of variation of the z-score probabilities of each ordinal pattern, averaging their absolute values across all ordinal patterns. These averaged coefficients of variation for each style are also shown as horizontal bars in Fig. 4. In the matrix visualization, shades of green indicate patterns exceeding the overall incidence, while shades of pink represent patterns below this incidence. Each row corresponds to the ordinal fingerprint of a style, linking the prevalence of ordinal patterns to the visual characteristics of each style.

Although a comprehensive analysis of all styles is beyond the scope of our study, several noteworthy global features exist. For example, Divisionism and Pointillism show a lower incidence of types A, D, and H, as well as other patterns featuring identical

values along rows or columns (types G and I). In contrast, more discontinuous types, such as B, C, and E, are more prevalent. Despite the limitations of our dataset, these ordinal features seem to reflect the textured visuals and often indistinct edges characteristic of Divisionist and Pointillist compositions. Light and Space, in turn, display almost the opposite behavior, with a higher prevalence of continuous types (such as A, D, and H) and a lower incidence of more discontinuous types (such as types B, C, and E). Notwithstanding the constraints of our data, these ordinal features align with the geometric and well-defined shapes typical of this style. These two styles also differ markedly regarding the degree of dispersion in the prevalence of ordinal patterns, as quantified by their average coefficients of variation, which are significantly higher for the Light and Space style. Most styles, indeed, show higher average coefficients of variation that tend to increase as they approach the overall prevalence of patterns (Fig. S4). In addition to the inherent oversimplification involved in representing artworks solely through two-by-two ordinal patterns, this higher variability may also relate to the heterogeneous quality of our dataset and the challenges of obtaining reliable style annotation, which depend on specialized art and art history expertise often lacking in crowdsourcing annotators. This variability further reflects the intrinsic ambiguity of attributing a unique style to artworks, as style evolution tends to be gradual, and paintings may incorporate elements from multiple styles (43). Additionally, certain styles are more narrowly defined by

