



OPEN

## Social signals predict contemporary art prices better than visual features, particularly in emerging markets

Kangsan Lee<sup>1,5</sup>, Jaehyuk Park<sup>2,5</sup>, Sam Goree<sup>3,5</sup>, David Crandall<sup>4</sup> & Yong-Yeol Ahn<sup>4</sup>

What determines the price of an artwork? This article leverages a comprehensive and novel dataset on art auctions of contemporary artists to examine the impact of social and visual features on the valuation of artworks across global markets. Our findings indicate that social signals allow us to predict the price of artwork exceptionally well, even approaching the professionals' prediction accuracy, while the visual features play a marginal role. This pattern is especially pronounced in emerging markets, supporting the idea that social signals become more critical when it is more difficult to assess the quality. These results strongly support that the value of artwork is largely shaped by social factors, particularly in emerging markets where a stronger preference for "buying an artist" than "buying an artwork." Additionally, our study shows that it is possible to boost experts' performance, highlighting the potential benefits of human-machine models in uncertain or rapidly changing markets, where expert knowledge is limited.

Imagine yourself sitting at an obscure auction, bidding on an old portrait attributed to someone obscure. How much would you be willing to bid? Now, consider the same painting being auctioned at the Christie's in New York, which is allegedly by Leonardo Da Vinci himself. How much would you be willing to pay? In 1958, the painting "Salvator Mundi", previously attributed to someone in Da Vinci's studio, was sold for a mere \$60. However, after it was re-attributed to Da Vinci himself, the *same painting* was sold for \$450 million at Christie's New York in 2017, becoming the most expensive painting ever sold<sup>1</sup>. This extreme example highlights how an exceptionally renowned artist's reputation can dwarf the innate qualities of the artwork. This tension between innate beauty and external fame has been two contrasting frames to examining arts. Is the value of artworks primarily based on the artist's fame and reputation, or are such examples more exceptional than typical? To what extent can we assess the value of a piece without any knowledge of its intrinsic quality?

Assessing the value of unique cultural products, such as artwork, is notoriously difficult. As each work of art possesses unique and subjective aesthetic qualities, creating a standardized measure of its quality would be almost impossible. However, cultural valuation is not only an individual's choice of artistic characteristics but also a highly socialized action<sup>2-5</sup>.

*Social constructionism* emphasizes the significance of social signals from the creators and distributors of art in determining its value<sup>4,6-8</sup>, rather than focusing on specific characteristics of the artwork itself. This can result in a weak connection between the quality of cultural products and their success in cultural markets<sup>9</sup> and the important role of social networks in success<sup>10</sup>.

This is particularly evident in modern abstract or contemporary art, where there is less emphasis on figurative sophistication and a relatively shorter period of time for a consensus to emerge<sup>11-13</sup>. As a result, different social signals from both the artist and the market are considered significant determinants in the valuation process within the global art world<sup>2,3,14-17</sup>. In this perspective, the valuation may vary across markets based on the audiences' level of understanding and shared values.

However, some argue that artistic quality has an unchanging value, as seen in *formalism*, a theory in the study of art and aesthetics. Formalist theories<sup>18,19</sup> suggest that the value of an artwork is a result of its *artistic form*—the

<sup>1</sup>Division of Social Science, New York University Abu Dhabi, Abu Dhabi, UAE. <sup>2</sup>School of Public Policy and Management, Korea Development Institute, Sejong-si, Republic of Korea. <sup>3</sup>Department of Computer Science, Stonehill College, Easton, MA 02357, USA. <sup>4</sup>Luddy School of Informatics, Computing, and Engineering, Indiana University, Bloomington, IN 47408, USA. <sup>5</sup>These authors contributed equally: Kangsan Lee, Jaehyuk Park and Sam Goree. ✉email: ks.lee@nyu.edu

formal visual qualities of the artwork itself, such as color, line, composition, and shape, over its social, historical, or cultural context. Thus, the valuation process should remain consistent across different markets since the forms do not change<sup>20,21</sup>. Formalists believe that an artwork's ability to express and convey meaning is the most essential factor in determining its value<sup>5</sup>. However, there has been no attempt to systematically quantify and compare the sizes of these two effects—social signals and artistic forms—on the market value of artworks.

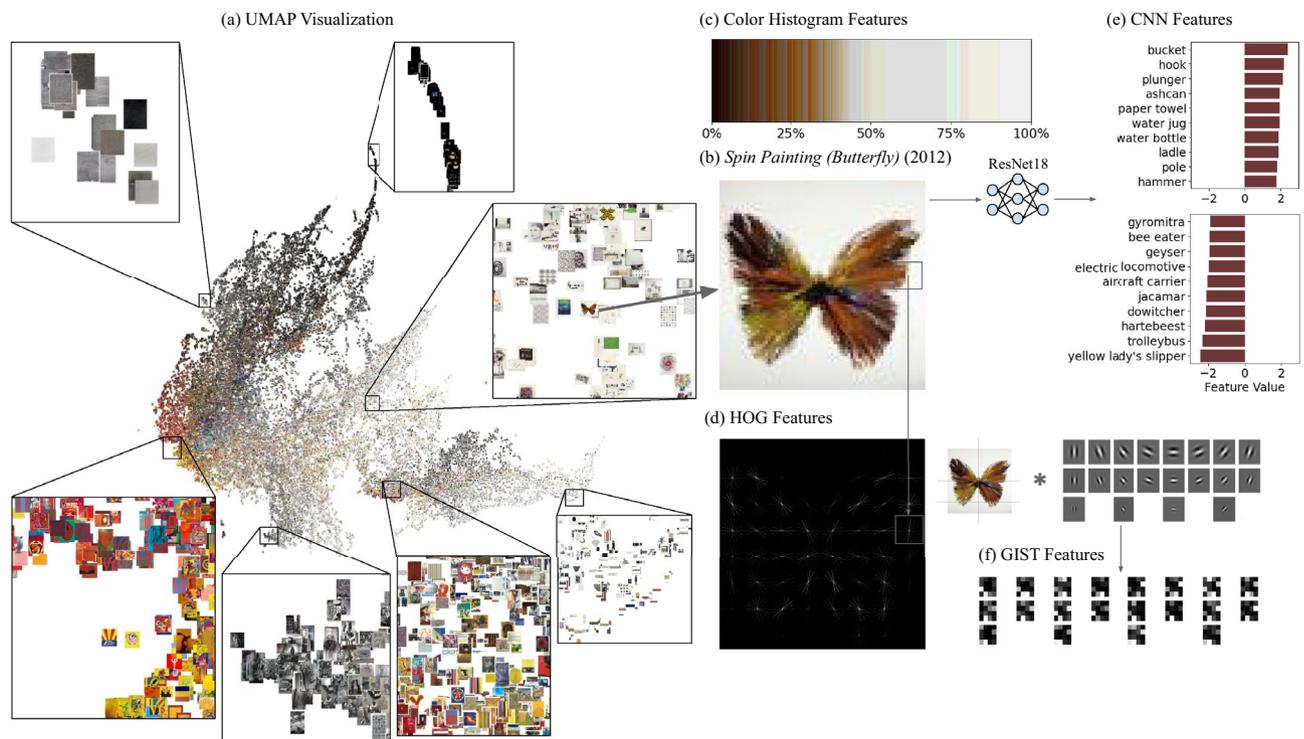
This study aims to revisit and expand upon classical debates in valuation by quantitatively comparing the effectiveness of predicting market values based on the *visual* characteristics of the artwork versus *social* (*non-visual*) information about the artist and market. In addition, we further examine whether the impact of visual and social factors varies across established and emerging markets. Our aim is to determine whether social conditions, such as diverse audiences, influence the impact of social factors on cultural valuation, or if the form of artwork provides a universal experience to audiences.

