
Unveiling fake artworks: using automatic classification based on freely available information

Arts and the
Market

Liliana Wuffli

University of Lausanne, Lausanne, Switzerland, and

André Berchtold

Institute of Social Sciences, University of Lausanne, Lausanne, Switzerland

Received 1 December 2025

Revised 25 February 2026

19 April 2026

Accepted 11 August 2026

Abstract

Purpose – In the dynamic art market, fake artworks pose a significant threat to cultural heritage preservation. This article aims to introduce a novel approach to enhancing art fraud prevention through the application of computer-based art classification techniques.

Design/methodology/approach – The study analyzes attributes such as subject matter, size, date, and signature drawn from documented sources, including sales receipts, auction records, and online descriptions, and examines how art buyers can assess the risk of fakes based on these features. We conducted a comparative analysis of Félix Vallotton's paintings. The methodology involved compiling a detailed database, including characteristics from 1,704 original pieces and 210 fakes, to establish a robust classification and regression tree (CART) model.

Findings – Our results reveal that specific combinations of characteristics effectively distinguish original from fake artworks. The findings demonstrate the effectiveness of our approach for Vallotton's works and suggest its potential for standardization and application to other artists, providing a reliable and independent preventive measure against art counterfeiting.

Originality/value – Our study is original for at least two reasons. First, we were able to analyze a database comprising not only authentic works by Félix Vallotton but also more than 200 fakes, which gave us an excellent point of comparison. Second, we based our analysis on characteristics of artworks that are accessible to everyone, not just specialists, making our work directly accessible to the general public.

Keywords Art market, Fakes, Prevention, Art authentication methods, Classification, Félix Vallotton

Paper type Research article

Introduction

The infiltration of fake artworks into the art world through dishonest practices such as fraud poses a significant challenge to the preservation of cultural heritage. Legislators in various countries emphasize the importance of a proactive approach by art market participants as a crucial requirement to mitigate their vulnerability to art counterfeiting [1]. For instance, Article 146 of the Swiss Criminal Code and Article 248.1 of the Spanish Criminal Code, which delineate the components of fraud, emphasize the importance of victims exercising due diligence. According to these provisions, fraud occurs when an individual intentionally seeks to illicitly enrich themselves or others through deceptive tactics, thereby prompting the victim to act to the detriment of their financial interests or those of a third party. This article aims to assess the effectiveness of authentication methods in preventing fakes, while also introducing a novel approach using statistics to raise awareness about this issue among individuals in the art world.

Fraud in the art market does not concern only specialists and wealthy collectors, who can be considered direct victims, but also affects almost everyone, and is therefore a real problem for society (Charney, 2009). This problem particularly affects museums, many of which are funded by public money and whose collections are viewed by visitors from all walks of life. If these museums unknowingly possess and exhibit fakes, then all citizens are affected, not only because their tax money has financed the purchase of these fakes, but also because they are misled in



Funding: The authors received no financial support for this research.

Conflict of interest: The authors have no conflict of interest to disclose.

Arts and the Market
© Emerald Publishing Limited
e-ISSN: 2056-4953
p-ISSN: 2056-4945

DOI 10.1108/AAM-12-2025-0093

exhibitions by viewing such fakes, and they can be considered indirect victims (Bazley, 2010). To illustrate this point, we refer to a well-known case in which several museums in the United States unknowingly held and exhibited artworks produced by the prolific forger Mark Landis. During an internal audit, Matthew Leininger, Head of Conservation at the Oklahoma City Museum of Art, discovered that all the works donated by Landis were in fact fakes. Leininger therefore contacted the American Alliance of Museums to make sure that no other museum had received gifts from Landis. Regrettably, several museums (e.g. the Art Institute of Chicago and the Art Museum at the University of Kentucky) held paintings fabricated and gifted as personal donations by Landis that were fakes (Wilkinson, 2013).