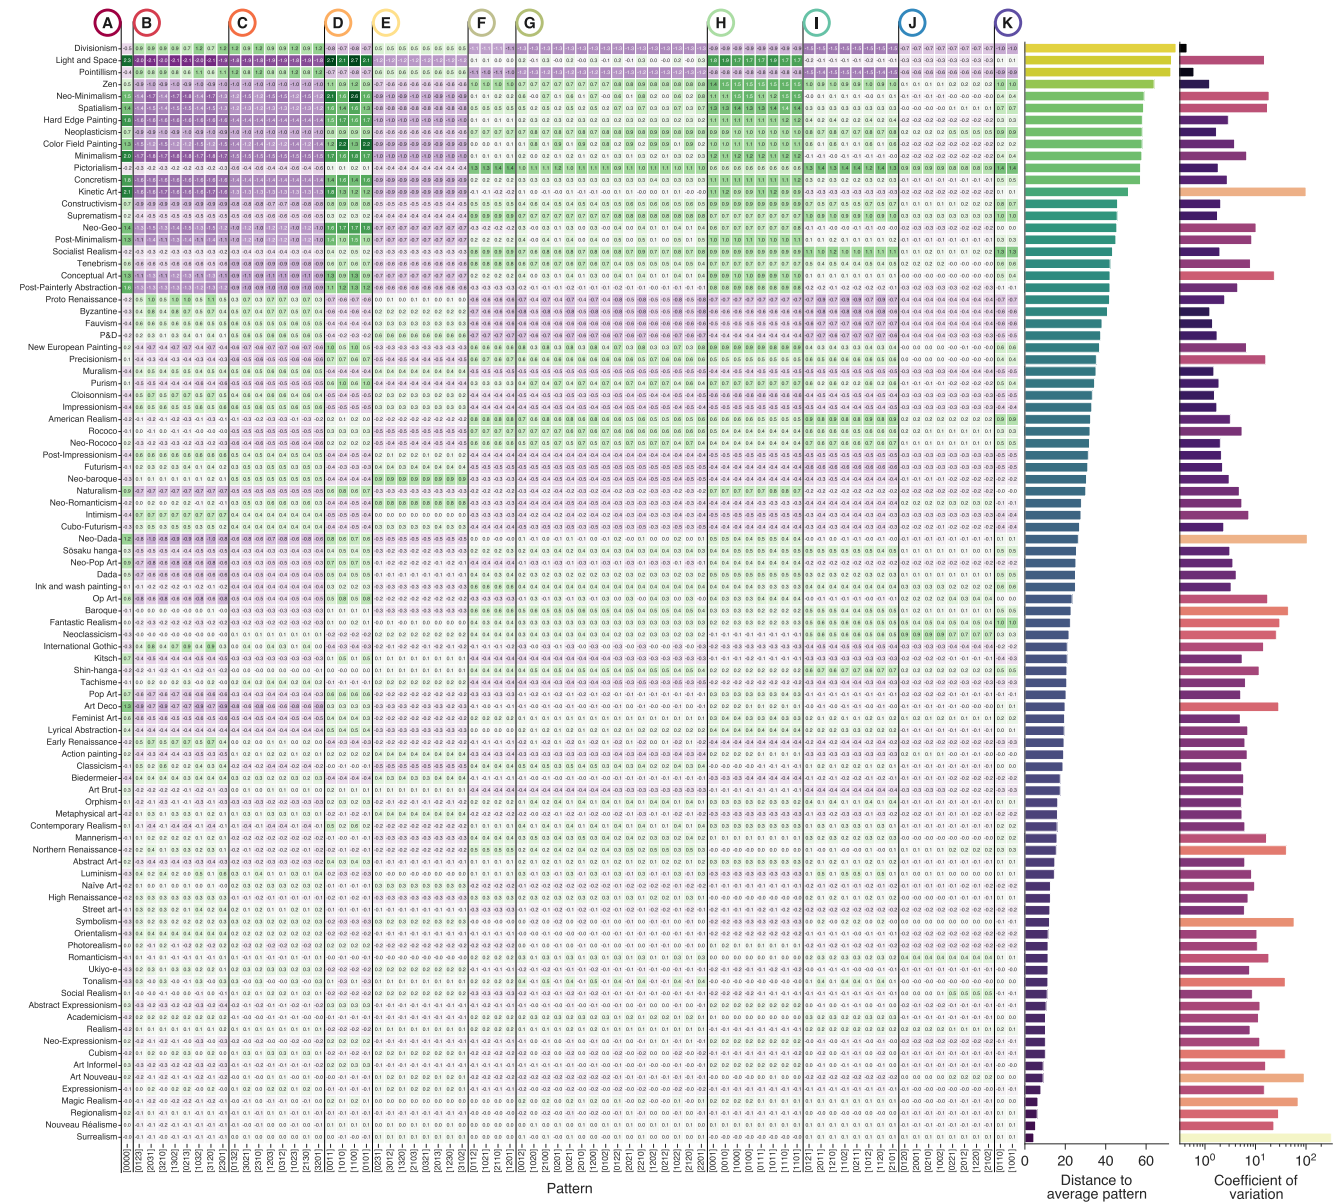


Fig. 4. Distinguishing artistic styles using z-score probabilities of two-by-two ordinal patterns. The lines in the matrix plot depict the z-score probability of each ordinal pattern for 92 artistic styles, each containing at least 100 artworks (~90% of data). The z-score probability is calculated by subtracting the average probability of an ordinal pattern within a particular style from the overall average probability and dividing the result by the overall SD of that pattern's probability across all images in the dataset. Positive values (represented by shades of green) indicate patterns occurring more frequently than the overall average, whereas negative values (represented by shades of purple) denote patterns occurring less frequently than the overall average, with deviations expressed in SD units. Styles are arranged in descending order based on their distance to the average pattern, determined by the absolute sum of the z-score probabilities, which is visualized in a horizontal bar plot on the right of the matrix plot. The right-most panel shows the average of the absolute values of the coefficients of variation of the ordinal probabilities (in z-scores units) across all ordinal patterns.

their “family resemblance,” whereas others are not. For example, Dada, Neo-Dada, and Surrealism are not generally regarded as cohesive esthetic doctrines (57, 58) and encompass heterogeneous works featuring a wide range of forms and techniques, which may also explain their high coefficients of variation.

We further investigate the possibility of using the incidence of ordinal patterns to predict the artistic style of paintings through a machine learning approach. In order to have enough training examples, we consider the 100,707 paintings from the 20 most common styles in our dataset, each represented by at least 1,500 artworks. Our task is to determine the style of these images using their distributions of ordinal patterns ($P = \{p_i; i = 1, \dots, 75\}$) as predictive features. Additionally, we compare the accuracy of this

approach against other ordinal quantifiers that do not consider patterns associated with identical values, namely: the complexity-entropy plane (23, 59, 60), the Fisher-Shannon plane (61), the smoothness-structure plane (53), the incidence of the three pattern types defined by Bandt and Wittfeld (53) (types I, II, and III), the probability of all 24 patterns disregarding ties, and the same set of 24 patterns obtained randomly resolving rank ties. These approaches are detailed in the Ordinal patterns without ties section. We also aggregate the probabilities of all 75 patterns into their 11 categories and consider the incidence of these pattern types as predictive features for a painting's artistic style. As outlined in the Machine learning artistic styles section, we rely on the extreme gradient boosting (XGBoost) (62) algorithm

to predict the artistic style of paintings using the incidence of ordinal patterns and all other sets of ordinal features. It is important to note that, beyond previously discussed limitations related to style definition, the simplification inherent in reducing paintings to ordinal patterns, and the heterogeneous quality of our dataset, the primary objective of this analysis is not to develop a state-of-the-art style classification system comparable with deep learning approaches (37, 40–43). Instead, our goal is to assess the informational value of accounting for identical pixel intensities in two-by-two patterns and the effectiveness of categorizing these patterns into 11 distinct types.

Figure 5A presents the average accuracy of our predictions and compares these results with those from baseline classifiers that either make uniformly random predictions (uniform) or generate random predictions respecting the style distribution (stratified). All sets of ordinal features achieve accuracy levels significantly higher than those of the baseline classifiers, indicating that ordinal patterns carry stylistic information about paintings, as also discussed by Sigaki et al. (23). Notably, using all 75 ordinal patterns as predictive features results in a significantly enhanced accuracy compared to all other sets of ordinal features. This improvement

is particularly pronounced when compared with the complexity-entropy, Fisher–Shannon, and smoothness-structure planes, as well as the incidence of types I, II, and III, and the set of all 24 ordinal patterns with rank ties resolved randomly. These former approaches all approximate an accuracy of 0.19, whereas using all two-by-two ordinal patterns achieves 0.28. Furthermore, the smoothness-structure plane, the incidence of types I, II, and III, and the set of all 24 ordinal patterns with rank ties resolved randomly show only marginally higher accuracy than the complexity-entropy and Fisher–Shannon planes. This result demonstrates that permutation entropy combined with statistical complexity or the Fisher information measure effectively compresses information from ordinal distributions, despite the meticulous theoretical underpinnings of the smoothness-structure plane and patterns types I, II, and III (53). Additionally, randomly resolving rank ties yields lower accuracy than the set of 24 ordinal patterns without ties, which achieves an accuracy of 0.25. While identical values are not explicitly addressed in this latter approach, their occurrences are deterministically handled based on the positions of equal values, thereby retaining additional ordering information rather than randomly resolving ties. Interestingly, although

