

Quantitative Analysis of Art Market Using Ontologies, Named Entity Recognition and Machine Learning: A Case Study

Dominik Filipiak¹, Henning Agt-Rickauer², Christian Hentschel², Agata Filipowska¹, Harald Sack²

¹ Department of Information Systems,
Poznań University of Economics,
Al. Niepodległości 10, 61-875, Poznań, Poland
{dominik.filipiak, agata.filipowska}@kie.ue.poznan.pl,
WWW home page: <http://kie.ue.poznan.pl>

² Hasso-Plattner-Institut
Prof.-Dr.-Helmert-Straße 2-3
14482 Potsdam, Germany
{henning.agt-rickauer, christian.hentschel, harald.sack}@hpi.de,
WWW home page: <http://hpi.de>

Abstract. In the paper we investigate new approaches to quantitative art market research, such as statistical analysis and building of market indices. An ontology has been designed to describe art market data in a unified way. To ensure the quality of information in the knowledge base of the ontology, data enrichment techniques such as named entity recognition (NER) or data linking are also involved. By using techniques from computer vision and machine learning, we predict a style of a painting. This paper comes with a case study example being a detailed validation of our approach.

Key words: art market, Semantic Web, linked data, machine learning, information retrieval, alternative investment, digital humanities

1 Introduction

Due to the constantly growing interest in the alternative investment area, the art market has become a subject of numerous studies. By publishing sales data, many services and auction houses provide a basis for further research in terms of the latest data analysis trends. A closer look at available data shows missing information or inconsistency in many cases, though. An intense effort (see the next section) has been observed among scientists carrying out research on auction markets, especially in the field of constructing indexes. To the best of our knowledge, the problem of data quality has not been raised often in that field. To tackle this issue, we propose mixing standard econometric analysis with the usage of the latest solutions known from the computer science domain.

To address the issue of data quality (influencing the indexes being built) regarding paintings sold in auction houses, a framework for quantitative art market research using ontologies, named entity recognition (NER) and machine learning has been introduced in this paper. This can be reached by combining various data sources and by employing recent developments in data science, such as semantic annotation of text and automated visual analysis. Visualising various trends or indicating economic incentives influencing the market is a possible outcome of this framework.

The remainder of this paper is organised as follows. The next section is a review of literature relevant to our framework. After that, we shed a light on the approach by explaining methods behind it. Since this paper introduces a new art market ontology, the following section is fully devoted to it. The next section contains a case study, being a detailed example of the usage of our method on a single Monet's painting. A short summary with future work closes this paper.

2 Related work

Art Market Research. Quantitative analysis of art market data has been a subject matter for many studies, such as portfolio diversification [1] and measuring the volatility of the market [2]. The vast majority of conducted research relies on building art market indices, which are to show price movements of a standard artwork in a given period (typically on a year basis). The most popular types of indices for art market research are built on top of a hedonic regression (HR) and repeat-sales regression (RSR) [3]. No matter which type has been chosen, the problem of availability of data always arises. Models based on RSR rely on prices of artworks sold at least twice [4]. The famous Mei Moses Fine Art Index¹, an example of the employment of RSR, considers only lots sold in the two biggest auction houses - Christie's and Sotheby's. Due to its nature, repeat-sales regression can operate on significantly smaller datasets. Art is considered to be a long-term investment, so it may be a challenging task to collect a decent dataset of lots sold at least twice - in some auction houses it is even strictly forbidden. Models based on HR take into account all sold lots and rely on the relation between (typically) a hammer price of a given lot and all its features. An example of this model is the Two-Step Hedonic Regression [5]. HR-based calculations are prone to feature selection bias, however. To overcome the limitation of taking into account only sold lots, a probit estimation may be employed [6].

