



How to raise artwork prices using action rules, personalization and artwork visual features

Laurel Powell^{1,2}  · Anna Gelich^{3,4} · Zbigniew W. Ras^{1,3}

Received: 31 December 2020 / Revised: 27 June 2021 / Accepted: 12 July 2021 /
Published online: 21 August 2021
© Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

This work explores the development of action rules for changing the prices of works of contemporary fine art. We used LISP-Miner to generate action rules related to artwork profiles and developed attributes covering artist descriptions and visual features of the artwork. We focus heavily on developing a method for partitioning a dataset to produce an increase in the coverage of the rule sets.

Keywords Art analytics · Data mining · Action rules

1 Introduction

Artists in the contemporary market for fine art must make many decisions on how to present their work for commercial sale. While these choices are crucial to artists' long term commercial success, they are given limited guidance on how to present their work online to its best advantage. Can the price that an artist asks, regarding a work, be altered by changing minor aspects that their potential customers see? What drives the price for a work of fine art?

The challenge of objective valuation is worsened by the complex nature of pricing in the art market. There are a number of theories on how pricing in the art market works. As said in Velthuis (2005),

✉ Laurel Powell
lpowel28@uncc.edu

Anna Gelich
agelich@uncc.edu

Zbigniew W. Ras
ras@uncc.edu

¹ College of Computing and Informatics, University of North Carolina at Charlotte, Charlotte, NC, USA

² Catawba College, Salisbury, NC, USA

³ Polish-Japanese Academy of Information Technology, Warsaw, Poland

⁴ College of Arts + Architecture, University of North Carolina at Charlotte, Charlotte, NC, USA

In a market that has to reconcile a fierce opposition between commercial and artistic values, and that has to commodify goods whose essence is considered to be non-commodifiable, dealers, artists, and collectors find ways to express and share non-economic values through the economic medium of pricing.

This work proposes a method for developing a set of action rules that can be used to suggest changes for artists that will improve the sales potential of their work. Ultimately, these rules could contribute to the construction of a decision support system for artists. These rules reclassify artworks from one price level to another higher price level. This work discusses possible sets of flexible attributes and proposes a basic set of stable attributes (Ras & Wiczorkowska, 2000). Stable attributes are generally aspects of the relevant item that are unlikely or difficult to change. In the context of this work, these represent the physical aspects of the artwork and fundamental qualities such as its medium. Flexible attributes are features that can be altered by an interested party to change its classification. In the context of this work, these represent aspects of how the artist presents themselves and their work to potential customers. This work is an extension of Powell et al. (2020c), which was presented at ISMIS 2020. It adds greater context and more experiments to further develop the original work.

2 Related work

Applying computing techniques to artworks has received considerable attention from the academic community. Some researchers have addressed the challenge of creating a recommender system that chooses artworks similar to a given one for a user (Saleh & Elgammal, 2015), while others are focused on automatically tagging a painting with emotions based on the colors represented (Kang et al., 2018).

Art price valuation has received limited attention. One interesting project in the realm of price prediction was Hosny's work (Hosny et al., 2014). As discussed in Bailey (2017), one notable trend explored in Hosny's work was that certain colors, specifically blacks, whites, and grays, are more likely to have a higher sales valuation. Different approaches have been used for price prediction, such as the approach used by Liu and Woodham (2019) to value Rothko paintings. In Galbraith and Hodgson (2018), hedonic regression, a statistical modelling method based on characteristics of the object in question, combined with past sales was used to predict art value.

Action rules were first proposed in Ras and Wiczorkowska (2000) as a method of reclassifying objects from one group to another group by changing values of their flexible attributes. They have been used for medical data, such as in Tarnowska et al. (2017), Mardini and Raś (2019), and Hajja et al. (2014) and for business purposes, as in Ras and Wiczorkowska (2000) and Tarnowska et al. (2020), and social networks (Kalanat & Khanjari, 2020). Some have explored expansions on the original methodology, such as considering the cost and feasibility of the implementation of action rules as discussed in Tzacheva and Raś (2005) and Tzacheva et al. (2017). The cost of an action rule, which was initially proposed in Tzacheva and Raś (2005), represents the cost of changing flexible attribute values listed in the rule from their initial values to other specific values.

This work uses LISp-Miner's Ac4ft-Miner tool, introduced in Rauch and Šimůnek (2009), discussed in Nekvapil (2009), and further developed in Rauch et al. (2019). The Ac4ft-Miner tool is an implementation of the GUHA method for action rule generation (Rauch & Šimůnek, 2009). In this method of action rule construction, objects resembling

traditional association rules are extracted from the dataset and used to form contingency tables (Rauch & Šimůnek, 2009). Two of these objects, with matching stable attributes, are used to build the action rules. One rule, termed here as the “before rule”, has the stable and flexible attributes of the object before any actions are taken, and the other rule, called the “after rule” here, has the desired set of flexible attributes and matching stable ones. Rules generated in this manner consist of an antecedent and succedent, association rules in the form $\varphi \approx \psi$ where the symbol \approx represents the association between φ and ψ (Rauch & Šimůnek, 2009). These associations are then used to create action rules, called G-Action Rules, through the examination of dependencies and similarities across the rules generated in the previous step (Rauch & Šimůnek, 2009). This strategy of action rules discovery is similar to the one presented in Tsay and Raś (2005).

As a hypothetical example, using the features addressed in this work, an association may be found between an artwork being in a realistic style with a very short description and this combination implies a very low price. Another association might be found between an artwork being in a realistic style with an average length description and this implies a considerably higher price. In this case, the style would be a stable attribute, and the length of the description would be a flexible attribute. As raising the price of an artwork is advantageous to the artist, we term the lower priced rule as the before rule and the higher price as the after rule. Assuming that both rules have sufficient support and confidence, it would be possible to list as a rule that if an artwork is in a realistic style if the length of the description changes from very low to average then the price will rise from one level to another.

3 Dataset

The dataset used in this work is an expanded version of the one used in Powell et al. (2019), Powell et al. (2020a), and Powell et al. (2020b). Our work expands on the ideas discussed in these works by adding action rules. The following section addresses the construction of the dataset as originally discussed in the preceding works.

We use a dataset of artworks and artist information collected from the website Artfinder.com (Artfinder.com, 2019). Artfinder represents artists from all over the world and has diverse styles and subjects available. Artfinder has paintings, prints, photographs, sculptures, and drawings posted for collectors. Using a single source for artworks allows for more consistent definitions of tags. The information was scraped using Beautiful Soup (Beautiful Soup, 2019) to parse the pages, Apache Selenium (Selenium, 2019) to work with dynamic webpages and Javascript, and Python to fetch pages and all other major steps. It contains approximately 200,000 artworks from approximately 3,300 artists.