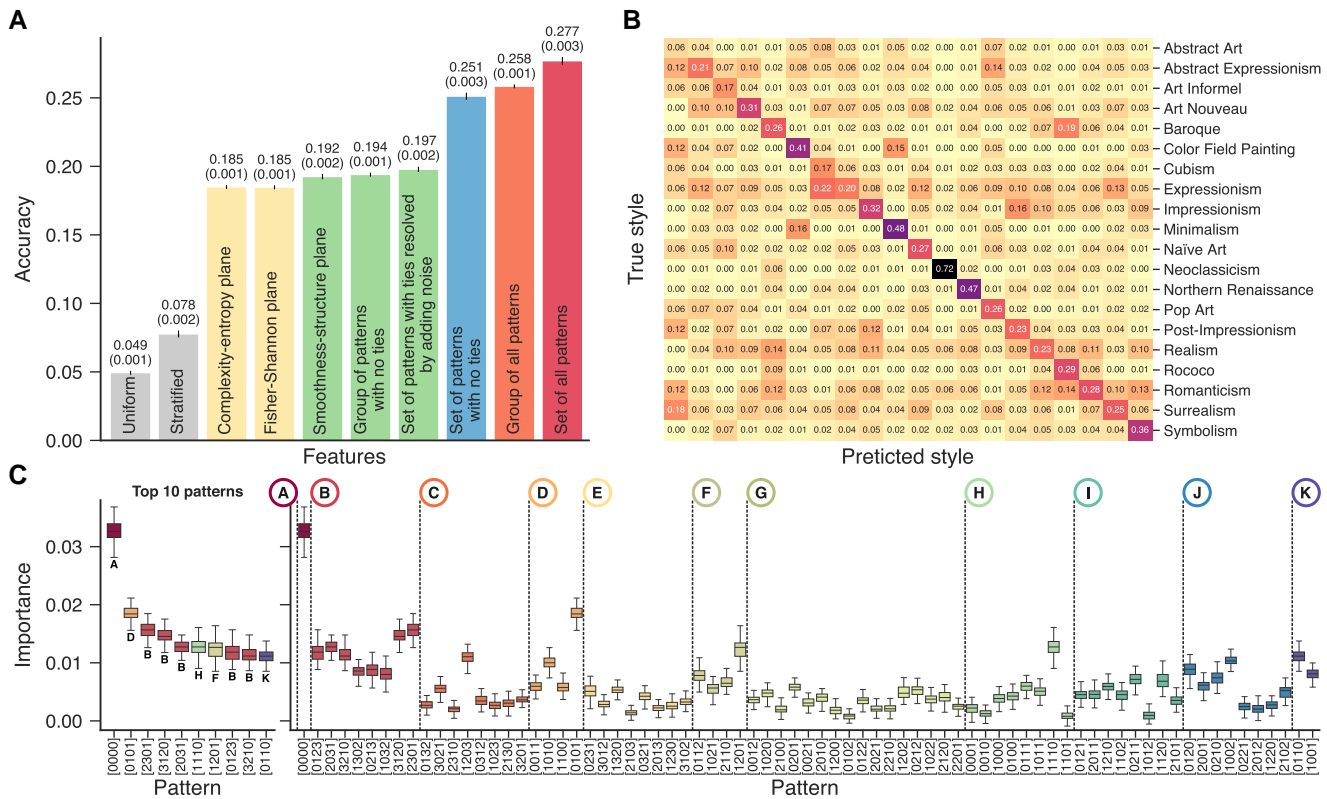


Fig. 5. Machine learning artistic styles using two-by-two ordinal patterns. A) Accuracy of classifiers trained to predict artistic styles from paintings using varied ordinal features. The dataset includes images from 20 styles, each with at least 1,500 artworks, divided into training (80%) and test sets in a stratified manner. Bars represent the average accuracy across ten independent train-test splits and training procedures, with small error bars showing the SD of accuracy. Numbers above bars indicate average accuracies, with values in parentheses representing SD. Gray bars denote accuracy from dummy classifiers using uniformly random (uniform) and frequency-based (stratified) predictions. Yellow bars depict XGBoost predictions from two pairs of quantifiers derived from the standard ordinal distribution (patterns without ties): the complexity-entropy plane ($H \times C$) and the Fisher–Shannon plane ($H \times F$). Green bars illustrate XGBoost predictions from the smoothness-structure plane ($\kappa \times \tau$), probabilities of pattern groups without ties (q_1 , q_2 , and q_3), and the ordinal distribution with ties resolved by adding small noise to the original images. The blue bar indicates XGBoost predictions from the standard ordinal distribution (patterns without ties). The orange bar represents predictions based on probabilities associated with the 11 pattern types (q_A , q_B , q_C , ..., and q_K). Finally, the red bar shows XGBoost predictions using probabilities related to all 75 two-by-two ordinal patterns. B) Confusion matrix for the XGBoost model trained with the full set of ordinal probabilities and applied to the test set. C) Box-plots displaying the permutation importance of each two-by-two ordinal pattern in predicting artistic styles. Permutation importance measures accuracy degradation after randomizing probabilities of an ordinal pattern among all images. The left panel highlights the ten most important patterns, and the right panel displays the importance of each ordinal pattern. Horizontal lines within boxes indicate medians, boxes denote the interquartile range, and whisker bars represent extreme observations.

comprising less than half the dimensionality of the set of 24 ordinal patterns without ties, the incidence of the 11 pattern types yields the second-highest accuracy. This finding indicates that these pattern groups effectively compress the information encoded in all 75 two-by-two ordinal patterns. However, despite the symmetries within each pattern type, the lower performance of classifiers trained with these grouped patterns suggests that each ordinal pattern provides unique and critical information for identifying the style of a painting.