Knowledges Bases for Art Market Data. Our work proposes the use of semantic web technologies and linked open data principals [7] to store information about artworks sold in auctions. While the use of ontologies in the auction domain is rather small, different works exist in the art domain. The openART ontology [8] was developed to describe a research dataset about London's artworld. Its focus is on events related to artworks in that time, but some parts deal with

¹ <http://www.artasanasset.com>

sales data. Europeana is a research initiative to provide access to millions of cultural heritage objects (e.g., books, films, paintings). The Europeana Data Model (EDM) is a semantic web-based framework to describe cultural heritage objects. The DBpedia ontology [9] also covers parts of artwork descriptions and can be used in art research [10]. Following the linked data principals we link these ontologies where necessary and reuse existing properties. In order to populate our auction knowledge base we apply state-of-the-art information extraction methods [11], in particular named-entity recognition and disambiguation [12, 13, 14].

Computer Vision Tools for Art Market Data. Typically, computer vision tools focus on the classification of real-world objects and scenes (e.g. "sunset", "faces" and "car"). Classification of paintings into different styles and art epochs has gained relatively little attention in the past. Several authors have evaluated the aesthetic quality of photographs. Datta et al. [15], designed algorithms to extract visual characteristics from images that can be used to represent concepts such as colourfulness, saturation, rule-of-thirds, and depth-of-field, and evaluated aesthetic rating predictions on photographs. The same methods were applied to a small set of Impressionist paintings [16]. Murray et al. [17] introduced the Aesthetic Visual Analysis (AVA) dataset, annotated with ratings by users of a photographic skill competition website. Marchesotti and Peronin [18] proposed a method to predict aesthetic ratings by using data gathered from user comments published on the website. The attributes they found to be informative (e.g., "lovely photo," "nice detail") are not specific to image style though. Several authors have developed approaches to automatically classify classic painting styles, including [19][20]. These works consider only a handful of styles (less than ten apiece), with styles that are visually very distinct, e.g., Pollock vs. Dalí. These datasets comprise less than 60 images per style. Mensink [21] provides a larger dataset of artworks, but does not consider style classification. Finally, the authors in [22] publish a dataset of 85,000 paintings annotated with 25 genre labels ranging from Renaissance to modern art (e.g. "Baroque", "Rococo" and "Cubism"). The authors present an approach for automatic classification of paintings into art epochs using latest computer vision methods and report per-class accuracies ranging from 72% ("Symbolism", "Expressionism", "Art Nouveau") to 94% ("Ukiyo-e", "Minimalism", "Color Field Painting").

3 Our Approach

Since numerous auction house websites are publishing data about sold lots in a well-structured manner, this information may be gathered using web crawlers. The quality and availability of information about a given artwork is the subject matter of this paper. There is a so-called *garbage in, garbage out* rule in data analysis and the art market is not different. Therefore, the process of collecting the data should be planned carefully. This should involve data refinement, due to the possible errors in collected data. A typo in the author's name may result in assigning artworks to two different creators. These problems can be mitigated by data cleansing techniques. On the other hand, some lots have missing values

(for example, the style of an artwork is often not provided). We explain how to deal with this problem in further sections. Having gathered the data, standard linear regression (estimated by Ordinal Least Squares) may be used to indicate the relation between painting's price and its qualities:

$$\ln P_{it} = \alpha + \sum_{j=1}^z \beta_j X_{ij} + \sum_{t=0}^{\tau} \gamma_t D_{it} + \varepsilon_{it} \quad (1)$$

where $\ln P_{it}$ is the natural logarithm of a price of a given lot $i \in \{1, 2, \dots, N\}$ at time $t \in \{1, 2, \dots, \tau\}$; α (intercept), β and γ are regression coefficients for estimated characteristics included in the model. X_{ij} represents hedonic variables (numeric and dummies, explained in the next paragraph) included in the model, whereas D_{it} stands for time dummy variables. The last parameter, ε_{it} represents the error term.