Each artwork posted on Artfinder has a dedicated page describing the work. This page consistently has at least one photograph of the work and has its descriptive information, as well as the tags that the artist selected. These tags are used by potential customers to search for works. Some of the analysis methods discussed below use the primary image of the artwork from the product page. A small number of artwork images could not be retrieved, so they are omitted from the following analysis.

Artists on Artfinder create profiles of themselves discussing their achievements and backgrounds. The profiles have places for artists to provide their biographies, links to their social media, and descriptions of achievements in their lives. A number of artists have reviews posted from customers and are rated on a star system with five stars being the best. The majority of artists with one or more reviews have ratings of four out of five stars or

higher and very few have any low reviews. The average score of all artists in the dataset with any star rating is 4.904. This makes the star ratings received by artists of very low value as predictive features.

After the initial scraping, the dataset was extended by processing both the artwork images and the associated text to form new sets of features. This processing is detailed in a later portion of this paper. Both the originally scraped data and the new features were stored as a collection of database tables. These tables were queried and selected attributes and tuples were converted into comma separated files for use in LISp-Miner using Python. For example, a set of selected features from three records have been reproduced in Table 1. These records were inserted into LISp-Miner, discretized, and used for rule generation. The set shown uses all the flexible features, the base set of stable features that were used each time, and three additional stable features of three prominent colors.

4 Methodology

This work focuses heavily on exploring the relative value of different sets of features. These sets are examined by testing combinations of features with varying conditions in LISp-Miner. LISp-Miner takes csv files as input. Each file was made using a subset of the available features and a portion of the available tuples. Different methods of partitioning the dataset were also explored. Images and unstructured text required processing before they

Table 1 Sample records

PRICE	MEDIUM	ARTISTIC_STYLE	SUBJECT
150	Acrylic	Abstract	People and portraits
4000	Oil	Impressionistic	Landscapes, sea and sky
249.94	Watercolor	Impressionistic	Landscapes, sea and sky
HEIGHT	WIDTH	BIOWC	DESCWC
8.5	11	193	43
40	40	632	16
20.08	14.17	226	38
FACEBOOK	TWITTER	INSTAGRAM	REVIEWS_PRESENT
Facebook	Twitter	Instagram	Reviews Present
Facebook	Twitter	Instagram	No Reviews
Facebook	No Twitter	No Instagram	Reviews Present
BIO_POS	BIO_NEG	DESC_POS	DESC_NEG
0.247	0.012	0.165	0.093
0.038	0	0.128	0
0.151	0	0.071	0
EDGE_PERCENTAGE	COLOR1	COLOR2	COLOR3
5.35	white	black	white
6.73	black	black	brown
5.3	gray	gray	gray

could be used as input for LISp-Miner. The forms of processing used, such as reducing the image to a group of colors or the text to a simple numeric sentiment score, is further discussed below. For each combination, LISp-Miner discretized the data and output the G-Action rules. In order to explore the impact of feature selection, a variety of thresholds for minimum support and confidence were used in rule construction. Lastly, the rule sets were examined to assess their coverage of the input tuples and their average level of confidence.

The decision feature considered here, which represents the succedent of the G-Action Rule, is the listed asking price of the artwork. Artists could attempt to make changes to how their artworks are presented to the public in order to increase their sales price and should be cautious about making changes that could potentially lower their sales price. The prices were discretized into ten groups using LISp-Miner (Rauch et al., 2019). The cuts were placed automatically to create partitions containing an approximately equal number of tuples. The price attribute was partitioned into ten levels, referred to in later sections as levels one through ten. Price is treated as a flexible attribute.

This work defines a set of stable attributes that are used in the antecedent portion of the G-Action Rule. Rules can be used to describe a work of art. The features chosen for the rules take inspiration from the sets used in Pawlowski et al. (2019). The chosen medium of the artwork is used as a stable feature. When the artwork was posted, the artist tagged it with a medium for search purposes. In order to keep the number of options within reasonable bounds, the mediums were discretized into seven categories and one “NA” category. The artistic style of the work as well as the artist’s subject is similarly used as a tag and as a stable attribute. As discussed in Pawlowski (Pawlowski et al., 2019), the size of a work is very relevant to determining its price, so the length and width of the works were discretized into five bins of roughly equal size using the LISp-Miner (Rauch et al., 2019) discretization tool and used as stable attributes. Additionally, the presence or absence of visible reviews was used as a stable binary feature. The review scores of the artists are consistently very high if they are present at all. This limits the utility of using the score as a quality metric. Lastly, the percentage of edges in the work was utilized as a stable attribute after being discretized. To calculate this, Canny edge detection, which was developed by John Canny in Canny (1986), was used from the OpenCV set of tools (Team, 2017) on the primary artwork image. A small percentage of artworks had errors when an attempt was made to retrieve the photograph, so those tuples were removed from consideration. This algorithm determines if an individual pixel represents an edge or not based on the amount of color variation surrounding it. The number of edge pixels and non-edge pixels are then counted, and the number of edge pixels was used as a stable feature. This can be considered as giving a rough idea of the amount of detail or amount of color variation in the work. The list of basic stable attributes used each time is given below. Others were added and will be discussed later.

- Artistic Style (categorical)
- Artistic Subject (categorical)
- Medium (categorical)
- Height (numeric, discretized into categorical)
- Width (numeric, discretized into categorical)
- Artist Has Visible Reviews (boolean)
- Percentage of the Artwork Representing Edges (numeric, discretized into categorical)

The set of flexible attributes explored, part of the antecedent of the rules, centers around how a work will be perceived by a consumer. How long is the artist’s biography? Do they have a presence on social media? What is the tone of their writing? These are easily

changeable for an artist hoping to improve their sales. These changes have a very low cost, so stakeholders may be more willing to try recommended rules. The full list of flexible features used is given below.