Here, we use a novel dataset of 34,200 auction sales records, including images, artists' attributes, and market information, encompassing 590 living contemporary artists spanning 17 years (1996–2012) across 23 countries. This comprehensive approach goes beyond the findings of previous studies on cultural valuations and its elusive quality metrics<sup>10,22</sup> by analyzing the most systematically-collected dataset of contemporary arts and teasing out the contribution from visual and social features. We employ two machine learning models in predicting auction prices: one that utilizes visual features and the other that utilizes socioeconomic features. Then, we compare the contribution of visual versus social factors between established and emerging markets. The visual features are extracted from publicly available artwork thumbnails, while publicly available metadata for artists and market features are used for social features. We divide the transactions based on auction locations and select established markets, such as the USA, UK, France, and Germany, based on their historical significance and market size, and emerging markets, such as China, the Middle East, Brazil, Korea, and others. We concentrate on living contemporary artists and their artworks as they are more comparable to each other, but their shorter history compared to the works of old masters presents a greater challenge in their valuation<sup>13,23–25</sup>.

## Results

### Buying an artwork: computer vision analysis of visual features

We first construct an XGBoost model predicting the prices of artwork only utilizing the visual features extracted from the artworks' images. To extract the visual features, we apply multiple algorithms that are widely used in computer vision and pattern recognition fields, such as GIST<sup>26</sup>, HOG<sup>27</sup>, color histograms<sup>28</sup>, and deep learning-based features from a model trained on the ImageNet dataset<sup>29–31</sup>, while adding hand-tuned measures of colorfulness and complexity (presented in Fig. 1 and see “Materials and methods” for details). Although these features



**Figure 1.** Our visual feature extraction process. (a) A UMAP dimensionality reduction of our total feature space, with several regions at high resolution showing organization according to visual characteristics. (b) One example, *Spin Painting (Butterfly)* (2009). (c) An area plot of the color histogram features. (d) HOG features, shown by the lightness of gradient lines in each region. (e) Bar graph of CNN features, which correspond to ImageNet classes. (f) Heatmap of GIST features for each Gabor filter shown above.

capture various characteristics of the visual forms that humans rely on to recognize objects and scenes, the total number of visual features generated by all algorithms is excessively large and redundant. Hence, we employ principal component analysis (PCA) to reduce the feature space for visual features. Subsequently, we utilize the XGBoost algorithm with these reduced 100 features (see “Materials and methods” for details) to predict the price of the artwork<sup>32</sup>.

As shown in Fig. 1, our visual features successfully capture the fundamental characteristics shared among artworks. The 2-dimensional representation of our dataset’s artworks using UMAP (Fig. 1a) provides insight into the contemporary art landscape by showcasing visual similarities, ranging from color composition to structural composition. Moreover, we test to determine whether the selected visual features could predict other characteristics of the artwork or artist, such as genre or gender. Employing the XGBoost classifier, we discovered that the visual features exhibited the ability to predict the genre of artwork; however, they proved to be insignificant in predicting other artist-level characteristics in contemporary arts (Table 2).

While the visual features extracted by our algorithm capture various aesthetic aspects of artworks, they have a limited impact on predicting the price of artworks. The results of our model, presented in Table 1 are measured in terms of  $R^2$  value, which represents the variance of the estimated price explained by visual features. Our initial model, focusing solely on visual features, achieves the  $R^2$  of 0.055, indicating a marginal level of explanatory power in determining the price of artworks. Although this model outperformed the mean prediction of the training set, it fell short by only explaining a mere 5.5% of the price variation.

Our results suggest that, as one might expect, there may not be universal visual features that signal the “quality” or “value” of contemporary artwork, especially ones that can be captured by figurative distinctions from low-level visual feature extraction techniques.

### Buying an artist: metadata analysis of artists and markets

What would happen if we only use socioeconomic metadata, such as artist-related and market-related data, without any visual information about the artworks themselves? This approach—predicting the auction prices without any artwork-level features, especially visual ones—estimates to what extent the valuation is a social construct. Our features used in this model have mostly been tested in previous studies and they include thirty artist-level features such as age, gender, number of exhibitions, private/public acquisitions, and previous price levels, as well as eight market-level features, such as market size, art market growth, and the reputation of auction houses where each auction occurred (see “Materials and methods” for details).

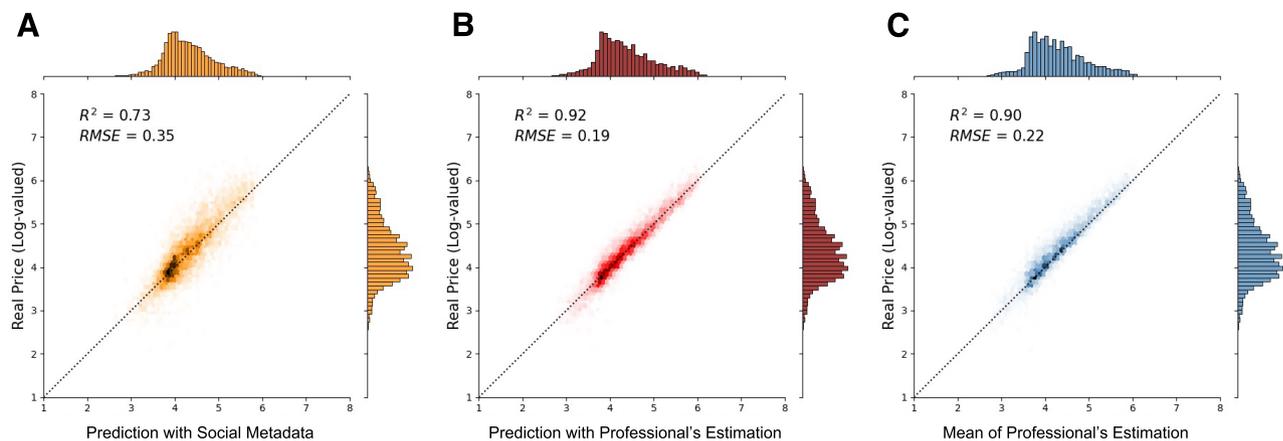
As a baseline, we first examine the estimated price ranges provided by art professionals in auction houses, taking into account all relevant factors including visual and social factors. These estimates are determined by auction houses that have exclusive access to extensive information about artists, artworks, and the markets and collectors<sup>8,33</sup>. The estimation provided by these professionals is a significant market signal because they *set expectations* about the value of an artwork. As a result, professional estimates may act as a self-fulfilling prophecy, particularly given the high valuation uncertainty<sup>9</sup>. Our findings indicate that these estimates are indeed highly predictive of the actual sales price, with about 90% of the variance explained (Fig. 2C).

Model	Price (test set $R^2$ )
Baseline (mean regressor)	– 0.002
All visual features	0.055
Color features	0.056
CNN features	0.009
HOG features	0.029
GIST features	0.048
Colorfulness and complexity	0.012

**Table 1.** Test set performance for computer vision analysis methods in predicting the price of an artwork, compared to a baseline model which predicts the mean of the training set. The visual features proved to be informative, with color histogram and GIST features being particularly useful. However, it was found that these two features were largely redundant with each other.