We also calculate the confusion matrix for the classifier trained with all two-by-two ordinal patterns, as shown in Fig. 5B. As expected, the overall accuracy of ~28% achieved with our ordinal features is modest compared to more sophisticated machine learning approaches, such as convolutional neural networks, which attain accuracies above 70% in similar tasks (40–43). Nonetheless, our approach demonstrates that the 75 ordinal patterns contain critical information for classifying styles that are not present in the standard distribution of ordinal patterns nor in its usual entropy measures. Moreover, the confusion matrix exhibits a well-defined diagonal structure, indicating that the most frequent predictions across all styles align with the actual styles of images. For instance, the trained algorithm correctly identified Neoclassicism, Minimalism, and Northern Renaissance artworks in 72, 48, and 47% of the predictions, respectively. However, our model struggles with Abstract Art, correctly identifying this style in only 6% of predictions. The low performance with Abstract Art highlights the difficulties in assigning a single style to works encompassing a diverse range of techniques, which often lack distinctive, consistent features and can vary widely among artists and even within a single artist's oeuvre. Similar observations were made by Karjus et al. (26), who used compression ensembles to classify styles and noted that certain styles are more frequently misclassified, suggesting that the structure of these errors may hold interpretive value. Additionally, our algorithm tends to misclassify more styles that exhibit high coefficients of variation in the prevalence of ordinal patterns and low distance to the average pattern (such as Cubism and Art Informel), while styles with low coefficients of variation and greater distance from the average pattern (such as Minimalism and Color Field Painting) are more frequently correctly identified.

To better understand the role of each ordinal pattern in predicting styles, we evaluate the importance of all patterns by measuring their contribution to the overall accuracy of the model (see Machine learning artistic styles section). Figure 5C shows the importance measure for each ordinal pattern grouped into their 11 categories, as well as the top 10 patterns that most significantly contribute to the overall accuracy. There is no evident association between pattern types or overall prevalence of patterns (see also Fig. 1) and their importance in classifying artistic styles. Notably, the ranking of the most important patterns includes the pattern [0110], which belongs to the least prevalent type (K). Another remarkable finding is the disproportional importance of the type A pattern [0000], which alone accounts for 12% of the overall accuracy. It is equally intriguing to note that five type B patterns are the only patterns without identical values in the top 10 importance rank, the remaining five patterns are from types with identical values. Combined with the low classification performance in models trained with ordinal patterns in which identical values are randomly resolved, these findings further corroborate the significance of adequately addressing identical values when investigating ordinal patterns in images.

In the final aspect of our investigation, we explore the potential of quantifying the temporal evolution of art paintings by

examining changes in the prevalence of ordinal patterns. This approach builds on the findings of Sigaki et al. (23), who quantified the evolution of art paintings using the complexity-entropy plane. However, beyond previously discussed challenges related to the quality and resolution of digital reproductions available in Wikiart, it is crucial to recognize that tracking these patterns over time with this dataset is inherently limited due to its underlying biases. As noted in Refs. (25, 26, 43), the Wikiart collection largely consists of Western European artworks, reflecting the historical canon and thus providing a skewed representation of the global evolution of art. Moreover, it displays notable biases towards male artists (25), as well as towards 20th-century and post-World War II works, with gaps in coverage for the 18th century and early 19th century (26, 43). These biases restrict the diversity of perspectives captured and narrow the temporal scope of our analysis. Consequently, conclusions about art's evolution will only become more robust through the development of balanced, high-quality datasets that are yet to be realized.

With these limitations in mind, we group paintings in our dataset by composition date and calculate the average value of s within time intervals encompassing a comparable number of artworks. Figure S5A depicts the evolution of the average value of s , with different background colors delineating three major art historical periods: Medieval Art, the Renaissance, Neoclassicism, and Romanticism, which developed until the 1850s; Modern Art, initiated with the onset of Impressionism in the 1870s and marked by the emergence of avant-garde styles such as Cubism, Expressionism, and Surrealism in the early 20th century; and Contemporary/Postmodern Art, typically associated with the rise of Pop Art in the 1960s (63). Each period exhibits distinctive typical distances to the average ordinal pattern, corroborating the observation that artistic styles also display different values of s . Notably, the evolution of art paintings shows a general trend of increasing divergence from the average ordinal pattern. This divergence accelerated during the transition from Modernism to Postmodernism, a pattern also observed in the complexity-entropy plane of Sigaki et al. (23). In addition to the average value of s , we calculate the absolute sum of the coefficients of variation for each ordinal pattern within the same time intervals. Figure S5B indicates that the evolution of art, as depicted by changes in s , is neither smooth nor homogeneous, with substantial variations in the prevalence of ordinal patterns. This variability is particularly pronounced in the earliest period in our dataset and after the 1930s. While the high variability in the earliest period partially reflects the broad range of epochs it encompasses—aggregating artworks from diverse historical contexts—the pronounced variability during the transition from Modernism to Postmodernism aligns with interpretations of this shift as marked by stylistic disruptions rather than a gradual evolution, leading to a fragmented, diverse artistic landscape in which multiple styles and approaches coexist (26, 64, 65). We further examine individual changes in the prevalence of ordinal patterns by calculating their z -score probabilities relative to the average prevalence across all images, as shown in Fig. S5C, and the coefficients of variation of each ordinal pattern, as presented in Fig. S5D (on logarithmic scale). These results corroborate the trends observed when aggregating all patterns to estimate the values of s but also highlight some interesting shifts in the prevalence of ordinal patterns, primarily marked by an increase in the incidence of patterns A, D, and H and a decrease in type B in more recent works (despite early artworks also exhibiting positive z -scores for types D and H). However, once again, the often high

variability in the prevalence of ordinal patterns underscores that these shifts do not represent a smooth development.