- Word Count of the Artist Biography (Bio. WC) (numeric, discretized into categorical)
- Word Count of the Artwork Description (Desc. WC) (numeric, discretized into categorical)
- Social Media (Considered Together as SM)
 - Artist Listed a Facebook Profile (boolean)
 - Artist Listed a Twitter Profile (boolean)
 - Artist Listed an Instagram Profile (boolean)
- Positive Sentiment Level of Biography (Bio. Ps.) (numeric, discretized into categorical)
- Negative Sentiment Level of Biography (Bio. Ng.) (numeric, discretized into categorical)
- Positive Sentiment Level of Artwork Description (Desc. Ps.) (numeric, discretized into categorical)
- Negative Sentiment Level of Artwork Description (Desc. Ng.) (numeric, discretized into categorical)

The number of words in the biography of the artist and in the artwork description are used as predictive features. As was discussed in Powell et al. (2019), Powell et al. (2020a), and Powell et al. (2020b), the word count does have some utility as a price predicting feature. In other arenas of online sales, the length and wording of a description can have a bearing on the saleability of an item. In (Rawlins & Johnson, 2007), the authors analyzed sales on Ebay and determined that the length of the description could have an impact on the price.

Similarly, many collectors find artists through social media (The Hiscox Online Art Trade Report, 2018). It has become an increasingly important tool for helping artists be discovered by collectors. Can adding or removing a link to a profile change the opinion of a collector?

Determining the sentiment of the text was done using VADER, the ‘Valence Aware Dictionary for sEntiment Reasoning’, which is part of the Python Natural Language Toolkit (Bird et al., 2009; Hutto & Gilbert, 2015). The objective was to determine the polarity of the sentiment of the text. This strategy uses a set of words and phrases to determine the text’s sentiment, called an ‘opinion lexicon’ (Aggarwal, 2018). VADER uses a combination of an opinion lexicon and a set of rules (Hutto & Gilbert, 2015). The lexicon includes thousands of candidate terms rated by humans on a scale, and the rule set considers the impact of capitalization and negation (Hutto & Gilbert, 2015). Each piece of text was given a score for its level of positive sentiment, its level of negative sentiment, and its level of neutral sentiment. Each attribute was discretized, and they were used as flexible attributes. The sentiments found were largely neutral, with only a few pieces of text containing strongly emotional language.

Lastly, the use of color in the artwork was explored as an additional stable feature in the rule sets. The pixels in the artwork were clustered using Open-CV (Team, 2017) across the RGB dimensions using K-Means to determine the ten most frequently appearing colors. Out of this set of ten, the centroids of the largest cluster and second largest cluster were tested against eleven reference colors using the CIEDE2000 color difference measure (Sharma et al., 2004), which was implemented using (Taylor, 2014). The reference colors are based on the idea of universal basic colors which was first proposed in Berlin and Kay (1969).

While this work has been challenged and the colors proposed have been refined such as in Lindsey and Brown (2006) and Roberson and Hanley (2007), the concept of basic colors is useful for classification. The eleven reference colors are white (R 255, G 255, B 255), gray (R 128, G 128, B 128), black (R 0, G 0, B 0), red (R 255, G 0, B 0), orange (R 255, G 128, B 0), yellow (R 255, G 255, B 0), green (R 0, G 255, B 0), blue (R 0, G 0, B 255), purple (R 128, G 0, B 128), pink (R 255, G 192.0, B 203), and brown (R 63.8, G 47.9, B 31.9). The RGB values, other than brown, are taken from Colors by name (2020) and the brown comes from Labrecque and Milne (2012).

4.1 Additional stable attributes

While utilizing the basic color has value as a classification feature, many of the distinctions between colors and their subtleties are lost. Therefore, other methods of using colors were used. The following section discusses two different methods of mapping colors to specific emotions.

First, the pleasure, arousal, and dominance scales, referred to here as PAD, were used as stable attributes. The PAD matrix is a three-dimensional space with independent axes for pleasure (ranging from pain to extreme pleasure), arousal (ranging from sleepiness to excitement), and dominance (ranging from powerless to control) (Mehrabian, 1978; Russell & Mehrabian, 1977). Relationships between the saturation and brightness of colors to the PAD scale were found in Valdez and Mehrabian (1994). Interestingly, this work determined that hue has a much less clear relationship with this scale (Valdez & Mehrabian, 1994). Valdez found equations reflecting these relationships using regression analysis.

$$Pleasure = 0.69Brightness + 0.22Saturation \quad (1)$$

$$Arousal = -0.31Brightness + 0.60Saturation \quad (2)$$

$$Dominance = -0.76Brightness + 0.32Saturation \quad (3)$$

These equations were used on the colors extracted from each artwork using K-Means to create a weighted average score for each work. This was then discretized using Lisp-Miner and treated as a stable attribute.

Another method of quantifying color examined is the activity, weight, and heat scale, referred to here as AWH, originally proposed in Ou et al. (2004). This scale was designed to reduce the variations in the emotional impact of color based on culture or gender (Ou et al., 2004). These measures attempt to combine other descriptors of color (Ou et al., 2004). Color activity could be described as a measure of a color's freshness, cleanliness, modernity, or passivity (Ou et al., 2004). Color weight is a description of its hardness, heaviness, and whether it could be described as masculine or feminine (Ou et al., 2004). Heat was included due its importance in other studies as determined by Ou et al. (2004).

The work of Ou et al. (2004) derived the following equations for calculating these features. The CIELab values of L^* , C^* , h , a^* and b^* are used.

$$Color\ Activity = -2.1 + 0.06 \left[(L^* - 50)^2 + (a^* - 3)^2 + \left(\frac{b^* - 17}{1.4} \right)^2 \right]^{1/2} \quad (4)$$

$$Color\ Weight = -1.8 + 0.04(100 - L^*) + 0.45 \cos(h - 100^\circ) \quad (5)$$

$$Color\ Heat = -0.5 + 0.02(C^*)^{1.07} \cos(h - 50^\circ) \quad (6)$$

As discussed above, the dominant colors of each artwork were extracted using K-Means and applied to the above equations to determine a weighted average for the work. These were then discretized.

4.2 Block averaging and clustering

The patterns of darkness and light in an image are among the most immediately obvious visual features. Further stable attributes considered were the average brightness of the image, determined using the average grayscale value of the work, the standard deviation of that brightness, and the average and standard deviation of the red, green, and blue color channels. In the following tables, the mean and standard deviation of the grayscale image is referenced as “gray”; the mean values of the red, green and blue channels are abbreviated as “RGB M”, and the standard deviations of the same channels are abbreviated as “RGB S”.

Another method of defining a work is by analyzing these patterns. The sample photo of each work was sliced into a three-by-three grid with an approximately equal number of pixels in each partition. This resembles the method used in Lombardi (2005) to quantify artwork features.