Suppose there is a non-numeric parameter, like the presence of the signature. This information can't be included in a model in that form. To overcome this problem and to not lose the information, so-called dummy variables are introduced. They are characterized by the amount of possible levels l (in this case $l = 2$). A well-defined ontology (see section 4) can help to summarize the possible levels. For the each level, an explanatory variable takes "1" if the condition is true (there is a signature) or "0" otherwise (lack thereof). Coefficients of dummy variables equal the average difference in impact on the model between these cases. Another example is the year of sale (D_{it}). For each sold lot, this variable is equal to "1" only if a given painting i was sold in a period t (otherwise it is equal to "0"). Supposing only paintings sold in 2014, 2015 and 2016 are considered, those years become dummy variables for representing time in the equation. Supposing a given observation is sold in 2016, only the last one takes the value of "1", the rest is equal to "0".

Although it will be hard to have more predictors than observations (so called $p > n$ setup) in our case, the correct selection of used explanatory variables is crucial and the overall result depends on it. Measures like statistical significance can be helpful here. It has to be remembered that one must not include l levels in the model, because this leads to multicollinearity. A simple way to avoid this situation is to use $l - 1$ levels.

The most important dummy variable is the one representing the year of sale. By calculating the coefficients, one can build an index to *measure* the art market and compare lots sold in a given year to other forms of assets (such as stocks). At the same time, this value is *separated* from the painting's qualities, such as its author or price. This allows to measure the impact of the given year on the hammer price without examining sold lots and regardless of their features (which typically involves experts with domain knowledge). A standard way to calculate the hedonic price index for a period t is:

$$Index_t = e^{\gamma_t} \quad (2)$$

where γ s are taken from the equation (1). A more sophisticated way to calculate the index, like the Two-Step Hedonic Index considers following calculations:

$$Index_{t+1} = \frac{\prod_{i=1}^n (P_{i,t+1})^{1/n} / \prod_{i=1}^m (P_{i,t})^{1/m}}{\exp \left[\sum_{j=1}^z \beta_j \left(\sum_{i=1}^n \frac{X_{ij,t+1}}{n} - \sum_{i=1}^m \frac{X_{ij,t}}{m} \right) \right]} \quad (3)$$

No matter which way of building a hedonic index is used, the quality of information (i.e. observations in our case) always seems to play an important role. A higher number of explanatory variables allows examining the importance of each one more accurately. This also leads to a higher coefficient of determination (R^2) – it is not a good indicator of the quality of the model, though².

To improve the generalization of the model, the number of explanatory variables should be reduced. This can be achieved by e.g. univariate feature selection, recursive feature elimination or hill climbing solutions. Performing those methods makes sense only if a decent number of variables is available. As it was stated previously, some observations suffer from a lack of data. Therefore, we decided to extend the amount of available information.

The goal of data collection is the processing of different kinds of data sources and the extraction of data values and descriptions regarding lots sold in auctions. The extraction is performed in two steps. First, we crawl auction house websites and download all data on sold lots. Values that can be directly extracted using standard techniques (e.g., regular expressions) are stored as RDF data using our art auction ontology. Secondly, during data enrichment, extracted information are processed with named entity recognition and natural language processing in order to link them to external data sources, such as DBpedia or Europeana. This enables the acquisition of additional variables for art market analysis.

Finally, information automatically extracted from the visual domain may provide further information on the art object at hand. By using techniques from computer vision and machine learning we predict the style of a painting in order to further enrich extracted metadata. Our approach uses features extracted from a Deep Convolutional Neural Network (CNN) trained on the well-know ImageNet [23] dataset. The model architecture is based on the CNN architecture winning the 2012 ImageNet challenge [24] with some minor modifications (for the sake of increased training performance, the size of the full connect layer is reduced from 4,096 to 2,048 neurons). We follow the approach presented in [25] by taking the pen-ultimate (fully connected) layer of the trained CNN as a feature extraction layer yielding a 2,048 dimensional feature vector per image which is $L2$ -normalized. Model training is conducted in a One-vs.-Rest-approach using a Support Vector classifier with a linear kernel. We optimize the cost parameter in a three-fold stratified cross-validation.