Feature	Visual accuracy	“Dummy” accuracy
Gender	0.861842	0.861988
Education	0.809503	0.809941
Elite school	0.733626	0.733771
Genre	0.472076	0.321053

**Table 2.** Prediction accuracy of a model based on visual features for several metadata characteristics. Each prediction is compared to a “dummy” classifier which always predicts the majority class (e.g. gender male, tier 1, country USA). Visual features are unhelpful to predict metadata characteristics other than genre. Please see Supplementary Materials for AUC and F1 scores.



**Figure 2.** A machine learning model can predict the price of contemporary artwork without any content information and can improve upon the prediction of the auction houses. We show the correlation between the actual auction prices and the estimated prices with  $R^2$  and Root Mean Square Error (RMSE) by (A) XGBoost model without professionals' estimates, (B) XGBoost model with professionals' estimates as an added feature, and (C) Arithmetic Mean of Professional's Minimum and Maximum Estimates (base-10 logarithmic scale).

In contrast to the results from our visual feature-based prediction (about 5.5% of variance explained), the XGBoost model that only uses social metadata surprisingly performs better (Fig. 2A), explaining about 73% of the variance. While this is still lower than the auction house's estimates, it significantly outperforms existing models that typically range between 50 and 52%<sup>23,24</sup> by a wide margin. Even considering previous studies highlighting the social construction of cultural valuation<sup>6,7,11</sup>, it is noteworthy that explaining 73% of the variation solely based on information about artists and markets is surprisingly high. One way to look at this result is that, because the model only uses information that exists prior to the creation of a given artwork, the model can predict 73% of the price variation *even before it is created*.

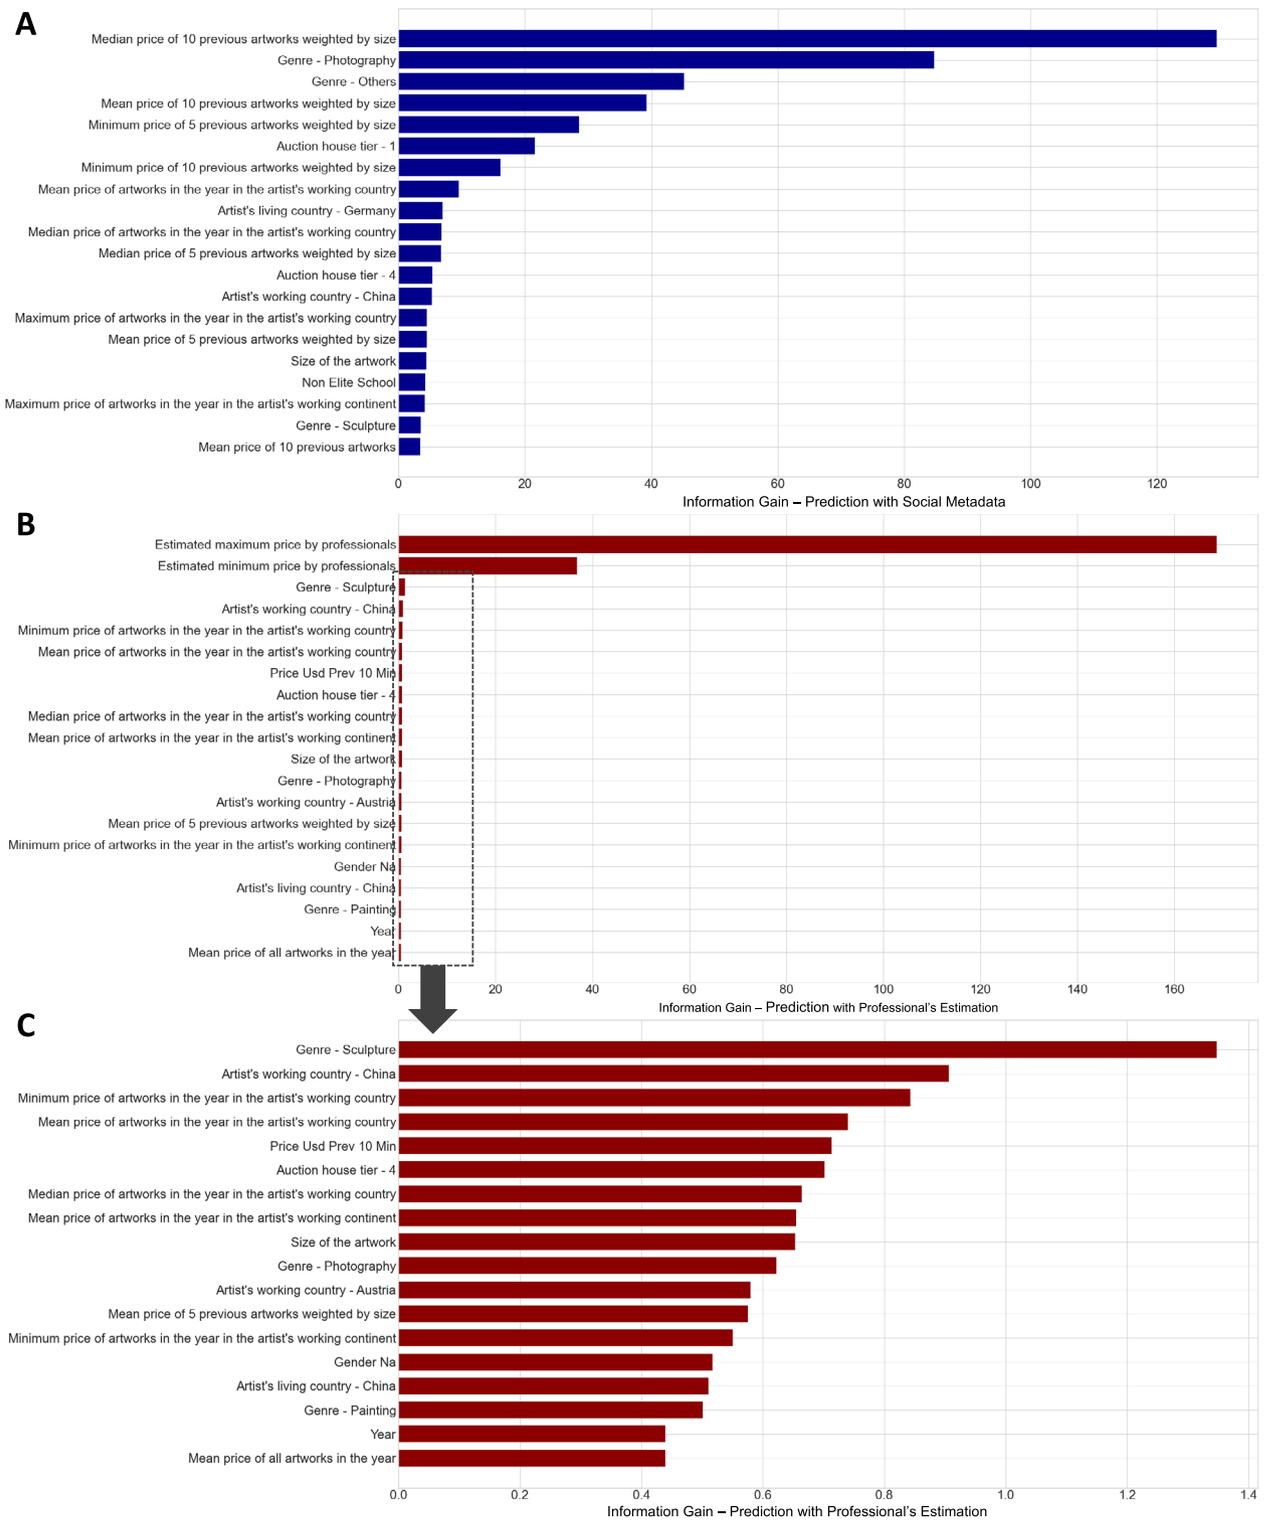
Furthermore, when we add the professionals' estimates as a feature in our model along with the metadata, the model could improve the professionals' estimates by explaining about 92% of the variance (Fig. 2B). This implies that while professionals' estimates are highly informative and have a self-fulfilling effect, there is still scope for improvement through machine learning. It also suggests that professionals' estimates can be further improved by leveraging only *publicly available information*.