The methodology presented here aims to demonstrate how art buyers can identify the risks of counterfeiting by examining various characteristics such as subject matter, size, date, and signature placement. These attributes, which remain consistent across both original and fake artworks, can be extracted from accompanying documents in the art market. While individually these features may not serve as obvious indicators of fakes, we hypothesize that certain clusters of them can effectively differentiate between original and fake artworks. To validate this hypothesis, we analyzed a set of real data describing both original and fake artworks using standard classification methods. This approach differs from studies that focus on identifying individual traits specific to original or fake artworks (e.g. Sablatnig *et al.*, 1998; Teegen, 2002; Lombardi, 2005; Montani, 2015; Elgammal *et al.*, 2018).

Our case study focuses on the work of Félix Vallotton (1865–1925), a painter who, despite taking many precautions, has nevertheless fallen victim to numerous attempts at counterfeiting. We chose this case because of the vast amount of data we were able to access, not only concerning authentic works, but also counterfeits. While the results presented in this article pertain specifically to Félix Vallotton's work, the methodology is standard and can be replicated for other artists using a new dataset containing characteristics of original and fake artworks of the examined artist.

Art authentication and fake detection

Art authentication relies on three essential pillars: stylistic connoisseurship, historical provenance research, and scientific expertise. These methods, employed individually or in combination, provide robust evidence supporting the authenticity of an artwork. Additionally, they can be employed to detect art counterfeits (Bazley, 2010). Are authentication and detection mutually exclusive or complementary functions? To address this question, we will explore why experts opt for one method over another. Subsequently, we will ascertain which method is suitable for prevention purposes.

Stylistic connoisseurship entails a meticulous visual inspection of an artwork conducted by an expert eye. It is a method of evaluation rooted in an informed perception and interpretation of the unique form and craftsmanship characteristic of the artist. A trained eye requires practice, sensitivity to quality, and the ability to visually reconstruct the artist's work. Visual memory provides a means to internalize comparisons guided by perception and the ability to categorize objects (Ebitz, 1988). Therefore, these skills are strongly conditioned by an expert's intellectual disposition and individual experience. To carry out a stylistic analysis, an expert in art authentication must perform two interdependent tasks. One is to identify features that correspond to the artist's characteristic style or unconscious personal traits. The other is to detect various errors committed by counterfeiters, such as negligence, omissions, and inconsistencies. When the examiner concludes that certain specific features of the artwork are identical to those of the original work and that there are no obvious errors, such a positive assessment may provide evidence of the work's authenticity. Otherwise, it turns into a negative evaluation, implying an art counterfeit. An intriguing observation is that, when compared to other authentication methods, stylistic connoisseurship stands out as the least objectively rooted in scientific objectivity. Scientific objectivity, in this context, pertains to the absence of personal bias. It does not exclude that scientists may have a certain degree of subjectivity, for

example, in interpreting results or selecting samples. However, an important premise is that the results should be objective and independent of the personal skills or experiences of the test-takers (Reiss and Sprenger, 2020). In the case of stylistic connoisseurship, the method of providing knowledge that relies on the examiner's trained eye depends on the intellectual skills, training, and experience of the person who conducted the analysis. Therefore, this method is scientifically subjective, but it holds significance due to its exclusivity, with the required expertise being rare and acquired by only a select few.

In contrast, historical provenance research analysis exhibits a higher degree of scientific objectivity compared to stylistic analysis. This method involves establishing the history of an artwork from its creation to the present day. Provenance information can be found in primary sources, such as receipts and bills of sale that document ownership transfers, and in secondary sources, such as catalogs and artist monographs (Yeide *et al.*, 2001). Like stylistic analysis, historical documentation analysis serves a dual purpose. It can authenticate a work of art by reconstructing its chain of ownership, while inconsistencies in the chronological sequence may cast doubt on its authenticity. Typically, historical analysis complements stylistic connoisseurship and is often conducted by the same experts.