The position in the three-by-three grid of the brightest and darkest blocks were considered as stable attributes. This is abbreviated as “block” in the subsequent tables. In Fig. 1, an example of this method of representing an image is shown.

Next, the average brightness level of each block was calculated. Then, all works by the same artist were averaged block by block. This creates a single representative set of values that represents how their specific work tends to handle darkness and light. This was then used to generate clusters that would place artists with visual similarities together. More clusters were generated using RGB values in addition to darkness and light alone, but the resulting clusters contained almost the exact same sets of artists as the gray-based clusters, so they were not used further. This method is referred to in later tables as “Gray Blocks”. The process was then repeated with the average percentage of edges per block. This method is referenced as “Edge Blocks”. This creates values that represent the overall composition of an artist’s body of work.

Using the LISp-Miner Ac4ft-Miner tool, sets of rules were generated for combinations of prices, flexible attributes, a single set of stable attributes used throughout for consistency, and additional stable attributes. As the goal is to move prices from lower to higher price points, rules were generated to transition items from lower price levels to higher price levels.

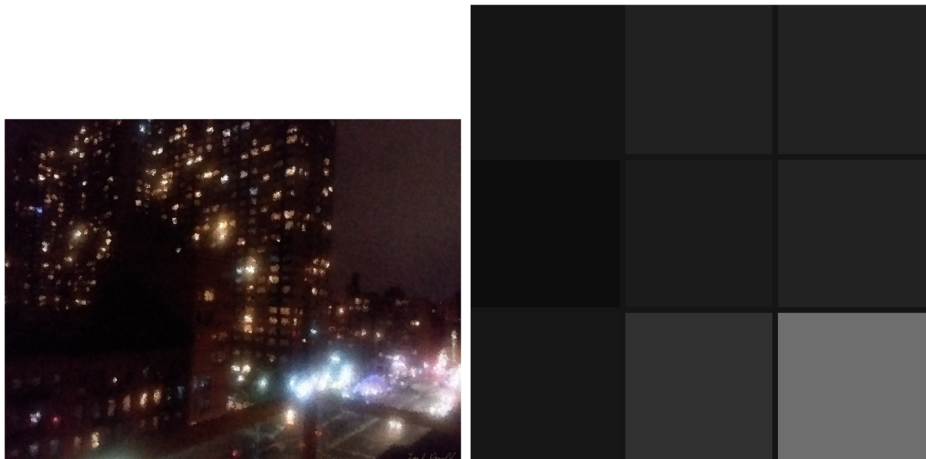


Fig. 1 A cityscape and its representation as blocks

Rules generated using this method have support and confidence scores for both the before-state association rule and after-state association rule that are used together to construct the action rule.

5 Results

The first set of rules generated had no minimum confidence for a rule and had a requirement that the rule must apply to 50 or more tuples out of the dataset. Rules were generated for each flexible attribute alone, the social media features alone, and the social media features in combinations. This work compares the results of different flexible attributes alone to explore which have the greatest impact on the confidence and support of the resulting rules. Each set of rules contains an average of 1,700 action rules. At least one rule applies to approximately 97% of the input tuples at the associated price level. This is termed as the coverage of the rule set and can be used as a measure of its applicability. However, the confidence in these rules is quite low. As all the features discussed here have an extremely low cost, exploring low confidence rules is not a concern. To make the changes these rules suggest, an artist would only need to rewrite their artwork descriptions or change their profile. While some individual rules do have high confidence, the average confidence for the before rule is approximately 13%.

In response to the low confidence of the initial sets of rules tested, additional rule sets were generated with greater constraints on the generation process. Rules were only generated to move each artwork up to the next higher price level. The social media flexible features (abbreviated in tables as SM) were used to generate rules as a group. The word counts (abbreviated in tables as WC) of both the artist's biography and description (abbreviated as Bio. and Desc., respectively) and text polarity features (abbreviated as Ps. for positive sentiment and Ng. for negative sentiment) were all used separately. To assess if overly strict support requirements were causing low confidence, the minimum support level for a rule was lowered from 50 to two for both the before and after rules, but the minimum confidence for a rule to be considered was changed to 60% for the before and after rules. Since the recommended changes would be very simple and low risk, a very low threshold for support was not a concern. The minimum number of attributes for the stable portion of the rule is one and the maximum number is five. The rule set generated here had a considerably lower average coverage at 9.375% but a considerably higher average confidence for the before rules at 79.16%, with the after rules at 79.2%. Considerably fewer rules were generated, and the average number of rules per group was 1,155. Table 2 displays the coverage of each rule set for the selected price levels and flexible attribute. This represents the percentage of objects at the lower price level that had at least one applicable rule to raise that object to the selected higher price level.

The level of support varies dramatically depending on the price level being addressed. The exact values of the prices referenced in the next table are as follows: (<12.97 - 58.28), (58.28 - 95.90), (95.90 - 130.04), (130.04 - 188.13), (188.13 - 250), (250 - 350), (350 - 490), (490 - 742), (742 - 1351.24), (1351.24 - >1,000,000). For example, the coverage of the rules sets that move artworks from the lowest prices, level one, to the next lowest price level, level two, ranges between 19.24% and 38.89% across all the different sets of flexible attributes. Notably, the social media attributes had markedly worse performance than the others. The attributes that changed an artwork's description had less consistently high coverage than the attributes that addressed an artist's biography. At higher values for

Table 2 Coverage of rules generated using base stable features

	SM	Bio. WC	Desc. WC	Bio. Ps.	Desc. Ps.	Bio. Ng.	Desc. Ng.
Prices 1->2	19.24	38.89	30.26	37.05	26.56	35.86	19.45
Prices 2->3	4.31	16.35	13.03	17.04	11.97	9.70	5.04
Prices 3->4	3.60	11.91	9.52	11.42	9.40	7.78	3.52
Prices 4->5	2.51	7.87	5.80	8.02	5.67	4.41	2.35
Prices 5->6	2.21	7.31	5.91	7.82	5.54	5.18	1.64
Prices 6->7	1.30	8.20	6.91	7.96	5.77	3.41	3.54
Prices 7->8	1.34	7.28	4.80	6.34	4.41	3.47	1.54
Prices 8->9	1.94	7.79	6.59	8.91	6.03	5.14	2.05
Prices 9->10	4.77	14.25	12.18	13.17	10.87	11.35	5.21

price, coverage decreases. This may be due to the size of the shifts necessary to move prices from one tier to another at the higher levels.