Then, how does our model predict the price so well? The top 20 features (see Fig. 3A) with the largest information gain in the stand-alone model without professional estimation align well with the three factors of the social construction theory of valuation. First, the most crucial set of features is the *past*—price levels of previous artworks sold by the same artist while controlling for the size and genres of artworks. For example, the average of the past 10 works is approximately 50% correlated with the price, and 49% correlated with the mean expert's estimate (yielding an R-squared value of 0.25 compared to our final model's 0.73). This is a higher rate than other similar features, such as the average of 10 random previous works which is approximately 44% correlated with both the price and the expert assessments. This finding provides strong evidence of the Matthew effect<sup>34,35</sup> that the past market reputation of the artist and their existing recognition are critical in determining the price of new transactions<sup>36,37</sup>. As the contemporary art world increasingly relies on numeric forms of valuation as a result of commensuration and financialization<sup>17,38</sup>, previous price levels are commonly used to assess an artist's relative status and make comparisons and reinforce the Matthew effect<sup>2,39</sup>. The significance of previous price levels in Fig. 3A, as captured by the mean and median prices of artworks by the same artist, underscores how auctioneers use past pricing patterns to estimate price ranges<sup>17</sup>.

Second, the machine learning model highlights the significance of the auction venue, specifically the tier of the auction house. This finding aligns with Bourdieu's theory of cultural capital and the conversion of cultural value into economic capital<sup>7,40</sup>. According to this perspective, collectors of contemporary art are well aware of the high-status and influential auction houses, which are perceived as a significant pricing signal<sup>13,17</sup>. Each auction house has its bidding ratios and a different price of return on similar artworks<sup>41</sup>, leading artists to compete for inclusion in the top auction houses in major cities. Such social hierarchy and prestige of auction houses can be self-perpetuating and can influence the valuation. The importance of venues in artistic performance is also evident in primary markets, as exhibition networks project the social and artistic reputation of artists<sup>10</sup>.

Finally, the remaining features in Fig. 3A and the most important features in Fig. 3C, which are the rest of the features except the professional price estimations in Fig. 3B, are more closely related to *where* the auction occurred, reflecting the puzzle (and complexity) posed by the recent expansion of the global art market. The contemporary art market is currently experiencing unprecedented and rapid expansion in both Asian and global markets<sup>17,25</sup>. Tier 4 art auctions—outside of major auctions in the US or Europe as well as those in China—draw attention to the significance of transaction locations. This suggests that the sales location is an important factor that cannot be explained by professional estimates alone. This highlights the emergence of new markets in the global art landscape where the new rules are still being shaped (Fig. 3C).

The rapid global expansion of the art market and the emergence of new art markets (i.e., China and the Middle East) have created greater interconnectivity across borders<sup>42</sup>, but have also presented challenges to the traditional



**Figure 3.** Top 20 features providing the largest information gain in estimating the auction price of an artwork. XGBoost model with social metadata only (A), with social metadata and professionals' estimates overall (B), and the magnified version of the second model (C).

art world in terms of valuation and consumption<sup>43</sup>. Recent studies highlighted not only the ongoing process of conventional gatekeeping in new markets by reproducing the existing norms and conventions of valuation but also the traditional professionals' lack of understanding about these new markets<sup>17,44,45</sup>.

### Between established and emerging art markets

Building on our finding of the significance of auction locations in predicting artwork prices, we conduct another test to evaluate the predictability of our models, both visual and metadata, and professional estimation across different markets. We categorized the auction locations into two groups: established and emerging markets. Established markets comprise the primary art hubs, such as the United States, the United Kingdom, France, and Germany, due to their historical importance and the size of their art markets. Emerging markets include all other countries, except those in Europe, such as China, Brazil, Korea, and Qatar.

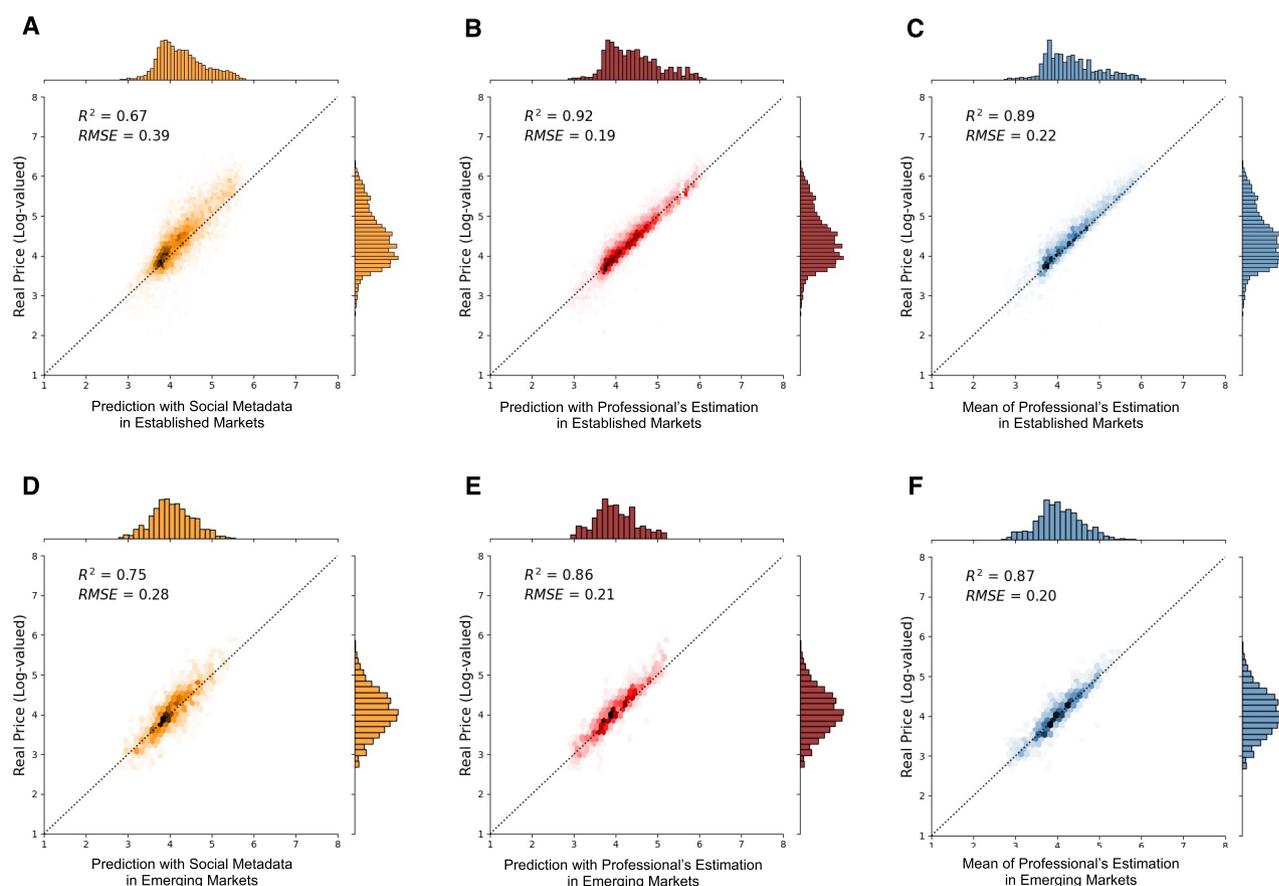
We first expect that there will be less variation in the predictions of visual features across markets, given the consistent visual quality. For social features, however, we anticipate greater variation between established and emerging markets, as distinct audiences across different countries may rely on shared social signals to varying degrees. In addition, we aim to determine whether our model can more accurately predict price patterns in emerging markets, despite their high uncertainty. This depends on whether the socially constructed valuation process is stronger in these new markets, as previous studies have suggested<sup>16,46</sup>.