When experts have doubts about the authenticity of an art object, they may turn to scientific expertise, using specific tools to detect anomalies. In essence, using, for example, computed tomography or spectroscopy, they seek characteristics that deviate from the typical style or materials of the artist or the period in question. For example, the incoherence between the composition of the pigments used in the production of a painting and the materials available during the artist's lifetime may confirm suspicions of counterfeit. For this purpose, there is a spectrum of scientific methods available in laboratories. However, when the technical analysis does not identify suspicious elements, the evidence is insufficient to confirm authenticity without the help of other methods. Reeves (2011) discussed the rights and risks of experts in courts, as later elaborated by Neuhaus (2014) in the context of art authentication. In this regard, Reeves (2011) pointed out that art experts usually view technical/scientific analysis as impractical because they are convinced that it can help detect fakes but cannot establish authenticity. Conversely, when suspicious evidence is uncovered, conclusions drawn using technical means become much less dependent on other methods.

Although technical methods are scientifically objective, they typically operate only to confirm doubts about authenticity rather than to refute them, with the exception of computer-assisted art classification approaches based on machine learning. These technical methods may operate similarly to stylistic analysis, as will be discussed further below.

Innovative technology for art fakes detection

The widespread use of information technology tools has made it possible to devise new authentication processes through the application of machine learning and other advanced computational techniques. This intersection of technology and art authentication is often referred to as computer-assisted art authentication (Bazley, 2010).

Scholars worldwide are developing effective methods for categorizing artists' styles. For instance, a study considered brushstrokes on portraits to enable semi-automated identification of artists (Sablattig *et al.*, 1998). Various machine learning methodologies have been explored to analyze and visualize style relationships within the realm of art (Lombardi, 2005). Another comparative analysis delved into the influence and interrelationships among artists, enriching our understanding of artistic evolution (Saleh *et al.*, 2014). Furthermore, the challenge of determining the authorship of paintings has encouraged innovative computational solutions. For instance, researchers have proposed a novel method harnessing multi-view local color features extracted from artworks to accurately ascertain the epoch in which a painting was created (Chen *et al.*, 2017).

The literature review also highlights studies dedicated to developing methodologies for detecting art counterfeiting (Craddock, 2009). Within these studies, researchers analyze specific characteristics of artworks, such as brushstrokes and color attributes. A study

conducted by researchers at Rutgers University, USA, used machine learning on a large dataset of fakes and originals (Elgammal *et al.*, 2018). Inspired by “pictology”, their approach employs statistical analysis by machines instead of the human eye. The dataset included 297 digitized works from books, downloaded images, and a set of 87 fake drawings resembling Picasso, Matisse, and Schiele’s styles. A new algorithm based on deep neural networks was developed to compare and classify handcrafted works. This innovative method reflects a shift towards utilizing advanced computer technology to detect art fakes, as compared to more standard approaches such as the one used to differentiate authentic Roman fibulas from copies (Teegen, 2002). Teegen proposed a method based on measurable parameters such as size, weight, and degree of corrosion, and principal component analysis.

Combining different approaches, the Art Fraud Insights consultancy project launched by art fraud investigator Colette Loll was aimed at predicting fraud in online marketplaces based on auction listing information. Art historians and fraud researchers collected data on suspected, faked, and forged artworks, focusing on vulnerable categories such as paintings, prints, and drawings priced between \$1,000 and \$5,000. Their analysis combined key indicators with advanced image recognition and documentation to develop a predictive model for fraud detection in online marketplaces (Loll, 2016).

In general, computer-assisted art authentication techniques are not typically viewed as standalone approaches to authentication. They are often proposed as supplementary methods to complement existing authentication techniques (e.g. Elgammal *et al.*, 2018). However, these techniques possess the potential to function as independent preventive measures. Their outcomes, expressed as probabilities of a work being fake, could serve as an alternative preventive tool accessible to non-specialists; however, to our knowledge, no study has yet achieved this goal. This is likely due to two critical limitations regarding applicability of such methods for non-professional users. Firstly, mathematical and computational challenges are involved in both using and understanding results from computer algorithms and softwares. Second, complex features such as brushstrokes or color attributes are difficult to capture and interpret without expert knowledge.