To expand on the set of stable attributes, another set of rules was generated that added the dominant color of the work to the list of stable attributes. This set was generated using a randomly selected subset with approximately 100,000 tuples. LISp-Miner's discretization tools were used to create a new set of partitions similar to those used previously for this set. The same restrictions on rule generation were repeated. Rules must have a minimum confidence of 60% for the before and after rules, and rules must have a minimum support of two. Slightly more rules were generated per set than in the previous variation with an average number of rules per group of 1,716. This set did have a slightly higher average coverage of 16.408% and a similar average before confidence of 78.65%, with an after confidence of 78.73%. As with the previous set of rules, the coverage shifts across different price levels. The coverage of each rule group changes depending on the selected flexible attributes and the selected price levels as demonstrated in Table 3.

Rules were developed using new criteria to improve on the coverage of the rules generated above. New stable attributes were considered, support and confidence were still restricted, and rules could now have up to two flexible attributes and two stable attributes. Minor adjustments were made to the manual discretization, so that the attribute medium

Table 3 Coverage of rules generated using base stable features and main color

	SM	Bio. WC	Desc. WC	Bio. Ps.	Desc. Ps.	Bio. Ng.	Desc. Ng.
Prices 1->2	25.45	55.11	47.57	54.12	44.77	48.22	26.05
Prices 2->3	8.93	26.44	22.25	25.33	21.92	17.04	8.05
Prices 3->4	5.62	20.50	17.12	18.83	15.88	12.03	7.27
Prices 4->5	5.09	13.57	12.66	14.68	11.66	7.89	5.64
Prices 5->6	3.81	14.71	11.52	13.42	11.15	7.75	3.77
Prices 6->7	4.21	14.38	12.30	15.09	10.65	8.16	5.20
Prices 7->8	4.08	13.76	11.53	13.66	11.03	7.04	4.22
Prices 8->9	4.94	16.46	13.41	17.09	14.57	10.16	5.01
Prices 9->10	12.51	27.51	26.67	27.37	23.80	21.58	11.51

grew from eight categories, including an “NA” category, to ten with an additional “NA” category. The maximum number of tuples used for rule generation was limited to 50,000. With these modifications, rules were generated with the coverage levels shown in Table 4.

No set of stable attributes have notably higher coverage than any other consistently. Many rules were generated, for example, over 3,000 rules were formed using only base features moving from price level one to price level two. Adding a second flexible attribute that could be changed led to a very dramatic increase in the number of rules formed.

A sample G-Action rule from this set is “Artistic_style (Surrealistic) & Width (very low) && Bio_pos (very high) & Desc_pos (lower) ->Price(1), Artistic_style (Surrealistic) & Width (very low) && Bio_pos (avg) & Desc_pos (higher) ->Price(2)”. In this example, the before rule is “Artistic_style (Surrealistic) & Width (very low) && Bio_pos (very high) & Desc_pos(higher) ->Price(1)”, and the after rule is “Artistic_style (Surrealistic) & Width (very low) && Bio_pos (avg) & Desc_pos (higher) ->Price(2)”. This may be read as, when Artistic_style is Surrealistic and Width is very low if Bio_pos changes from very high to avg and Desc_pos changes from lower to higher then Price will change from 1 to 2.

To raise the accuracy of the models and gain better coverage, the clusters based on artist style discussed in a previous section were used to form rules. While ten clusters were generated, there were significant size discrepancies. One cluster is a very distant outlier with only a single prolific artist. Another contains over 1,759 artists. To generate the following sets of rules, the maximum number of works considered was 50,000, but fewer works were used if the cluster was very small.

The addition of personalization noticeably improves the coverage of the rule sets. This is not consistent for all clusters, such as the sixth edge-based cluster, likely due to the small number of represented artists. Clusters with more artists, such as the fourth edge-based cluster, have much better coverage. However, the pattern of lower coverage in middle price ranges and lack of dramatic shifts in coverage due to stable attributes is consistent. In Tables 5 and 6, the average coverage percentage of the ten clusters is displayed.

Some interesting patterns emerge when examining this information. The coverage is consistently higher when examining the average coverage levels of the partitions over a random selection of tuples. More interesting is how the different partitioning methods shift in their level of coverage depending on the price ranges being considered. The gray block partition has higher coverage for the middle range values, especially when moving from level 3 to level 4, level 5 to level 6, and level 6 to level 7. By contrast, the edge-blocking method of partitioning handles the upper and lower ranges better. The rules were generated separately,

Table 4 Coverage with new attributes

	BF	AWH	PAD	Col 1	Col 2	Gray	Block	RGB M	RGB S
1->2	64.05	64.95	65.37	66.31	67.98	64.61	65.97	64.81	65.09
2->3	22.95	24.03	24.13	25.29	28.31	24.03	26.23	24.25	24.27
3->4	19.36	20.58	20.34	21.60	23.40	19.94	22.16	20.48	20.84
4->5	11.37	11.99	11.95	13.51	15.04	11.67	13.71	11.85	12.27
5->6	9.48	9.96	10.14	11.65	13.47	9.70	10.71	10.06	9.98
6->7	8.96	9.78	9.52	11.06	13.18	9.48	9.98	9.66	9.68
7->8	10.04	10.52	11.00	11.88	13.82	10.18	11.10	10.56	10.46
8->9	10.94	11.40	11.48	13.47	15.43	11.26	11.94	11.72	11.28
9->10	23.20	23.96	24.44	26.42	28.92	23.84	25.24	24.12	24.02

Table 5 Aggregate coverage of gray block clusters

	BF	AWH	PAD	Col 1	Col 2	Gray	Block	RGB M	RGB S
1->2	76.29	78.53	80.49	77.60	78.20	78.70	79.24	78.88	79.64
2->3	51.94	57.14	59.39	54.42	56.04	55.34	57.77	56.88	57.63
3->4	51.11	55.87	54.86	53.20	54.96	54.66	55.07	55.45	55.84
4->5	41.27	49.19	49.90	44.28	47.64	46.60	47.74	47.51	46.72
5->6	36.25	45.05	45.06	39.21	42.92	44.07	45.35	43.00	45.05
6->7	37.82	44.48	44.55	40.70	43.50	43.13	44.89	44.41	43.88
7->8	39.70	49.21	47.76	44.47	47.37	46.49	48.33	48.08	48.11
8->9	43.40	50.93	49.55	46.36	49.30	49.01	50.87	50.25	50.63
9->10	63.28	69.45	68.51	65.96	69.33	67.61	68.88	68.90	68.29

so there may be slight differences in what each individual partition considers as set 3. The change in coverage shows potential for further exploration.