Initially, we verify that price predictions based on visual forms continue to be marginally significant across markets, thereby explaining the marginal but consistent interpretation of artistic forms across various markets as presented in Table 3. The difference in prediction for visual features between established and emerging markets is negligible compared to that of metadata features, which show stronger but different predictive power.

Figure 4 demonstrates the predicted variance comparisons across three sets of models, the XGBoost model with social metadata only, the XGBoost model incorporating professional estimations as additional features, and the model comprising solely of professionals' estimations. Notably, these models were tested on separate samples, established markets and emerging markets.

The results show that while the professional estimation predicts better in established markets, our machine-learning model offers increased explanatory power in emerging markets. This result supports the social constructionism finding in the previous tests.

The less effectiveness of professional estimation in emerging markets is attributed to the slower acquisition of expertise by professionals. In emerging markets, this process is not yet fully established, owing to the



**Figure 4.** Accuracy comparisons of our XGBoost models without professionals' estimations (A,D), our models with professionals' estimations (B,E), and professionals' estimations (C,F) between established markets (USA, UK, France, and Germany) and emerging markets (otherwise). While professionals predict the artworks sold in established markets more accurately than those in emerging markets ( $R^2 = 0.89$  vs.  $0.87$ ), our model without professionals' estimations predicts the artworks sold in emerging markets better than those in established markets ( $R^2 = 0.75$  vs.  $0.67$ ). Together with professional estimation, the result shows a better prediction in the established market.

Market	Model	Price (test set $R^2$ )
Established	Baseline (mean regressor)	- 0.002
	Visual features	0.053
	Metadata (no prof. est.)	0.667
	Metadata (with prof. est.)	0.916
Emerging	Baseline (mean regressor)	- 0.002
	Visual features	0.056
	Metadata (no prof. est.)	0.750
	Metadata (with prof. est.)	0.859

**Table 3.** Performance of visual and metadata features split by established and emerging markets. As we anticipated, the lack of significance of visual features remains consistent across markets.

distinctiveness of new audiences and the level of market maturity. For example, during the Contemporary Spring auction in 2015, Shanghai Christie's auctioneers noted that "Many people here are fairly new to art auctions and we still need to learn about them and create a relationship to understand what they want and to return" (Author interview, April 24, 2015). Professionals tend to build relationships with local collectors and artists over an extended period, facilitating a shared understanding of the value of artworks<sup>11</sup>. While professionals possess both formal and tacit knowledge and experience of established market conditions<sup>47</sup>, they are still in the process of familiarizing themselves with new audiences and engaging with them to attract more collectors in emerging markets<sup>17</sup>.

In addition, the results from emerging markets exhibit that socially constructed consumption is a major driver of its variation, which leads to a strong preference for buying artists. Collectors in these markets have relatively limited experience and knowledge about the subtle variation across the artworks of an artist, leading to a greater dependence on social signals as a means of coping with market uncertainty. This is due to the fact that professional guidance regarding artwork-based distinctions is not yet fully elaborated and shared in emerging markets.

In other words, established markets exhibit more variance, which cannot be solely explained by publicly available information about artists and markets. Through their extensive experience and close relationships with professionals, collectors in established markets not only comprehend differences in the prestige among artists but also gain access to the knowledge of subtle distinctions across artworks, even those created by the same artist. They often use this knowledge to express their distinctive cultural capital and distinguish themselves from other collectors<sup>47</sup>. Having said that, established markets are better equipped to appreciate and acknowledge the distinctiveness of artworks, and to recognize differences across artworks of the same artist aligned with professional valuations. In contrast, emerging markets place greater emphasis on social prestige at both the artist and market levels, which can be captured more efficiently, faster, and effectively using machine learning algorithms.

## Materials and Methods

### Data Collection

We selected the top 1,000 living contemporary artists based on their rankings by ArtFacts.net in January 2015, which was the time when we collected our data. It is worth noting that our focus in this study is on living contemporary artists for two main reasons. Firstly, if we were to include deceased artists, there would be a significant variation in terms of: a) the age of their artwork (i.e. how much time has passed since their death); b) random historical events; and c) the social consensus, which could all affect the price of their artwork at any given point. Secondly, there is a high probability of a reduction in uncertainty when artists die, as their impact on the market dynamics of their works is restricted or removed. Among art connoisseurs, the death of an artist restricts the supply of new artwork to the markets, allowing for the stabilization of the artist's identity and prestige<sup>3</sup>. Over time, there tends to be a common understanding about the deceased artist, which generally increases their value. However, it should be noted that this generalization may not hold true for relatively unsuccessful artists, for whom the information may be limited and inaccessible.

After selecting the famous living artists, we collected relevant information manually from their webpages, as well as publicly available popular art databases such as ArtNet.com, ArtPrice.com, Blouin Art Index, and Artfacts.net. We also collected auction sales data between 1996 and 2014 from Blouin Art Index and validated it by cross-checking with publicly available auction house data at the time, such as Christie's and Sotheby's. As art auction prices serve as a proxy for the market value of an artwork<sup>3,41,3323</sup>, we excluded artists who did not sell any artworks at auction houses, as well as those without publicly available information on their careers, such as exhibition records or artistic recognition awards.

At the same time, to empirically test the importance of visual qualities in artworks, we used publicly available thumbnail images of artwork associated with our auction records. High-resolution digital images of artwork are not usually available to the public due to the risk of unauthorized reproduction. Therefore, we collected publicly available thumbnail images. The availability of images online excluded 20,116 records, which dropped many of the transactions in 2013–14 when the images were not available. To avoid a highly skewed sample of later data, we cut off our sample in 2012. While some may argue that using low-resolution images leads to a loss of crucial visual information, it's important to note that no photographic resolution is ever high enough to fully capture the artistic content of a physical artifact, and different resolutions will result in different digital representations. Therefore, we utilized the largest amount of publicly available data with a consistent resolution.