Discussion and conclusions

We have conducted a comprehensive analysis of artworks using ordinal patterns extracted from two-by-two-pixel partitions of images. Initially, we have categorized the 75 possible ordinal patterns into 11 categories based on symmetry considerations and smoothness metrics. These categories were hierarchically organized based on the number of unique symbols and the smoothness of their patterns. Subsequently, each of the nearly 140 thousand images in our dataset was represented by a discrete probability distribution, reflecting the relative frequency of each ordinal pattern. By integrating expertise from the art domain, our method demonstrated the potential of using the distribution of ordinal patterns to facilitate quantitative analyses of various visual aspects of artworks, ranging from the dominance of vertical and horizontal structures to the prevalence of sharp edges and well-defined elements.

Our analysis identified a universal pattern in the occurrence of ordinal patterns across painting images, which led to our classification into 11 distinct pattern types. The average prevalence of these patterns within each group remains approximately constant but monotonically decreases as the roughness of the patterns increases among types sharing the same number of symbols. We have also observed that the incidence rank of pattern types reflects a nontrivial interplay between their smoothness and the decreased likelihood of finding pixels with identical intensities in two-by-two image partitions. Although a clear hierarchy exists in the occurrence of ordinal patterns in paintings, we have verified that patterns with lower prevalence exhibit significant variability compared to their mean probability of occurrence. This finding shows that tiny changes in pattern incidence can lead to substantial relative variations, prompting us to establish a standardized measure for the prevalence of ordinal patterns. Moreover, when summed absolutely, this measure provides a distance measure relative to the average pattern incidence.

We have used this standardized metric to define and rank the ordinal fingerprint of artistic styles based on their proximity to the overall incidence of ordinal patterns. The average incidence of ordinal patterns captures low-level features that appear to correlate with certain aesthetic attributes of styles. For instance, Divisionist and Pointillist compositions, which often depict textured visuals with indistinct edges, show a lower incidence of patterns with identical pixel intensities, whereas Light and Space artworks, characterized by geometric and well-defined shapes, display a higher occurrence of these patterns. Styles also exhibit notable variability in the prevalence of ordinal patterns, highlighting not only the limitations in representing artworks solely through two-by-two ordinal patterns but also the heterogeneity of our dataset, challenges in achieving reliable style annotation, and the inherent ambiguity in assigning a single style to artworks. Despite these complexities, our analysis demonstrated that the prevalence of ordinal patterns is predictive of artistic styles in paintings, and that by adequately accounting for patterns of equal pixel intensities, we have substantially enhanced the performance of style classification. Notably, patterns characterized by identical pixel intensities are among the most important features for style classification, with the type A pattern alone contributing to 12% of the overall accuracy. Although our overall accuracy of ~28% is modest compared to deep learning-based approaches (37, 40–43), it underscores that simple, interpretable, low-level

ordinal patterns provide meaningful insights into the features of artworks, particularly when patterns with identical pixel intensities are adequately considered.

Our investigation also explored the temporal changes of art paintings in our dataset, examining shifts in the prevalence of ordinal patterns and their deviation from the average pattern. The results revealed a general trend wherein art paintings increasingly diverge from the average pattern, punctuated by brief intervals in which this trend is reversed. This evolution is neither smooth nor homogeneous; instead, it is marked by substantial variations in the prevalence of ordinal patterns, particularly in artworks produced after the 1930s and during the transition from Modernism to Postmodernism. These findings suggest that the evolution of art does not follow a unilinear path. Rather, shifts in the prevalence and variability of ordinal pattern prevalence, particularly during transitional periods, appear to align with stylistic disruptions, resulting in a fragmented and diverse artistic landscape where multiple styles and approaches coexist (26, 64, 65).

Our work is not without limitations, most of which have been previously acknowledged; however, it is useful to summarize them here. First, our dataset, sourced from Wikiart, consists of a large yet heterogeneously compiled collection of digital reproductions, primarily reflecting Western paintings, which limits its generalizability and introduces variability in image quality and metadata accuracy. The diverse resolutions of artworks, especially larger pieces, further impact our analysis, as high-resolution images capture fine details that can influence the prevalence of certain ordinal patterns. Although we did not observe any clear bias in the association between image size, pixel density, and distance from the average ordinal pattern, caution remains essential, particularly in analyses of nuanced features in individual artworks. Additionally, style classification is inherently complex, with ambiguities in style attribution, variability in stylistic consistency across certain artistic movements, and limited representation within the dataset of non-Western and diverse gender perspectives. The temporal evolution of art is also challenging to track accurately due to dataset biases, as Wikiart predominantly features 20th-century Western works, which may distort broader historical interpretations. Lastly, we must acknowledge that the simplification involved in reducing the rich complexity of paintings to ordinal patterns at a single spatial scale inherently overlooks multiscale details essential to the artworks' esthetic and structural nuances. Future research, ideally utilizing high-resolution reproductions, could extend our use of ordinal patterns across multiple scales (66), similar to wavelet transforms, which have proven valuable in art investigations, including dating, artistic identification, authentication, image separation, cradle removal, and brushstroke analysis (28–31, 33).

Despite these limitations, our research demonstrates that, even amidst the visual complexity of art paintings, simple ordinal patterns emerging from two-by-two-pixel partitions encode critical information capable of suggesting artistic styles, fostering in-depth analysis of their esthetic features, and even revealing the dynamical behavior of art.

Methods

Data

The dataset comprises 137,364 digitized paintings from the WikiArt visual arts encyclopedia, as originally introduced by Sigaki et al. (23). It spans nearly a millennium (1031–2016) and includes works by 2,391 artists representing 154 artistic styles. For

the temporal analysis, 33,724 artworks lacking composition dates are excluded (Fig. S5). Images are stored in JPEG format with 24-bit RGB encoding. Color layers are converted to grayscale for analysis using the standard luminance transformation (67). See the [Supplementary material \(Section 1A\)](#) for more details.