Groundtruth data is taken from the “WikiArt.org – Encyclopedia of fine arts” project³ which contains a collection of paintings from different art movements, ranging from Renaissance to modern art. All paintings are manually labeled according to the respective art movement. We used the dataset to generate groundtruth data by randomly selecting 1,000 images for training and 50 images

² Due to the problem of overfitting.

³ WikiArt.org – Encyclopedia of fine arts, <http://www.wikiart.org>

for testing purposes. Since the number of images in the WikiArt.org paintings collection varies a lot (e.g., for some art movements less than 100 images are available while others provide more than 10,000 images), only those classes were selected that are supported by at least 1,050 images (we obtained a total of 22 classes).

4 Ontology

We designed an art auction ontology that is able to capture information gathered from different sources, such as auction house websites, linked data sets, and other art market databases. Currently, the ontology concentrates on the description of paintings that appeared in auctions, but the model is extensible for different kinds of artworks. We use Web Ontology Language (OWL) and Resource Description Framework Schema (RDFS) to specify the schema.

Figure 1 shows part of the ontology schema that is responsible for the sales data. The central entity is the *Lot*. In a lot a particular item is offered for sale, in our case *Artwork* items. The *Sale* event is organized by an *AuctionHouse* at a certain location. We link to the vCard ontology and reuse it for address description. The most important information for price indices is the *date* of the

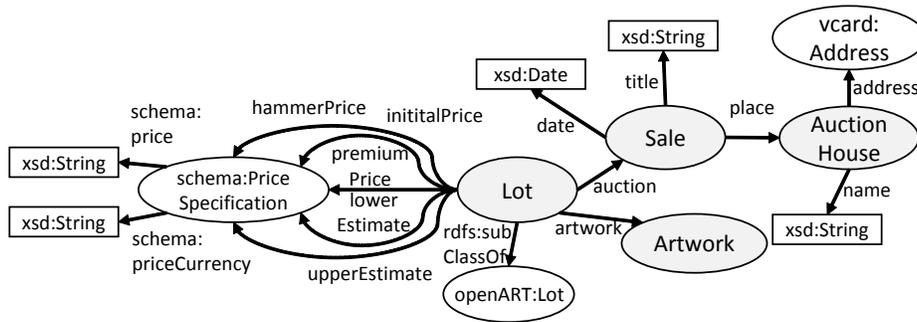


Fig. 1. Schema of the art auction ontology – sales data (excerpt).

auction and final selling price. A lot is associated with different types of prices: *upperEstimate* and *lowerEstimate* are an appraisal of what price the lot will fetch. The *initialPrice* is the suggested opening bid determined by the auctioneer. *hammerPrice* is the final price at which a lot was sold (without any fees and taxes), and *premiumPrice* includes fees and taxes, such as buyer’s premium and sales tax. *PriceSpecification* is done using the schema.org vocabulary (price value and currency). Additionally, the lot class is linked to the openART ontology [8] that provides additional properties for the artwork domain.

The lot has a reference to the artwork as shown in Figure 2. This part of the ontology is responsible for modeling the auction items and their creators in detail. The *Artwork* class is the general description for different types of artworks,

and has a creation date (in case it is known). The artwork class is linked to the DBpedia ontology [9] enabling reuse of its properties (*dbo* prefix). As mentioned before, our work currently operates on paintings data, thus one specialization *Painting* exists. Other kinds of specific artworks can be added later using the subclass relation. A painting is described by several datatype properties, such as