Feature Name	Description
<i>Width, Height, and Square-inch size</i>	Width, height, and square inch size of an artwork
<i>Genre - f{Genre Name}</i>	Genre of an artwork (Painting, Print, Photograph, Drawing and Others)
<i>Age</i>	Age of an artist
<i>Gender</i>	Gender of an artist
<i>Education</i>	Education level of an artist (Domestic/Abroad, Elite School or not, and Degree)
<i>Price - f {Mean / Median / Maximum}</i>	Mean / Maximum / Minimum / Median price of the artworks by an artist during the year when auction is occurred.
<i>{Minimum / Mean / Median} price of {5 / 10} artworks</i>	Minimum / mean / median price of the artworks by an artist during the last five and ten transactions
<i>{Minimum / Mean / Median} price of {5 / 10} artworks weighted by size</i>	Size-standardized minimum / mean / median price of the artworks by an artist during the last five and ten transactions
<i>Award</i>	Records of art awards won by artists
<i>Biennial</i>	Records of participation in art fair or biennial
<i>Ranking</i>	Global artistic ranking of the artist as determined from U.S. and European art sources
<i>Solo and Group Shows</i>	Number of solo and group exhibitions for an artist until the auction year, respectively
<i>Match - Genre</i>	Whether an artwork is identified as part of the major genre of the artist
<i>Match - Country</i>	Whether an artwork is sold in the working country of its artist
<i>Private Acquisition</i>	Number of private acquisitions of artist works by individual collectors
<i>Public Acquisition</i>	Number of public acquisitions by museums
<i>Artist's working country - {Country Name}</i>	Country where an artist is mainly working
<i>Artist's living country - {Country Name}</i>	Country where the artist is living
<i>Auction house tier - {1/2/3/4}</i>	Tier of auction houses are coded based on the market dominance (the frequency and the volume of auctions) and the reputation of the auction house. Tier 1 includes global auction houses in major cities, such as Christie New York, Sotheby London, and de Pury New York, and the rest of international branches from the same global auction houses are grouped as tier 3 while the other major stand-alone auction houses were coded as tier 2 such as Dorotheum, Artcurial, and Kunsthaus Lempertz. Tier 4 is for the rest, non-major auction houses.
<i>Year</i>	Year of auction
<i>{Continent / Country} - {Minimum / Mean / Median / Maximum}</i>	Global art market size and growth based on the mean/maximum/median price levels across 26 countries in six continents for the year of auction
<i>Estimated {minimum / maximum} price by professionals</i>	Minimum and maximum prices of an artwork estimated by professionals

**Table 4.** List of features included in our social metadata analysis model without artworks' information and professional's estimation, with a brief description of the features.

As a result, our final dataset includes 34,200 records of artwork of 590 living contemporary artists across 23 countries, spanning 17 years (1996 to 2012), including images of the art, artist-related features, and global art market features. For established and emerging market samples, the former includes 29,853 records across four countries (USA, UK, France, and Germany) while the latter includes 4347 records from 19 countries.

Considering the temporal structure of our dataset, the market transactions of artworks at auctions, we split the dataset into a training set with the earlier 27,360 transaction records (80% of our records) and a test set (the later 20%, 6840 transactions), which enable us to respect the temporal ordering and avoid data-leakage about future trends. This split simulates predictions about auctions between May 2011 and December 2012 using a model trained in May 2011 based on the past 15 years of data. For established and emerging markets, we maintain the same cutoff date, yielding training sets of 24,163 and 3197 works, and test sets of 5690 and 1150, respectively.

The datasets generated and analyzed during the current study are available in the Figshare repository with the accession number(s): <https://figshare.com/s/6fcf9b970ad4818e7047>.

### Feature Selection—Metadata

Our model includes eight market-level features—such as minimum, mean, median, and maximum prices of all other auctions in the continent and country where the target artwork is being auctioned, as well as auction house tier—and thirty artist-level features, which are listed in Table 4. For artwork, we only considered size and genre information. While the size can significantly modulate the price, we can safely assume that it does not inform us about the objective quality of the artwork. Thus, we standardized the prices by size to control for the potential price differences based on the size of each artwork. Larger pieces of artwork usually cost more due to the quantity of material involved (canvas, inks, paints, etc), but this does not necessarily reflect the quality of the artwork as well.

### Feature Selection—Image

We first use a vector of 8971 features to describe the art images (960 GIST features, 2915 HOG features, 4096 color histogram features, and 1000 CNN features). While we explored end to end deep learning-based methods for price prediction, their tendency to overfit led to poor validation set performance, which makes us focus our

inquiry on interpretable visual features. In this section, we describe how these features are computed and how they roughly capture four key formal characteristics of art: composition, shape, color, and recognizable objects.

- *Composition*: GIST features<sup>26</sup> are a classical computer vision technique for scene classification that is constructed using multiple scales and orientations of Gabor filters. They coarsely capture the composition of a scene by identifying directional frequency content, such as the horizon in landscapes.
- *Shape*: HOG features<sup>27</sup> are another classical computer vision technique originally introduced for human detection, but then used for shape detection more broadly. They measure the histogram of oriented gradients, a circular histogram of all the directions of edges, for each region in the image.
- *Color*: Color histogram features are the number of pixels with each RGB color value, binned into 16 ranges per dimension (i.e. yielding 4096 bins). The captured color histograms are useful for identifying the dominant colors used in a work and are also widely used for content-based image retrieval (e.g.<sup>28, 48</sup>). These features capture color information.
- *Objects*: CNN features use the final layer of a deep convolutional neural network image classifier, in this case ResNet18<sup>31</sup>, trained on the ImageNet dataset<sup>29</sup>. Since ImageNet has 1000 classes, there are 1000 features corresponding to the image classes. While the image classes are defined based on photographic data which is quite different from artwork thumbnails, the latent space learned by such a model can still be useful for prediction.

We observe that these features, especially the color histogram, are quite sparse and often highly correlated. To reduce the redundant dimensionality and avoid overfitting, we perform principal component analysis, keeping 100 visual components for full dataset analysis and 10 components for analysis when split by market sector, to avoid overfitting.

In addition to these relatively low-level visual features which correspond to specific characteristics like edge directions and colors, we compute two high-level features: predictions of perceived colorfulness and complexity. We predict perceived colorfulness using the average distance from each pixel to the image mean on the a-b plane in Lab color space (adapted from<sup>49</sup>). We also predict perceived complexity using the average value of a Canny edge detector on the saturation channel in HSV color space<sup>50,51</sup>. We find that these features are both useful for visualizing an artist's collected work, as well as informative to the model.

### Extreme Gradient Boosting (XGBoost) Regressor

We use Extreme Gradient Boosting (XGBoost) regressor for our prediction models, which has won various machine learning competitions<sup>32</sup> as well as has been applied to a wide range of applications from predicting mortality of patients<sup>52,53</sup> to evaluating personal credit<sup>54</sup>. XGBoost is a high-performance machine learning algorithm based on the gradient-boosting decision tree. XGBoost can efficiently construct boosted trees and run in parallel—either regression trees or classification trees—and optimize the value of the objective function across the trees.

The major advantage of XGBoost is its scalability in diverse situations with a fast learning process<sup>32</sup>. The model works by combining a set of weaker machine-learning algorithms to obtain an improved machine-learning algorithm as a whole<sup>55</sup>. Also, XGBoost's recursive tree-based decision system provides great interpretability potential. Although the internal model mechanisms of black-box modeling strategies are still difficult to interpret in general, the importance of each individual feature can be determined by its accumulated use in each decision step in trees. The metric calculating the relative importance of each feature is valuable to estimate features that are the most enhancing the outcomes of a model, especially when they are related to meaningful valuation parameters.

For each dataset and feature set, we perform hyperparameter search over maximum tree depth and learning rate using a validation set with size equal to the corresponding test set sampled at random from the training set.

### Discussion

The proliferation of “big data” and enhanced computational capabilities for data analysis have enabled the empirical reassessment of classical questions regarding human behavior and culture. One such debate is the valuation of art, where advocates of the artistic forms assert that the worth of an artwork is determined by its significant visual forms to individual viewers, whereas social constructionists posit that the value of a unique object is influenced by collective social and cultural processes. After analyzing the relative importance of social and visual features in predicting auction prices of artworks, the findings demonstrate that social factors of artists and markets play a significantly large role in the valuation of contemporary artworks rather than visual forms. This effect is particularly pronounced in emerging markets relative to established ones.