These challenges led us to consider alternative characteristics that, while less technical, are more accessible to the public. As a result, we turned to those commonly used in provenance research. This idea was inspired by the guidelines of the American Association of Museums for provenance research, which emphasize the importance of subtle details in artworks, revealing crucial information about their authenticity (Yeide *et al.*, 2001). These details include key aspects such as style, subject matter, signature, materials, dimensions, and even the frame. This information is typically readily available from sources such as sales receipts, auction catalogs, and online descriptions. We propose a novel approach based on statistical classification techniques to analyze the characteristics frequently assessed in provenance research. Our aim is to rationalize and enhance the detection process, particularly for non-professional users who may lack the specialized expertise required for traditional authentication methods. From this perspective, the approach seeks to democratize access to art authentication methods, thereby enabling a broader range of stakeholders to assess authenticity risks effectively within a rapidly evolving digital landscape.

Our research question was to determine the extent to which a statistical analysis of observable characteristics of artworks could help to discriminate between original and fake artworks. More specifically, we hypothesized that (H1) artworks can be grouped into distinct, homogeneous clusters based on publicly available features, and that (H2) some of these clusters are predominantly composed of original artworks, while others are primarily made up of fakes.

Methodological framework

Our case study focuses on Félix Vallotton (1865–1925), a Swiss-French painter whose artistic legacy, methodically preserved and protected by the Vallotton Foundation (<https://felixvallotton.ch/>), has nonetheless been impacted by the presence of fake artworks.

The meticulous notes and diaries of the artist have served as invaluable tools for authentication purposes and played a fundamental role in the creation of the catalogue raisonné in 2005 by the experts of the Vallotton Foundation (Ducrey and Poletti, 2005). Like other foundations, the Vallotton Foundation is a private institution under no obligation to detect or report fake artworks. In most cases, an art market participant requests the Foundation's expertise. In exceptional cases, for example, when a fake appears at a trusted art dealer (such as an auction house or renowned gallery), experts may take the initiative to give their opinions. However, as the exclusive institution responsible for issuing certificates of authenticity for Vallotton's artwork, the Foundation accurately preserves information on fakes. Its experts keep records comprising letters, photos, and correspondence documenting several hundred instances of fake paintings, drawings, and prints. Broadly, there are three categories of fakes: artworks bearing forged signatures or monograms (labeled as "fake"), artworks whose authenticity is doubted by experts ("doubtful"), and artworks without the artist's signature but previously attributed to Vallotton ("hypothetical"). This extensive collection served as a basis for our study. With the Foundation's consent, we were granted access to these materials, ensuring that all confidentiality concerns were resolved through formal agreements.

To conduct comparative analyses, we assembled a comprehensive database comprising details of 1,704 original artworks and 210 fake artworks attributed to Vallotton's oeuvre, integrating information from the catalogue raisonné and the archive files held by the Foundation. The analysis phase aimed to achieve three objectives: conducting a comprehensive examination, establishing a practical model for comparison across elements in our database, and minimizing information loss by solely relying on the expertise provided by authentication experts on Félix Vallotton's works. To ensure coherence with both sources (catalogue raisonné and files concerning fakes), we selected and codified information that corresponded to those mentioned in the catalogue raisonné.

In addition to their degree of authenticity ("original", "fake", "doubtful", and "hypothetical"), artworks were described through the presence or absence of sixty-five characteristics, each denoting a specific observed trait. These variables were subsequently categorized into distinct categories:

- (1) Title – one variable: Presence of artwork title in the documents.
- (2) Years – two variables: This group delves into the temporal aspect, featuring variables "with date" and "illegible date".
- (3) Market – five variables: This group offers information on the country in which the painting is located (Switzerland, France, or England), and sale details (whether the painting was put up for sale and whether it was sold).
- (4) Signature/monogram - five variables: Variables in this group detail the presence, type, and location of the artist's signature or monogram. The first variable, "without signature", indicates works of art without the painter's signature. The second and third variables, "monogrammed" and "signature or seal", indicate works with the artist's monogram or seal, respectively. The fourth variable, "trace of signature/monogram", refers to the existence of only a trace of a signature on the canvas or other support. The fifth variable, "signature on the inside of the pattern or over shoulders of the person", indicates instances where the signature is situated in an atypical location, often over the shoulders of the figure in a portrait.
- (5) Place of signature/monogram – eight variables: This group specifies the location of the artist's signature or monogram on the artwork: "lower right", "lower left", "upper right", "upper left", "bottom center", "center right", "center left", and "top center".
- (6) Technical characteristics – twenty-two variables: This group encompasses dimensions, subjects, and techniques: "large", "medium", "small", "portrait", "portrait with

identified person”, “still-life”, “landscapes”, “private and public interiors with or without figures”, “great decorations”, “mythological, allegorical or biblical subjects”, “various”, “nudes”, “copy of another famous painter”, “oil”, “pastel”, “tempera”, “gouache”, “unknown technique”, “canvas”, “cardboard”, “wood”, “paper”.

- (7) Back of the painting – eight variables: The variable in this group provides information on the back of the painting: markings, stickers, seals, signatures/initials, date, information written on the frame, information written on the canvas, information under other painting.
- (8) Historical documentation – five variables: This group offers insights into ownership chronology: “without provenance”, “recent provenance”, “first provenance”, “restored”, “relining”.
- (9) Period – nine variables: The final group categorizes paintings into distinct periods, capturing the evolution of Vallotton’s artistic style over time: “before 1882”, “1882–1889 youth”, “1890–1900 Nabi period”, “1901–1910”, “1901–1905 landscapes”, “1905–1910 nudes”, “1909–1915 sunsets”, “1915–1917 war”, “1916–1925 sunsets”.

To facilitate statistical analysis, all variables underwent dummy coding, where “1” represented the presence of a particular attribute, and “0” indicated its absence.

We used classification and regression tree (CART) analysis to build predictive models for determining the authenticity of artworks based on their various characteristics (Breiman *et al.*, 1984; Strobl, 2013). The principle of CART is to identify at each stage the predictor variable that allows the best possible differentiation between the two categories of a dependent variable, in our case, whether a work of art is an original or a fake. Starting with the entire dataset, known as the root node, comprising both original and fake artworks, the objective is to divide it into a set of smaller groups or nodes, each containing a maximum number of either original or fake artworks. To achieve this objective, the CART algorithm attempts to divide the content of each node according to each of the predictors available in the database. The division performed by each predictor is then evaluated by measuring the impurity of the result, i.e. the degree of dissimilarity of the artworks in a given node. Various measures can be used for impurity, including entropy, but CART uses the Gini index for this purpose (Gini, 1921; Daniya *et al.*, 2020). Let C be the number of categories of the target variable (two categories in our case, original and fake artworks), and let $p(c)$ be the proportion of artworks in category c , then the Gini index for a specific node is

$$Gini = 1 - \sum_{c=1}^2 p(c)^2$$

The Gini index takes a value of zero when there is no impurity and the node contains only original or fake artworks. It then increases as impurity increases and reaches its maximum value when there is exactly 50% of each type of artwork in a node. For example, if a node contains 100% original artworks and therefore 0% fake artworks, then

$$Gini = 1 - (1^2 + 0^2) = 1 - 1 = 0$$

And the impurity of the node reaches its minimum value. Conversely, if a node contains 25% original and 75% fake artworks, then

$$Gini = 1 - (0.25^2 + 0.75^2) = 1 - 0.625 = 0.375$$

In practice, the Gini index is calculated for each node, then a weighted average taking into account the total number of artworks in each node is calculated to obtain the overall index for

the entire model. The division obtained with the predictor that minimizes the overall impurity of the entire model is retained, and the algorithm iterates, testing all predictors again on all nodes.