Rules were generated for all of the clusters with a few exceptions. As an example, the following rule was created for the first cluster of the brightness average, referred to as gray, “Subject(Flowers and plants) & Width(very high) && Bio_pos(avg) & Biowc(higher) ->Price(9)”, “Subject(Flowers and plants) & Width(very high) && Bio_pos(higher) & Biowc(very low) ->Price(10)”.

After noting the shifts in coverage depending on the partitioning method selected, the next step was to explore more methods of partitioning the dataset to examine what patterns would emerge.

Further rules were generated with slightly altered restrictions. The number of flexible attributes was restricted to one for this set of comparisons. The same set of stable and flexible attributes discussed previously was used, with the addition of the main three colors as stable attributes. The restrictions on confidence and support were repeated as in the previous example. All partitions have a maximum size of 50,000 tuples, but some, especially ones generated from edge or gray blocks, are smaller.

Several different methods of partitioning the dataset were explored. First, for the purposes of comparison, rules were generated with the mentioned method for four sets of 50,000 tuples. As shown previously, partitioning on visual features does have value. Partitions were generated using the gray and edge blocks as used above. Additionally, four

Table 6 Aggregate coverage of edge block clusters

	BF	AWH	PAD	Col 1	Col 2	Gray	Block	RGB M	RGB S
1->2	77.13	78.76	80.00	77.71	78.38	79.19	79.02	80.15	79.14
2->3	49.95	55.41	56.02	51.94	53.69	53.19	54.37	55.11	53.61
3->4	36.34	41.94	42.57	39.02	42.44	41.59	43.93	43.65	42.01
4->5	39.15	47.44	46.52	40.94	43.32	43.97	45.42	45.60	44.08
5->6	27.96	34.07	35.53	33.21	36.37	35.78	32.75	39.54	33.36
6->7	28.96	35.80	35.46	31.33	33.48	34.45	35.40	35.84	34.90
7->8	36.97	45.24	44.55	41.00	46.33	42.68	44.62	45.10	44.33
8->9	43.10	46.97	46.16	46.78	48.22	47.27	46.94	49.10	46.42
9->10	60.82	66.90	66.33	63.93	65.62	64.17	65.29	64.89	64.57

Table 7 Partition coverage averages

	Random	Edges	Follow	G. BI	E. BI	PAD	AWH
1->2	8.66	14.34	36.01	50.79	40.15	23.64	19.81
1->3	8.57	11.38	25.40	37.94	26.74	12.60	15.31
1->4	8.14	4.18	23.65	42.09	28.49	24.80	11.56
1->5	3.04	2.81	15.28	23.44	24.58	10.11	10.00
1->6	1.45	1.51	12.08	26.53	18.17	15.14	7.71
1->7	1.90	0.84	14.10	22.70	12.06	20.96	3.13
1->8	1.41	3.85	9.11	12.34	12.80	11.05	3.95
1->9	1.84	3.38	8.35	12.96	11.80	4.49	4.06
1->10	2.80	6.18	12.28	18.50	17.15	7.35	5.56
2->3	2.00	2.38	11.89	21.57	22.48	11.21	5.06
2->4	1.12	1.49	9.36	26.18	21.09	19.68	2.62
2->5	0.86	1.71	7.49	22.06	11.99	9.47	3.58
2->6	0.58	0.69	8.38	17.76	14.42	13.01	1.29
2->7	0.68	0.85	6.19	17.09	11.87	20.63	1.31
2->8	0.24	0.56	4.32	11.81	16.77	8.73	1.13
2->9	0.64	0.64	3.75	7.79	12.99	2.19	0.79
2->10	0.61	0.82	6.40	11.99	14.89	3.66	1.47
3->4	0.86	1.44	11.68	23.47	14.30	19.66	2.55
3->5	0.86	2.41	6.42	13.60	12.54	9.18	2.44
3->6	0.49	1.16	3.90	14.06	13.11	13.60	2.05
3->7	0.47	2.27	5.35	13.58	7.56	20.02	2.92
3->8	1.90	1.19	4.47	11.75	8.61	8.44	1.78
3->9	0.79	0.47	3.59	9.17	6.44	2.32	1.33
3->10	0.26	1.07	4.51	7.58	11.17	3.89	1.96
4->5	0.33	1.03	6.90	18.34	13.75	10.26	2.57
4->6	0.27	0.55	3.74	14.46	17.28	12.09	1.08
4->7	0.16	0.30	4.80	10.57	16.67	19.34	1.16
4->8	0.23	0.55	5.59	8.90	12.97	8.50	1.25
4->9	0.25	0.25	3.95	11.26	13.74	2.01	0.97
4->10	0.34	0.62	6.29	5.90	16.54	3.79	1.50
5->6	0.24	0.35	4.63	10.06	9.18	12.66	1.51
5->7	0.18	0.32	3.27	13.84	7.36	19.38	0.93
5->8	0.11	0.49	6.29	8.89	8.28	8.94	0.91
5->9	0.18	0.51	3.85	6.07	9.22	1.82	1.06
5->10	0.25	0.58	4.93	6.88	10.61	4.44	1.38
6->7	0.24	0.33	3.73	15.68	12.07	19.62	1.03
6->8	0.18	0.40	3.66	13.64	10.18	9.36	1.06
6->9	0.46	0.19	1.83	7.02	14.52	1.72	0.86
6->10	0.32	0.57	4.44	11.59	12.94	6.23	0.97
7->8	0.28	0.84	3.97	19.97	21.13	10.90	0.97
7->9	0.25	0.20	4.29	14.82	12.30	1.97	1.23

Table 7 (continued)

	Random	Edges	Follow	G. Bl	E. Bl	PAD	AWH
7->10	0.26	0.51	3.73	15.59	14.41	7.56	1.28
8->9	0.32	0.79	4.68	23.68	24.84	2.11	1.75
8->10	1.17	1.17	8.26	20.74	26.69	9.63	2.02
9->10	1.03	4.49	12.51	41.74	39.93	14.53	4.43

partitions were constructed based on the average number of edges in an individual work. This was added primarily to compare to the results when using the number of edges broken into blocks across an artist's posted body of work. To continue the use of visual features, five partitions were made using PAD (Pleasure, Arousal and Dominance) and five using AWH (Activity, Weight and Heat). This was made using k-means clustering on the PAD scores for an individual work. The process was completed using the AWH scores. Finally, six partitions were constructed based on the listed number of followers for an artist. This was meant to serve as a means of separating artists that are highly visible from those that are less widely viewed.