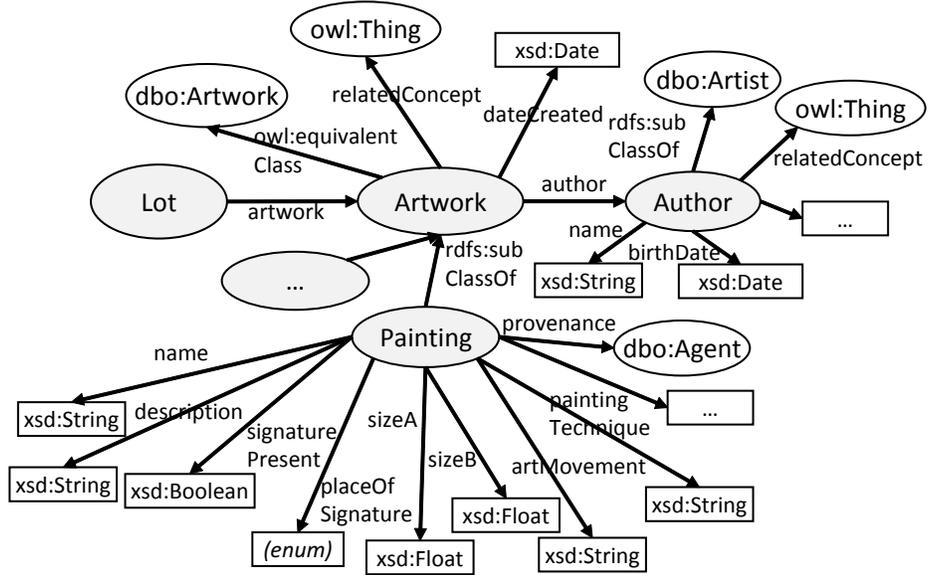


Fig. 2. Schema of the art auction ontology – artwork data (excerpt).

name and a natural language description. The ontology is capable of describing the painting’s size and information on the author’s signature (present or not, and the location of the signature). An item can be categorized using the *art-Movement* and *paintingTechnique* properties. We also provide properties to store art movement information that was classified automatically by image analysis (*artMovementClassified*, and *artMovementClassifiedConfidenceValue*, not shown in the Figure). A painting can be associated with provenance information describing previous owners, such as organizations and persons. Therefore, we link to the DBpedia *Agent* class. Similarly, exhibitions are described in which the painting appeared (not shown in the Figure).

An artwork item is usually associated with an *Author*. We inherit from the DBpedia Artist class that offers a wide range of person properties, such as birth date, death date and nationality, as well as artwork-specific properties (e.g., *influencedBy*, *works*). In case the author is not known the ontology offers an origin property to be able to describe nationality or other indicators of the painting’s source. Both the author and the artwork can be linked to other entities using the *relatedConcept* object property in order to store additional information, such as important cities, other artists or related artworks.

5 Case Study

This section applied our method on a real world example. We have chosen Monet's *Le Parlement Soleil Couchant*⁴, sold at Christie's New York for \$40,485,000 (hammer price plus buyer's premium) in 2015.

5.1 Data Collection

In our example, we illustrate the extraction of data values from Christie's website. We use techniques from information extraction [26] in order to transform data provided in an unstructured manner into a structured machine-readable format. In particular boilerplate removal (deletion of unnecessary HTML content), filtering, tokenization and regular expressions are applied to build a wrapper for the auction data. Figure 3 depicts an excerpt of the RDF visual representation that could be extracted from the example website. It shows that the painting "Le Parlement, soleil couchant" by Claude Monet was sold on 11 May 2015 for a premium price of USD 40,485,000.

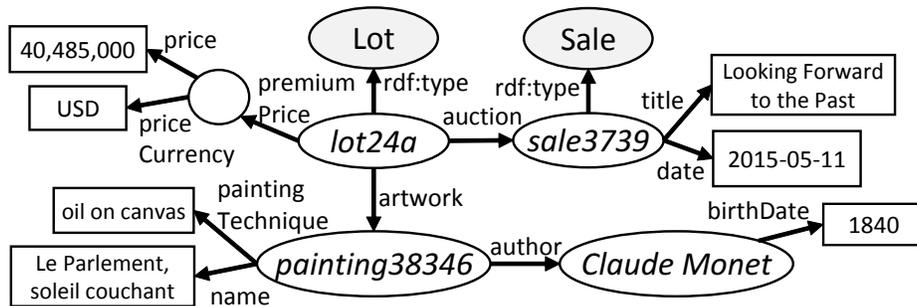


Fig. 3. Auction data in RDF format extracted from the auction website (excerpt).