Ordinal patterns without ties

We follow the seminal approach of Bandt and Pompe (52) to encode images using ordinal patterns. Images are sampled using sliding partitions; each partition is flattened, and its elements are replaced by their ranks to form an ordinal pattern. However, in contrast to our approach, identical values within a partition are handled by preserving their order of occurrence or adding a small noise term to resolve rank ties without altering other ordering relationships (alternative approaches have been proposed (68)). This procedure yields an ordinal probability distribution $P = \{p_i; i = 1, \dots, n\}$, where each p_i corresponds to the probability of finding one of the $n = (d_x d_y)! = 24$ possible ordinal patterns. Thus, our distinction of patterns with identical values introduces 51 additional ordinal patterns, raising the total to 75 possible patterns.

Using the standard ordinal probability distribution, we calculate the normalized Shannon entropy H (52), the statistical complexity (59, 66, 69–71) C , and Fisher information measure F . Furthermore, following Bandt and Wittfeld (53), the 24 ordinal patterns are classified into three types (I, II, and III) based on continuity, to then calculate the smoothness $\tau = q_1 - 1/3$ and the branching structure $\kappa = q_2 - q_3$, where q_1 , q_2 , and q_3 are the relative frequencies of types I, II, and III, respectively. The combination of H and C defines the complexity–entropy plane, the values of H and F yield the Fisher–Shannon plane, and the values of τ and κ determine the smoothness–structure plane. We implement all these ordinal methods using the Python module *ordpy* (72). See the [Supplementary material \(Sections 1B–E\)](#) for further details.

Machine learning artistic styles

We apply the extreme gradient boosting (XGBoost) (62) algorithm to classify the 20 most common styles in our dataset, with each style represented by at least 1,500 artworks. We partition the data in a stratified manner—80% for training and 20% for testing—and repeat this procedure across 10 independent realizations to compute mean and SD of accuracy. We investigate eight sets of ordinal features. Six sets use the Bandt and Pompe encoding (52), comprising the complexity–entropy, Fisher–Shannon, and smoothness–structure planes; the probabilities of the three pattern types; the probabilities of the 24 possible patterns without ties; and the corresponding probabilities obtained by resolving rank ties through the addition of small noise. The remaining two feature sets, original to our work, explicitly consider the occurrence of identical values within the image partitions: the probability associated with each of the 11 pattern types (q_A , q_B , q_C , ..., and q_K) and the probability of each of the 75 ordinal patterns. Finally, we assess the contribution of each of the 75 ordinal patterns using permutation feature importance (73–75) as implemented in scikit-learn (76). See the [Supplementary material \(Section 1F\)](#) for further details.

Low-dimensional projection of the ordinal pattern probabilities

We use uniform manifold approximation and projection (UMAP) (54) to reduce 75D ordinal pattern frequency vectors to two

dimensions. UMAP constructs a weighted graph (fuzzy simplicial complex) from a dissimilarity matrix and projects the data via a force-directed layout algorithm. Owing to inherent stochasticity, projections vary slightly between runs, rendering UMAP most useful for comparative visualization. We implement this approach using the Python package *umap* (77). See the [Supplementary material \(Section 1G\)](#) for further details.

Supplementary Material

[Supplementary material](#) is available at PNAS Nexus online.

Funding

This work is supported in part by funds from the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES), the Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq – grant 303533/2021-8), the Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET, Argentina), and the Slovenian Research Agency (grant nos. P1-0403 and N1-0232).

Author Contributions

M.M.T., A.A.B.P., L.Z., O.A.R., M.P., and H.V.R. contributed the conceptualization, data curation, methodology, software, investigation, validation, formal analysis, visualization, writing—original draft, and writing—review & editing.

Data Availability

The data that support the findings of this study are available as a [File S1](#).

References

- Birkhoff GD. *Aesthetic measure*. Vol. 38 Harvard University Press Cambridge, 1933.
- Beardsley MC. *Aesthetics from classical Greece to the present* University of Alabama Press, 2015.
- Taylor RP, Micolich AP, Jonas D. 1999. Fractal analysis of Pollock's drip paintings. *Nature*. 399(6735):422.
- Jones-Smith K, Mathur H. 2006. Fractal analysis: revisiting Pollock's drip paintings. *Nature*. 444(7119):E9–E10.
- Taylor RP, Micolich AP, Jonas D. 2006. Fractal analysis: revisiting Pollock's drip paintings (reply). *Nature*. 444(7119):E10.
- Taylor RP, et al. 2007. Authenticating Pollock paintings using fractal geometry. *Pattern Recognit Lett*. 28(6):695–702.
- Pedram P, Jafari GR. 2008. Mona Lisa: the stochastic view and fractality in color space. *Int J Mod Phys C*. 19(6):855–866.
- Jones-Smith K, Mathur H, Krauss LM. 2009. Drip paintings and fractal analysis. *Phys Rev E*. 79(4):046111.
- Hughes JM, Graham DJ, Rockmore DN. 2010. Quantification of artistic style through sparse coding analysis in the drawings of Pieter Bruegel the Elder. *Proc Natl Acad Sci U S A*. 107(4):1279–1283.
- Shamir L. 2012. Computer analysis reveals similarities between the artistic styles of Van Gogh and Pollock. *Leonardo*. 45(2):149–154.
- De la Calleja EM, Cervantes F, De la Calleja J. 2016. Order-fractal transitions in abstract paintings. *Ann Phys (N Y)*. 371(5):313–322.
- Alvarez-Ramirez J, Ibarra-Valdez C, Rodriguez E. 2016. Fractal analysis of Jackson Pollock's painting evolution. *Chaos Solitons Fractals*. 83:97–104.