In Table 7, the average coverage of the partitions is shown. The number of followers partition is abbreviated as Follow, the gray blocks partition is abbreviated as G. Bl, and the edge blocks partition is abbreviated as E. Bl. Rules were generated for increases in price from all levels, rather than just to the next highest step.

When examining the aggregate scores, a number of patterns emerge. First, it is evident that using a form of partitioning produces superior coverage rather than using randomly selected tuples.

Secondly, while there is some variation regarding which partition generates the best coverage, the gray block method and the edge block method partitions produce the highest average coverage for most price pairs. Interestingly, the PAD partition and follower partitions also perform well but not as well as the gray block and edge block partitions. The edge partition performs considerably less well than the edge block partitions. Also, the pattern of whether the edge or gray block partition performs better on a given price pair does not always match the results from the previous set of rules generated.

6 Conclusions and future work

This work only begins to address the potential for the use of action rules to improve artist sales in the market for contemporary fine art. Many potential avenues for further research exist in the development of features for action rules.

In the initial set of results, the very high rate of coverage is promising. However, the low confidence level, while high for specific rules, is on average quite low. This issue may be attributable to allowing too many low confidence rules. It may also be due to the partitions of the feature sets being overly broad. However, low confidence for the set of G-Action rules is not a barrier to their being useful to an artist. These rules are extremely low cost and easily implemented. Stakeholders may want to try them even on the small chance that the rules will provide an improvement. The second and third set of rules discussed demonstrates that it is possible to significantly raise the average confidence in the rules, but it is at the expense of lowering the average coverage of the rule sets.

The addition of personalization by clustering visually similar artists increased the coverage of many of the derived rule sets. By grouping similar artists, the coverage was consistently increased. However, the question of what method of partitioning is most effective remains open for further exploration. Often the method of developing features to represent an artist's overall visual style had the highest level of coverage. The edge and gray block methods group artists, rather than artworks. Interestingly, when comparing the results of partitioning based on the percentage of edges in a work as opposed to the block method, it demonstrates that partitioning based on the overall traits of the artist produces superior rule sets to partitioning based solely on the features of an individual work. This pattern continues to be visible in the high coverage of the partitions based on the gray blocks method and the number of people following a given artist. One exception to this pattern is that partitions formed using the PAD cluster method often had a high coverage.

Certain partitions produced higher coverage on particular sets of prices. However, which partitions perform better in a specific situation remains an avenue for later research. Similarly, the exploration of whether the performance of individual partitions changes depending on the prices or the selection of attributes is an interesting topic for further exploration.

One strong avenue for potential future research is the development of more features for rule generation. The development of new stable visual attributes did not have a dramatic impact, so flexible attributes, especially flexible attributes with a higher cost, should be considered. With higher accuracy, stakeholders may become willing to make more drastic changes. The development of more flexible attributes, as well as exploring the attributes discussed here in combination, has potential value for stakeholders in the art market. In addition to expanding the list of artwork features, artists have characteristics that may serve as additional stable attributes. A greater exploration of larger scale career changes an artist could make has great potential for research.

Acknowledgements This research is supported by the National Science Foundation under grant IIP 1749105. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- Aggarwal, C. C. (2018). *Machine learning for text*. Cham: Springer International Publishing. <https://doi.org/10.1007/978-3-319-73531-3>.
- Artfinder.com (2019). <https://www.artfinder.com/>.
- Bailey, J. (2017). Machine Learning for Art Valuation. An Interview With Ahmed Hosny. <https://www.artnome.com/news/2017/12/2/machine-learning-for-art-valuation>.
- Beautiful Soup (2019). <https://www.crummy.com/software/BeautifulSoup/>.
- Berlin, B., & Kay, P. (1969). *Basic color terms: their universality and evolution*. Berkeley: University of California Press.
- Bird, S., Klein, E., & Loper, E. (2009). *Natural language processing with python*, 1st edn. Massachusetts: O'Reilly Media Inc.
- Canny, J. (1986). A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(6), 679–698. <https://doi.org/10.1109/TPAMI.1986.4767851>. <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4767851>.
- Colors by name (2020). <http://colormine.org/colors-by-name>.
- Galbraith, J., & Hodgson, D. (2018). Econometric fine art valuation by combining hedonic and repeat-sales information. *Econometrics*, 6(3), 32. <https://doi.org/10.3390/econometrics6030032>. <http://www.mdpi.com/2225-1146/6/3/32>.
- Hajja, A., Raš, Z. W., & Wiczorkowska, A.A. (2014). Hierarchical object-driven action rules. *Journal of Intelligent Information Systems*, 42(2), 207–232. <https://doi.org/10.1007/s10844-013-0291-2>.