5.2 Data Enrichment

We pursue two strategies to obtain additional data on artworks and artists. First, we employ named entity recognition [13] on text values and natural language descriptions. The identified entities are then disambiguated [14] and linked to other semantic data sets. Secondly, we use image classification to retrieve the style (art movement) of the painting which is often not provided.

Named Entity Recognition The task of named entity recognition (NER) concerns the identification of words or phrases in natural language texts that refer to names. These names are associated with categories (e.g., person, location). NER on the already obtained name of the author and painting is relatively

⁴ <http://christies.com/lotfinder/paintings/claude-monet-le-parlement-soleil-couchant-5895978-details.aspx>

easy and used to verify the data collection task. We also apply NER on the long description of the lot (lot notes). From that we learn that London is the most important city related to the lot (mentioned 27 times), as well as the River Thames (natural feature, 10 times). Two organizations have been found: St. Thomas’s Hospital and Saint James’s Hospital (mentioned 16 times, also indirectly just with "hospital".) Paul Durand-Ruel could be identified as the most important person related to Monet.

Data Linking. Whereas NER can only identify position and category of a named entity, named entity linking performs disambiguation of the name and connects it to a specific instance in an existing knowledge base. Currently, we link to the DBpedia knowledge base [9]. In our case study, we are able to link the artist value (Claude Monet) to the DBpedia resource `Claude_Monet`⁵. The link allows the retrieval of additional structured data values (e.g., exact birth and death date, nationality, or author description). Furthermore, we are able to query additional information from the knowledge base using pre-defined properties from the DBpedia ontology and SPARQL language. For example, using the `dbo:movement` property we retrieve "Impressionism". In case this information or the artist is not available in DBpedia, we apply image analysis (see next Section). The other identified named entities (e.g., London, Thames) are linked to the lot using the `relatedConcept` property of our art auction ontology.

Art Movement Classification. Figure 4 shows the scores of our art movement classifiers when applied on Monet’s *Le Parlement Soleil Couchant*. For reasons of brevity we visualize only positive classification scores. As can be seen, the SVM classifier for the art movement “Impressionism” produces the highest classification score.

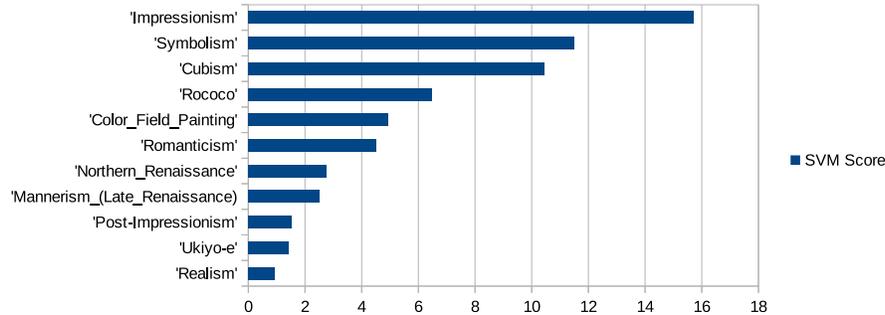


Fig. 4. Classifier scores for positively classified epochs.