It might not be surprising that visual features have limited influence on the valuation of contemporary art, given that the figurative differentiation of artworks is less important in this context. There is always a possibility that a better model with novel features will be able to predict the price better. Yet, the fact that the other model, with a mutually exclusive set of features, approaches the accuracy of expert prediction, reaching 75% of price variation explained in emerging markets, sets a tight upper limit on how much the visual features (or the “beauty”) of the artwork can explain the price. It is also important to note that this prediction was solely based on publicly accessible data, without any additional information about the artwork itself or insider knowledge.

This discovery underscores the notion that cultural consumption is primarily a socialized activity reinforced by social mechanisms<sup>6</sup> including the Matthew effects of existing recognition<sup>35</sup>, the social status of intermediaries (auction houses), and geographical locations of cultural markets.

Furthermore, the discrepancy between machine learning and professional estimation in established and emerging markets exhibits how socially constructed consumption exerts a significant influence on emerging market variation, leading to a stronger preference for “buying artists” and their prestige.

Lastly, although domain expertise remains critical in cultural industries, especially in established markets, as evidenced by our findings, this study’s results underscore the potential benefit of the human–machine learning model, integrating machine learning methods with expert knowledge to better understand the complex dynamics of cultural consumption. This approach has particular relevance for emerging markets, where professional expertise is less established.

However, our methodology does present certain limitations that warrant further exploration. First, since XGBoost models employ an extensive ensemble of decision trees, each individual tree determines price estimates based on a specific set of features. Consequently, the impact of each feature doesn’t follow a linear path but rather varies depending on the values of other features within the same tree. It is important to highlight that such flexibility and nonlinearity of certain features also reflect the complexities inherent in real-world valuation processes. In these contexts, the influence of one factor is often interrelated with, and affected by, other associated factors. However, our analysis of individual-level tendencies using the SHAP values<sup>56</sup> reveals that the directional impacts of our main features align consistently with findings from existing research. This correlation substantiates the reliability and coherence of our model (Supplementary Materials Fig. 2).

Second, our visual analysis, while not statistically significant, does not rule out the possibility that additional visual aspects can be significant to the valuation process in different ways. For example, a recent study on artworks’ strokes shows its importance for attribution and authentication<sup>57</sup>. The concept of the “aesthetic gap,” which refers to the disparity between information that can be directly derived from the image itself and information that requires inferring the emotional state and contextualized understanding of the image from a human viewer, has been discussed in the computer vision and multimedia retrieval communities<sup>58,59</sup>. It is possible that different approaches may capture the constructed meaning in contemporary art behind the various layers of an artist’s unique style, which can explain the relationship between visual features and valuations. Such investigations could result in better predictive performance on prices or other career patterns of contemporary artists.

Third, our prediction was solely based on publicly accessible data, without any data about insider knowledge. However, it is widely acknowledged that insider knowledge and informal relationships play a key role in the valuation of artworks, especially during the early stages of an artist’s career, influencing the pricing of their artworks<sup>11,60</sup>. This becomes particularly evident when there is a scarcity of prior transactions. As such, future research should focus on addressing the challenge of limited observations among emerging artists in nascent markets, aiming to explore the intricate connection between insider knowledge and the interpretation of status signals and artistic styles. This will significantly extend our understanding of cultural valuation in general, especially within the context of early-career producers and emerging markets.

Finally, one could argue that it is impossible to separate the influence of artist, artwork, and art market on price, as artwork-specific information is already included, for example, in the price history of artists, and artist identity can sometimes be determined based on visual information. While such leakage is possible, artworks are generally considered artistically unique in every way. The clear performance difference between visual and metadata models indicates that the information contained in the social metadata and visual features are mostly distinct.

Nevertheless, our study demonstrates a stark difference between the information conveyed by the visual and social aspects of artworks in influencing their market value, reinforcing the idea that the market value of cultural products is mostly determined by how they are perceived by others at a producer level, rather than their inherent characteristics at an object level. Our findings also suggest that audiences in emerging markets follow these patterns, thus making it easier to predict their market valuations using machine learning approaches.

## Data availability

The datasets generated and analyzed during the current study are available in the Figshare repository with the accession number(s): <https://figshare.com/s/6fcf9b70ad4818e7047>.

Received: 30 January 2024; Accepted: 29 April 2024

Published online: 21 May 2024

## References

1. Reyburn, S. How this leonardo’s mind-blowing price will change the art market—the new york times. (2017). <https://www.nytimes.com/2017/11/24/arts/design/salvator-mundi-leonardo.html> (accessed 24 Jan 2021).
2. Shin, D., Lee, K. & Lee, H. Neoliberal marketization of art worlds and status multiplexity: Price formation in a Korean art auction, 1998–2007. *Poetics* **43**, 120–148 (2014).
3. Penet, P. & Lee, K. Prize & price: The turner prize as a valuation device in the contemporary art market. *Poetics* **43**, 149–171 (2014).
4. Kopytoff, I. The cultural biography of things: Commoditization as process. *The Social Life of Things. Comm. Cult. Perspect.* **2**, 65–91 (1986).
5. Hutter, M. & Shusterman, R. Value and the valuation of art in economic and aesthetic theory. *Handb. Econ. Art Cult.* **1**, 169–208 (2006).
6. Becker, H. S. *Art Worlds* (Berkeley University, 1982).
7. Bourdieu, P. *Distinction: A Social Critique of the Judgement of Taste* (Harvard University Press, 1984).
8. McAndrew, C., Smith, J. L. & Thompson, R. The impact of reserve prices on the perceived bias of expert appraisals of fine art. *J. Appl. Econ.* **27**(2), 235–252 (2012).
9. Salganik, M. J., Dodds, P. S. & Watts, D. J. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* **311**(5762), 854–856 (2006).
10. Fraiberger, S. P., Sinatra, R., Resch, M., Riedl, C. & Barabási, A. L. Quantifying reputation and success in art. *Science* **362**(6416), 825–829 (2018).