- 13 Calleja EM, Zenit R. 2017. Topological invariants can be used to quantify complexity in abstract paintings. *Knowl Based Syst.* 126:48–55.
- 14 Castrejon-Pita JR, Castrejón-Pita AA, Sarmiento-Galán A, Castrejón-García R. 2003. Nasca lines: a mystery wrapped in an enigma. *Chaos.* 13(3):836–838.
- 15 Koch M, Denzler J, Redies C. 2010. $1/f^2$ characteristics and isotropy in the Fourier power spectra of visual art, cartoons, comics, mangas, and different categories of photographs. *PLoS One.* 5(8): e12268.
- 16 Montagner C, Linhares JMM, Vilarigues M, Nascimento SMC. 2016. Statistics of colors in paintings and natural scenes. *J Opt Soc Am A.* 33(3):A170–A177.
- 17 Perc M. 2020. Beauty in artistic expressions through the eyes of networks and physics. *J R Soc Interface.* 17(164):20190686.
- 18 Manovich L. 2015. Data science and digital art history. *Int J Digit Art Hist.* 1:13–35.
- 19 Yazdani M, Chow J, Manovich L. 2017. Quantifying the development of user-generated art during 2001–2010. *PLoS One.* 12(8): e0175350.
- 20 Ushizima D, Manovich L, Margolis T, Douglas J. 2021. Cultural analytics of large datasets from Flickr. *Proc Int AAAI Conf Web Soc Media.* 6(4):30–34.
- 21 Kim D, Son S-W, Jeong H. 2014. Large-scale quantitative analysis of painting arts. *Sci Rep.* 4(1):7370.
- 22 Lee B, Kim D, Sun S, Jeong H, Park J. 2018. Heterogeneity in chromatic distance in images and characterization of massive painting data set. *PLoS One.* 13(9):e0204430.
- 23 Sigaki HYD, Perc M, Ribeiro HV. 2018. History of art paintings through the lens of entropy and complexity. *Proc Natl Acad Sci U S A.* 115(37):E8585–E8594.
- 24 Valensise CM, et al. 2021. Entropy and complexity unveil the landscape of memes evolution. *Sci Rep.* 11(1):20022.
- 25 Lee B, et al. 2020. Dissecting landscape art history with information theory. *Proc Natl Acad Sci U S A.* 117(43):26580–26590.
- 26 Karjus A, Solà MC, Ohm T, Ahnert SE, Schich M. 2023. Compression ensembles quantify aesthetic complexity and the evolution of visual art. *EPJ Data Sci.* 12(1):21.
- 27 Lee K, Park J, Goree S, Crandall D, Ahn Y-Y. 2024. Social signals predict contemporary art prices better than visual features, particularly in emerging markets. *Sci Rep.* 14(1):11615.
- 28 Jafarpour S, et al. Stylistic analysis of paintings using wavelets and machine learning. In: 2009 17th European Signal Processing Conference. IEEE, 2009. p. 1220–1224.
- 29 Johnson CR, et al. 2008. Image processing for artist identification. *IEEE Signal Process Mag.* 25(4):37–48.
- 30 Polatkan G, Jafarpour S, Brasoveanu A, Hughes SM, Daubechies I. Detection of forgery in paintings using supervised learning. In: 2009 16th IEEE International Conference on Image Processing (ICIP). IEEE, 2009. p. 2921–2924.
- 31 Yin R, et al. Digital cradle removal in X-ray images of art paintings. In: 2014 IEEE International Conference on Image Processing (ICIP). IEEE, 2014. p. 4299–4303.
- 32 Cornelis B, et al. 2013. Crack detection and inpainting for virtual restoration of paintings: the case of the Ghent Altarpiece. *Signal Process.* 93(3):605–619.
- 33 Deligiannis N, Mota JFC, Cornelis B, Rodrigues MRD, Daubechies I. 2016. Multi-modal dictionary learning for image separation with application in art investigation. *IEEE Trans Image Process.* 26(2):751–764.
- 34 Sabetsarvestani Z, Sober B, Higgitt C, Daubechies I, Rodrigues MRD. 2019. Artificial intelligence for art investigation: meeting the challenge of separating X-ray images of the Ghent Altarpiece. *Sci Adv.* 5(8):eaaw7416.
- 35 Shamir L, Macura T, Orlov N, Eckley DM, Goldberg IG. 2010. Impressionism, expressionism, surrealism: automated recognition of painters and schools of art. *ACM Trans Appl Perception (TAP).* 7(2):1–17.
- 36 Arora RS, Elgammal A. Towards automated classification of fine-art painting style: a comparative study. In: *Proceedings of the 21st International Conference on Pattern Recognition (ICPR2012)*. IEEE, 2012. p. 3541–3544.
- 37 Wu T, et al. Painting analysis using wavelets and probabilistic topic models. In: 2013 IEEE International Conference on Image Processing. IEEE, 2013. p. 3264–3268.
- 38 Krizhevsky A, Sutskever I, Hinton GE. Imagenet classification with deep convolutional neural networks. In: Pereira F, Burges CJ, Bottou L, Weinberger KQ, editors. *Advances in neural information processing systems*. Vol. 25. Curran Associates, Inc., 2012.
- 39 Goodfellow I, Bengio Y, Courville A. *Deep learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- 40 Zujovic J, Gandy L, Friedman S, Pardo B, Pappas TN. Classifying paintings by artistic genre: an analysis of features & classifiers. In: *IEEE International Workshop on Multimedia Signal Processing, 2009. MMSP'09*. IEEE, 2009. p. 1–5.
- 41 Tan WR, Chan CS, Aguirre HE, Tanaka K. Ceci n'est pas une pipe: a deep convolutional network for fine-art paintings classification. In: 2016 IEEE International Conference on Image Processing (ICIP). IEEE, 2016. p. 3703–3707.
- 42 Mao H, Cheung M, She J. Deepart: learning joint representations of visual arts. In: *Proceedings of the 25th ACM International Conference on Multimedia*. ACM, 2017. p. 1183–1191.
- 43 Elgammal A, Liu B, Kim D, Elhoseiny M, Mazzone M. The shape of art history in the eyes of the machine. In: *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence, AAAI'18/IAAI'18/EAAI'18*. AAAI Press, 2018.
- 44 Sandoval C, Pirogova E, Lech M. 2019. Two-stage deep learning approach to the classification of fine-art paintings. *IEEE Access.* 7:41770–41781.
- 45 Hekkert P. 2006. Design aesthetics: principles of pleasure in design. *Psychol Sci.* 48(2):157.
- 46 Hubbard TL. Aesthetics and preferences in spatial and scene composition. In: *Spatial biases in perception and cognition*. Cambridge University Press, 2018. p. 222–240.
- 47 Warnke M. 1989. On Heinrich Wölfflin. *Representations.* 27: 172–187.
- 48 Wölfflin H. *Principles of art history: the problem of the development of style in later art* Dover, Mineola, 1950.
- 49 Darst DH. 1983. Renaissance symmetry Baroque symmetry and the sciences. *Diogenes.* 31(123):69–90.
- 50 Sypher W. *Four stages of renaissance style* Peter Smith Pub Inc, 1978.
- 51 Davis W. *A general theory of visual culture* Princeton University Press, 2018.
- 52 Bandt C, Pompe B. 2002. Permutation entropy: a natural complexity measure for time series. *Phys Rev Lett.* 88(17):174102.
- 53 Bandt C, Wittfeld K. 2023. Two new parameters for the ordinal analysis of images. *Chaos.* 33(4):043124.
- 54 McInnes L, Healy J, Melville J. 2018. UMAP: uniform manifold approximation and projection for dimension reduction, arXiv, arXiv:1802.03426, <https://doi.org/10.48550/arXiv.1802.03426>, preprint: not peer reviewed.