5.3 Data Analysis

Having enriched and refined a dataset, it is possible to perform the analysis en route to build the indices. The natural logarithm of the realized price

⁵ http://dbpedia.org/resource/Claude_Monet

(40,485,000) becomes the dependent variable. The currency is omitted, provided that all of the observations will be presented in USD. Table 1 shows the excerpt of possible explanatory variables for Equation (1). *Signature* is an example of a dichotomous variable. Nominal variables, such as *Author* can't be employed in the equation in the original form. Therefore, *Claude_Monet* becomes a dummy variable (equal to one) among expected other ones, for example *Pablo_Picasso* which takes zero as the value. Ordinal variables, such as *Year_of_sale* may be transferred to the set of dummies representing each year. Concepts found by NER in text descriptions are also included as dummies (like *London*). These ones related directly to the painter (such as Paul Durand-Ruel) can't be included, due to the presence of the *Author* variable group, which may lead to multicollinearity. Because of the possibly large number of variables, a selection process is a must.

Table 1. Explanatory variables (excerpt) for *Le parlement, soleil couchant*

Group	Variable	Value
Author	Claude_Monet	1
	Pablo_Picasso	0
	Claude_Monet	0
Dichotomous	Signature	1
Continuous	Size_A	81.3
	Size_B	93
Year	XIX	1
Technique	Oil_on_canvas	1
	Tempera	0
	London	1
Concept	River_Thames	1
	Berlin	0

Such an approach applied for various artworks and various artists would enable us to construct indices describing markets, and lead to comparisons of markets in case of choosing between investments. Gathering such values enables also for other analyses which are subject to further research.

6 Conclusions and future work

This paper presented a method for quantitative art market research using ontologies, named entity recognition and machine learning. By connecting concepts to a given artwork or using linking to another data source we presented the value of data enrichment in this particular case. To suit our needs for data storing and provide a possibility to create a standard for presenting art market data, the art market ontology was developed. As for overcoming problems with missing style

information, an artwork classifier was prepared, trained and employed in our research. A one-example use case illustrates the usage of the proposed solution. The next steps consider gathering data as a base for further research. Data refinement seems to be a challenge after that. It has to be decided which variables are playing a significant role regarding the description of a considered lot. This can be achieved by statistical analysis. Finally, we expect to build precise art market indices on a large-scale data sample.

Future work may consider extending the ontology to the other artworks (e.g. sculptures). This might even enable a possibility to carry out research on all of lots sold in auction houses – not only paintings. Adding other data enrichment sources is also a feature worth considering. Using an ontology enables the transfer of cultural knowledge or artworks itself more efficiently between different countries, so the further steps might consider publishing data and making a translational effort. An extensive study of the behavior or correlation of different indices on an enriched dataset is the next possibility.

References

1. Jurevičiene, D., Savičenko, J.: Art Investments for Portfolio Diversification. *Intellectual Economics* **6**(2) (2012) 41–56
2. Bocart, F.Y.R.P., Hafner, C.M.: Volatility of Price Indices for Heterogeneous Goods with Applications to the Fine Art Market. *Journal of Applied Econometrics* **30**(2) (mar 2015) 291–312
3. Etro, F., Stepanova, E.: The market for paintings in paris between rococo and romanticism. *Kyklos* **68**(1) (2015) 28–50
4. Ginsburgh, V., Mei, J., Moses, M.: The Computation of Prices Indices. In: *Handbook of the Economics of Art and Culture*. Volume 1. Elsevier (2006) 947–979
5. Kräussl, R., van Elsland, N.: Constructing the true art market index : a novel 2-step hedonic approach and its application to the german art market (2008)
6. Collins, A., Scorcu, A., Zanola, R.: Reconsidering hedonic art price indexes. *Economics Letters* **104**(2) (aug 2009) 57–60
7. Heath, T., Bizer, C.: *Linked Data: Evolving the Web into a Global Data Space*. 1st edn. Morgan & Claypool (2011)
8. Allinson, J.: Openart: Open metadata for art research at the tate. *Bulletin of the American Society for Information Science and Technology* **38**(3) (2012) 43–48
9. Lehmann, J., Isele, R., Jakob, M., Jentzsch, A., Kontokostas, D., Mendes, P.N., Hellmann, S., Morsey, M., van Kleef, P., Auer, S., et al.: Dbpedia—a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web* **6**(2) (2015) 167–195
10. Filipiak, D., Filipowska, A.: DBpedia in the Art Market. In Abramowicz, W., ed.: *Business Information Systems Workshops*. Volume 228 of *Lecture Notes in Business Information Processing*. Springer International Publishing, Cham (2015) pp 321–331
11. Etzioni, O., Fader, A., Christensen, J., Soderland, S., Mausam, M.: Open information extraction: The second generation. In: *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume One*. IJCAI'11, AAAI Press (2011) 3–10