11. Velthuis, O. *Talking Prices: Symbolic Meanings of Prices on the Market for Contemporary Art* Vol. 55 (Princeton University Press, 2005).
12. Karpik, L. & Scott, N. *Valuing the Unique: The Economics of Singularities, The Economics of Singularities* (Princeton University Press, 2010).
13. Velthuis, O. & Curioni, S. B. *Cosmopolitan Canvases: The Globalization of Markets for Contemporary Art* (Oxford University Press, 2015).
14. Eyck, T. A. T. & Busch, L. Justifying the art critique: Clement greenberg, michael kimmelman, and orders of worth in art criticism. *Cult. Sociol.* **6**(2), 217–231 (2012).
15. Yogev, T. The social construction of quality: Status dynamics in the market for contemporary art. *Socio-Econ. Rev.* **8**(3), 511–536 (2010).
16. Kharchenkova, S. & Velthuis, O. How to become a judgment device: Valuation practices and the role of auctions in the emerging Chinese art market. *Socio-Econ. Rev.* **16**(3), 459–477 (2017).
17. Lee, K. Neoliberal marketization of global contemporary visual art worlds: Changes in valuations and the scope of local and global markets In *Art and the Challenge of Markets Volume 2* 65–97 (Springer, 2018).
18. Bell, C. Art. In *Art and Its Significance: An Anthology of Aesthetic Theory* (ed. Ross, S. D.) (SUNY Press, 1994).
19. Fry, R. *An Essay in Aesthetics* (Routledge, 1982).
20. Bell, C. *Art* (Chatto & Windus, 1916).
21. Rose, S. *Art and form: From roger fry to global modernism, refiguring modernism* (Penn State University Press, 2019).
22. Lee, B. et al. Dissecting landscape art history with information theory. *Proc. Natl. Acad. Sci.* **117**(43), 26580–26590 (2020).
23. Galenson, D. W. *Painting Outside the Lines: Patterns of Creativity in Modern Art* (Harvard University Press, 2001).
24. Renneboog, L. & Spaenjers, C. Buying beauty: On prices and returns in the art market. *Manag. Sci.* **59**(1), 36–53 (2013).
25. Dempster, A. M. *Risk and Uncertainty in the Art World* (Bloomsbury Publishing, 2014).
26. Oliva, A. & Torralba, A. Modeling the shape of the scene: A holistic representation of the spatial envelope. *Int. J. Comput. Vis.* **42**(3), 145–175 (2001).
27. Dalal, N., & Triggs, B. Histograms of oriented gradients for human detection In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, Vol. 1, 886–893 (IEEE, 2005).
28. Zujovic, J., Gandy, L., Friedman, S., Pardo, B., & Pappas, T. N. Classifying paintings by artistic genre: An analysis of features & classifiers. In *2009 IEEE International Workshop on Multimedia Signal Processing 1–5* (IEEE, 2009).
29. Deng, J., et al. Imagenet: A large-scale hierarchical image database In *2009 IEEE Conference on Computer Vision and Pattern Recognition* 248–255 (IEEE, 2009).
30. Krizhevsky, A., Sutskever, I., & Hinton, G. E. Imagenet classification with deep convolutional neural networks. *Adv. Neural Inf. Process. Syst.* **25** (2012).
31. He, K., Zhang, X., Ren, S., & Sun, J. Deep residual learning for image recognition In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 770–778 (2016)
32. Chen, T., & Guestrin, C. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 785–794 (2016).
33. Mei, J. & Moses, M. Vested interest and biased price estimates: Evidence from an auction market. *J. Finance* **60**(5), 2409–2435 (2005).
34. Zuckerman, H. *Scientific Elite: Nobel Laureates in the United States* (Transaction Publishers, 1977).
35. Merton, R. K. The matthew effect in science: The reward and communication systems of science are considered. *Science* **159**(3810), 56–63 (1968).
36. Podolny, J. M. A status-based model of market competition. *Am. J. Sociol.* **98**(4), 829–872 (1993).
37. Podolny, J. M. *Status Signals: A Sociological Study of Market Competition* (Princeton University Press, 2010).
38. Espeland, W. & Sauder, M. Rankings and reactivity: How public measures recreate social worlds. *Am. J. Sociol.* **113**(1), 1–40 (2007).
39. Beckert, J. Where do prices come from? Sociological approaches to price formation. *Socio-Econ. Rev.* **9**(4), 757–786 (2011).
40. Tomlinson, J. Cultural imperialism. *The Wiley-Blackwell Encyclopedia of Globalization* (2012).
41. Mei, J. & Moses, M. Art as an investment and the underperformance of masterpieces. *Am. Econ. Rev.* **92**(5), 1656–1668 (2002).
42. Alexander, V. & Bowler, A. Art at the crossroads: The arts in society and the sociology of art. *Poetics* **43**, 1–19 (2014).
43. Adam, G. *Big Bucks: The Explosion of the Art Market in the 21st Century* (Lund Humphries, 2014).
44. Cattani, G., Ferriani, S. & Allison, P. D. Insiders, outsiders, and the struggle for consecration in cultural fields: A core-periphery perspective. *Am. Sociol. Rev.* **79**(2), 258–281 (2014).
45. Buchholz, L. What is a global field? Theorizing fields beyond the nation-state. *Sociol. Rev. Monogr.* **64**(2), 31–60 (2016).
46. Buchholz, L. Rethinking the center-periphery model: Dimensions and temporalities of macro-structure in a global field of cultural production. *Poetics* **71**, 18–32 (2018).
47. Wohl, H. Performing aesthetic confidence: How contemporary art collectors maintain status. *Socio-Econ. Rev.* **18**(1), 215–233 (2019).
48. Huang, J., Kumar, S. R., Mitra, M., Zhu, W. J., & Zabih, R. Image indexing using color correlograms In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition* 762–768 (IEEE, 1997).
49. Hasler, D., & Suesstrunk, S. E. Measuring colorfulness in natural images In *Human Vision and Electronic Imaging VIII*, Vol. 5007, 87–95 (International Society for Optics and Photonics, 2003)
50. Machado, P. et al. Computerized measures of visual complexity. *Acta Psychol.* **160**, 43–57 (2015).
51. Canny, J. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.* **6**, 679–698 (1986).
52. Yan, L. et al. An interpretable mortality prediction model for covid-19 patients. *Nat. Mach. Intell.* **2**(5), 283–288 (2020).
53. Liu, J. et al. Predicting mortality of patients with acute kidney injury in the ICU using xgboost model. *PLoS One* **16**(2), e0246306 (2021).
54. Li, H., Cao, Y., Li, S., Zhao, J. & Sun, Y. Xgboost model and its application to personal credit evaluation. *IEEE Intell. Syst.* **35**(3), 52–61 (2020).
55. Chen, X., Wang, Z. X. & Pan, X. M. Hiv-1 tropism prediction by the xgboost and hmm methods. *Sci. Rep.* **9**(1), 1–8 (2019).
56. Lundberg, S. M., & Lee, S. I. A unified approach to interpreting model predictions. *Adv. Neural Inf. Process. Syst.* **30** (2017).
57. Elgammal, A., Kang, Y., & Den Leeuw, M. Picasso, matisse, or a fake? Automated analysis of drawings at the stroke level for attribution and authentication. *Proc. AAAI Conf. Artif. Intell.* **32**(1), (2018)
58. Datta, R., Li, J., & Wang, J. Z. Algorithmic inferencing of aesthetics and emotion in natural images: An exposition In *2008 15th IEEE International Conference on Image Processing* 105–108 (IEEE, 2008).
59. Goree, S. What does it take to cross the aesthetic gap? The development of image aesthetic quality assessment in computer vision. In *International Conference on Computational Creativity* (2021).
60. Wohl, H. *Bound by Creativity: How Contemporary Art Is Created and Judged* (University of Chicago Press, 2021).

## Acknowledgements

We benefited from comments by Paul DiMaggio, Monika Kackovic, Karunakar Mannem, and Minsu Park. An earlier version of the paper was presented at the 2019 IC2S2, 2022 Arts and Data Conference at New York

University (New York USA). Y.Y.A. was supported by the Air Force Office of Scientific Research under award number FA9550-19-1-0391. J.P. was supported by the National Research Foundation of Korea (NRF) Grant through the Korea Government under Grant NRF-2022R1A5A7033499.

### Author contributions

K.L., J.P., and S.G. conceived the study and designed the research; K.L., J.P., S.G., D.C., and Y.A. wrote the manuscript; K.L., J.P., and S.G. performed the literature review; K.L. collected the data; K.L., J.P., and S.G. analyzed the data; J.P. and S.G. produced the visualizations; K.L., J.P., S.G., D.C., and Y.A. provided revisions.

### Competing interests

The authors declare no competing interests.

### Additional information

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1038/s41598-024-60957-z>.

**Correspondence** and requests for materials should be addressed to K.L.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024