- Hosny, A., Huang, J., & Wang, Y. (2014). The Green Canvas. <http://ahmedhosny.github.io/theGreenCanvas/>.
- Hutto, C., & Gilbert, E. (2015). VADER: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the 8th international conference on weblogs and social media, ICWSM 2014*.
- Kalanat, N., & Khanjari, E. (2020). Action extraction from social networks. *Journal of Intelligent Information Systems*, 54(2), 317–339. <https://doi.org/10.1007/s10844-019-00551-2>.
- Kang, D., Shim, H., & Yoon, K. (2018). A method for extracting emotion using colors comprise the painting image. *Multimedia Tools and Applications*, 77(4), 4985–5002. <https://doi.org/10.1007/s11042-017-4667-0>.
- Labrecque, L. I., & Milne, G. R. (2012). Exciting red and competent blue: the importance of color in marketing. *Journal of the Academy of Marketing Science*, 40(5), 711–727. <https://doi.org/10.1007/s11747-010-0245-y>.
- Lindsey, D. T., & Brown, A. M. (2006). Universality of color names. *Proceedings of the National Academy of Sciences*, 103(44), 16608–16613. <https://doi.org/10.1073/pnas.0607708103>.
- Liu, D., & Woodham, D. (2019). Using AI to Predict Rothko paintings' auction prices. <https://www.artsy.net/article/artsy-editorial-ai-predict-mark-rothko-paintings-auction-prices>.
- Lombardi, T. E. (2005). The classification of style in fine-art painting. Ph.D. thesis. <https://librarylink.uncc.edu/login?url=https://search-proquest-com.librarylink.uncc.edu/docview/305390005?accountid=14605>. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2020-02-21.
- Mardini, M. T., & Raś, Z. W. (2019). Extraction of actionable knowledge to reduce hospital readmissions through patients personalization. *Information Sciences*, 485, 1–17. <https://doi.org/10.1016/j.ins.2019.02.006>. <https://www.sciencedirect.com/science/article/pii/S0020025519301008>.
- Mehrabian, A. (1978). Measures of individual differences in temperament. *Educational and Psychological Measurement*, 38(4), 1105–1117. <https://doi.org/10.1177/001316447803800431>.
- Nekvapil, V. (2009). Using the ac4ft-miner procedure in the medical domain [online] [cit. 2020-01-03] <https://theses.cz/id/0abafc/?lang=en>.
- Ou, L., Luo, M. R., Woodcock, A., & Wright, A. (2004). A study of colour emotion and colour preference. part i: Colour emotions for single colours. *Color Research & Application*, 29(3), 232–240. <https://doi.org/10.1002/col.20010>.
- Pawlowski, C., Gelich, A., & Raś, Z.W. (2019). Can We Build Recommender System for Artwork Evaluation? In R. Bembek, Skonieczny, G. Protaziuk, M. Kryszkiewicz, & H. Rybinski (Eds.) *Intelligent Methods and Big Data in Industrial Applications* (pp. 41–52). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-77604-0_4.
- Powell, L., Gelich, A., & Ras, Z.W. In T. Mihálydeák, F. Min, G. Wang, M. Banerjee, I. Düntsch, Z. Suraj, & D. Ciucci (Eds.) (2019). *Developing artwork pricing models for online art sales using text analytics*. Cham: Springer International Publishing.
- Powell, L., Gelich, A., & Ras, Z.W. (2020). Applying analytics to artist provided text to model prices of fine art. In *Complex pattern mining: New challenges, methods and applications, studies in computational intelligence*, (Vol. 880 pp. 189–211). Springer.
- Powell, L., Gelich, A., & Ras, Z.W. (2020). Art innovation systems for value tagging. In *Encyclopedia of organizational knowledge, administration, and technologies. IGI Global*. <https://www.igi-global.com/book/encyclopedia-organizational-knowledge-administration-technology/242894>.
- Powell, L., Gelich, A., & Ras, Z.W. In D. Helic, G. Leitner, M. Stettinger, & Z. W. Ras (Eds.) (2020). *The construction of action rules to raise artwork prices*. Austria: Springer.
- Ras, Z. W., & Wiczorkowska, A. In D. A. Zighed, J. Komorowski, & J. Żytkow (Eds.) (2000). *Action-rules: How to increase profit of a company*. Berlin: Springer.
- Rauch, J., & Šimůnek, M. In J. Rauch, Z. W. Raś, P. Berka, & T. Elomaa (Eds.) (2009). *Action rules and the guha method: Preliminary considerations and results*. Berlin: Springer.
- Rauch, J., Šimůnek, M., Berka, P., Černý, Z., Dolejší, P., Karban, T., Kejkula, M., Lin, V., Máša, P., Spáčil, P., Strossa, P., Šebek, M., Šlesinger, J., Šubrt, M., Turunen, E., & Vanžura, F. (2019). <https://lispminer.vse.cz/index.html>.
- Rawlins, C., & Johnson, P. (2007). Selling on eBay: Persuasive communication advice based on analysis of auction item descriptions. *Journal of Strategic E-commerce*, 5(1&2), 75–81.
- Roberson, D., & Hanley, J. (2007). Color vision: Color categories vary with language after all. *Current Biology*, 17(15), R605–R607. <https://doi.org/10.1016/j.cub.2007.05.057>. <http://www.sciencedirect.com/science/article/pii/S0960982207014819>.
- Russell, J. A., & Mehrabian, A. (1977). Evidence for a three-factor theory of emotions. *Journal of Research in Personality*, 11(3), 273–294.

- Saleh, B., & Elgammal, A. (2015). Large-scale classification of fine-art paintings: Learning the right metric on the right feature. arXiv:1505.00855.
- Selenium (2019). <https://www.seleniumhq.org/>.
- Sharma, G., Wu, W., & Dalal, E.N. (2004). The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application*, 30(1), 21–30. <https://doi.org/10.1002/col.20070>.
- Tarnowska, K., Daniel, L., & Ras, Z.W. (2020). *Recommender System for Improving Customer Loyalty, 1st ed. 2020 edn. No. 55 in Studies in Big Data. Springer International Publishing : Imprint. Cham: Springer*. <https://doi.org/10.1007/978-3-030-13438-9>.
- Tarnowska, K. A., Ras, Z. W., & Jastreboff, P.J. (2017). Mining for actionable knowledge in tinnitus datasets. In G. Wang, A. Skowron, Y. Yao, D. Ślezak, & L. Polkowski (Eds.) *Thriving Rough Sets: 10th Anniversary - Honoring Professor Zdzisław Pawlak's Life and Legacy & 35 Years of Rough Sets* (pp. 367–395). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-54966-8_18.
- Taylor, G. (2014). python-colormath.
- Team, O. (2017). OpenCV.
- The Hiscox Online Art Trade Report (2018). Colors by Name. Tech. rep., ArtTactic 2018. <https://arttactic.com/product/hiscox-online-art-trade-report-2018/>.
- Tsay, L. S., & Raś, Z. W. (2005). Action rules discovery: system DEAR2, method and experiments. *Journal of Experimental & Theoretical Artificial Intelligence*, 17(1-2), 119–128. <https://doi.org/10.1080/09528130512331315855>. Publisher: Taylor & Francis.
- Tzacheva, A. A., Bagavathi, A., & Ayila, L. (2017). Discovery of action rules at lowest cost in spark. In *2017 IEEE international conference on data mining workshops (ICDMW)* (pp. 87–94). New Orleans: IEEE. <https://doi.org/10.1109/ICDMW.2017.173>. <http://ieeexplore.ieee.org/document/8215648/>.
- Tzacheva, A. A., & Raś, Z. W. (2005). Action rules mining. *International Journal of Intelligent Systems*, 20(7), 719–736. <https://doi.org/10.1002/int.20092>.
- Valdez, P., & Mehrabian, A. (1994). Effects of color on emotions. *Journal of Experimental Psychology: General*, 123(4), 394–409. <https://doi.org/10.1037/0096-3445.123.4.394>.
- Velthuis, O. (2005). *Talking prices: symbolic meanings of prices on the market for contemporary art*. Princeton: Princeton University Press. <http://www.jstor.org/stable/j.ctt4cgd14>.

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