- 55 Lause J, Berens P, Kobak D. 2024. The art of seeing the elephant in the room: 2D embeddings of single-cell data do make sense. *PLoS Comput Biol.* 20(10):e1012403.
- 56 Chari T, Pachter L. 2023. The specious art of single-cell genomics. *PLoS Comput Biol.* 19(8):e1011288.
- 57 Gascoyne D. *A short survey of surrealism* Cass, London, 1970.
- 58 Hopkins D. *Dada and surrealism: a very short introduction* Oxford University Press, Oxford, 2004.
- 59 Rosso OA, Larrondo HA, Martin MT, Plastino A, Fuentes MA. 2007. Distinguishing noise from chaos. *Phys Rev Lett.* 99(15):154102.
- 60 Ribeiro HV, Zunino L, Lenzi EK, Santoro PA, Mendes RS. 2012. Complexity-entropy causality plane as a complexity measure for two-dimensional patterns. *PLoS One.* 7:1–9.
- 61 Olivares F, Plastino A, Rosso OA. 2012. Contrasting chaos with noise via local versus global information quantifiers. *Phys Lett A.* 376(19):1577–1583.
- 62 Chen T, Guestrin C. XGBoost: a scalable tree boosting system. In: *KDD'16: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.* New York: ACM, 2016. p. 785–794.
- 63 Danto AC, Goehr L. *After the end of art: contemporary art and the pale of history.* Vol. 197 Princeton University Press, Princeton, 1997.
- 64 Foster H, Krauss RE, Bois Y-A, Buchloh BHD, Joslelit D. *Art since 1900: modernism, antimodernism, postmodernism* Thames & Hudson, 2016.
- 65 Mesoudi A, Thornton A. 2018. What is cumulative cultural evolution? *Proc Biol Sci.* 285(1880):20180712.
- 66 Zunino L, Ribeiro HV. 2016. Discriminating image textures with the multiscale two-dimensional complexity-entropy causality plane. *Chaos Solitons Fractals.* 91(5):679–688.
- 67 Van der Walt S, et al. 6 2014. scikit-image: image processing in Python. *PeerJ.* 2(2):e453.
- 68 Bian C, Qin C, Ma QDY, Shen Q. 2012. Modified permutation-entropy analysis of heartbeat dynamics. *Phys Rev E.* 85(2):021906.
- 69 Lopez-Ruiz R, Mancini HL, Calbet X. 1995. A statistical measure of complexity. *Phys Lett A.* 209(5–6):321–326.
- 70 Lamberti PW, Martin MT, Plastino A, Rosso OA. 2004. Intensive entropic non-triviality measure. *Physica A.* 334(1):119–131.
- 71 Martin MT, Plastino A, Rosso OA. 2006. Generalized statistical complexity measures: geometrical and analytical properties. *Physica A.* 369(2):439–462.
- 72 Pessa AAB, Ribeiro HV. 2021. ordpy: a Python package for data analysis with permutation entropy and ordinal network methods. *Chaos.* 31(6):063110.
- 73 Breiman L. 2001. Random forests. *Mach Learn.* 45(1):5–32.
- 74 Fisher A, Rudin C, Dominici F. 2019. All models are wrong, but many are useful: learning a variable's importance by studying an entire class of prediction models simultaneously. *J Mach Learn Res.* 20(177):1–81.
- 75 Molnar C. *Interpretable machine learning.* 2nd ed Lulu, 2022.
- 76 Pedregosa F, et al. 2011. Scikit-learn: machine learning in Python. *J Mach Learn Res.* 12:2825–2830.
- 77 McInnes L, Healy J, Saul N, Grossberger L. 2018. UMAP: uniform manifold approximation and projection. *J Open Source Softw.* 3(29):861.