12. Filipiak, D., Węcel, K., Filipowska, A.: Semantic Annotation to Support Description of the Art Market. In: Joint Proceedings of the Posters and Demos Track of 11th International Conference on Semantic Systems - SEMANTiCS2015 and 1st Workshop on Data Science: Methods, Technology and Applications (DSci15) co-located with the 11th International Conference on Sema, Vienna, CEUR Workshop Proceedings (2015) 51–54
13. Usbeck, R., Röder, M., Ngomo, A.N., Baron, C., Both, A., Brümmer, M., Ceccarelli, D., Cornolti, M., Cherix, D., Eickmann, B., Ferragina, P., Lemke, C., Moro, A., Navigli, R., Piccinno, F., Rizzo, G., Sack, H., Speck, R., Troncy, R., Waitelonis, J., Wesemann, L.: GERBIL: general entity annotator benchmarking framework. In: Proceedings of the 24th International Conference on World Wide Web, WWW 2015, Florence, Italy, May 18-22, 2015. (2015) 1133–1143
14. Daiber, J., Jakob, M., Hokamp, C., Mendes, P.N.: Improving efficiency and accuracy in multilingual entity extraction. In: Proceedings of the 9th International Conference on Semantic Systems (I-Semantics). (2013)
15. Datta, R., Joshi, D., Li, J., Wang, J.Z.: Studying aesthetics in photographic images using a computational approach. In: Computer Vision–ECCV 2006. Springer (2006) 288–301
16. Li, C., Chen, T.: Aesthetic visual quality assessment of paintings. Selected Topics in Signal Processing, IEEE Journal of **3**(2) (2009) 236–252
17. Murray, N., Marchesotti, L., Perronnin, F.: Ava: A large-scale database for aesthetic visual analysis. In: Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on, IEEE (2012) 2408–2415
18. Luca Marchesotti (Xerox), F.P.X.X.G.: Learning beautiful (and ugly) attributes. In: Proceedings of the British Machine Vision Conference, BMVA Press (2013)
19. Keren, D.: Painter identification using local features and naive bayes. In: Pattern Recognition, 2002. Proceedings. 16th International Conference on. Volume 2., IEEE (2002) 474–477
20. Shamir, L., Macura, T., Orlov, N., Eckley, D.M., Goldberg, I.G.: Impressionism, expressionism, surrealism: Automated recognition of painters and schools of art. ACM Transactions on Applied Perception (TAP) **7**(2) (2010) 8
21. Mensink, T., van Gemert, J.: The rijksmuseum challenge: Museum-centered visual recognition. In: Proceedings of International Conference on Multimedia Retrieval, ACM (2014) 451
22. Karayev, S., Trentacoste, M., Han, H., Agarwala, A., Darrell, T., Hertzmann, A., Winnemoeller, H.: Recognizing image style. In: Proceedings of the British Machine Vision Conference, BMVA Press (2014)
23. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. International Journal of Computer Vision (2015)
24. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems 25. Curran Associates, Inc. (2012)
25. Chatfield, K., Simonyan, K., Vedaldi, A., Zisserman, A.: Return of the devil in the details: Delving deep into convolutional nets. CoRR **abs/1405.3531** (2014)
26. Cimiano, P., Mädche, A., Staab, S., Völker, J.: Ontology learning. In: Handbook on Ontologies. (2009) 245–267