Once all possible divisions of the data have been identified, the result is presented in the form of a classification tree showing the role of each predictor in the final model. The terminal nodes contain subsets of data that can no longer be divided, either because they contain only one type of artworks (original or fake) or because there are no more predictors available that would improve the explanation of that subset. Figure 1 is an example of such a tree: The root node contains the complete sample of 1,704 original and 127 fake artworks. It is first divided according to the “provenance” predictor. Next, the subset of works “without provenance” is in turn divided according to the “date” predictor, and so on, until six terminal nodes appear at the bottom of the figure. Note that the same predictive variable can be used several times by a model. Figure 2 illustrates such a situation with the “title” predictor.

We applied the CART approach to build two distinct models. The first model, “Original vs Fake1” utilized artwork characteristics to differentiate between original artworks and known fake artworks, those identified as “fake” in the files of the Vallotton Foundation. This model enables us to concentrate on the most straightforward cases of authenticity versus identified fakes.

The second model, “Original vs Fake2”, aimed to distinguish original artworks from all categories of fake artworks in our database, encompassing “fake”, “doubtful”, and “hypothetical” artworks. This broader model offers a more extensive classification, capturing a wider spectrum of fake types and providing a deeper insight into the intricacies of the data.




Figure 1. Classification tree for the Original vs Fake1 model




Figure 2. Classification tree for the Original vs Fake2 model

The overall quality of each classification was evaluated through the computation of the sensitivity (percentage of correctly identified fakes) and the specificity (percentage of correctly identified original artworks) of the results. If these values give a first idea of the predictive quality of a model, it is essential to assess how the results of a statistical analysis are independent from the specific characteristics of the sample used (Picard and Cook, 1984). For this reason, we also performed a cross-validation study using 10,000 replications of the dataset (Blockeel and Struyf, 2002; Borra and Di Ciaccio, 2010). For each replication, the original dataset was randomly split into two equal-size subsamples, a training set and a validation set. Then the CART model was estimated on the training set and applied on the validation set. We provide sensitivity and specificity for both sets.

Finally, we tested our approach on two additional paintings that had not been used in the construction of our models, one assumed to be original artwork by Félix Vallotton and another for which experts had expressed serious doubts.

The open-source statistical environment R (R Core Team, 2021) was used for all computations, with the specific *rpart* package for the estimation of the CART models.

Results

Figure 1 shows the classification tree for the “Original vs Fake1” model. We analyzed a sample of 1,831 artworks, including 1,704 originals and 127 fakes. All terminal and decision nodes are labeled either “Original” if they contain a majority of original artworks, or “Fake1” if they

contain a majority of fakes. The two figures appearing in each box are the respective numbers of original and fake artworks in each node. The two lines starting from each decision node are labeled with the two categories of the predictor used to divide the data. For example, for the root node appearing on top of the figure, “with provenance” is written on the left-hand line going down to a terminal node, and “without provenance” is written on the right-hand line going to a second intermediary node containing 59 authentic artworks and 113 fakes. In this example, the predictor “provenance” was selected by the CART algorithm because it was the predictor that allowed for the greatest reduction in impurity according to the Gini index. This also means that the split based on the “provenance” predictor reduces impurity, resulting in more homogeneous groups than the root node.

Figure 2 presents the classification tree for the “Original vs Fake2” model. In this case, we analyzed a sample of 1,914 works (1,704 original and 210 fake artworks including the 127 fake artworks in model 1 as well as 70 hypothetical artworks, and 13 doubtful artworks).

Before interpreting the results, it is essential to highlight that a classification tree operates on a probabilistic basis, aiming to identify the most probable scenario rather than an exact one (Borra and Di Ciaccio, 2010). As a result, not all works classified in final nodes labeled as “Original” are truly original, and not all works classified in final nodes labeled as “Fake” are actually fake. This means that the predictors available to the CART model did not allow us to reliably distinguish all original works from fakes. For example, in the first node labeled “Fake1” in Figure 1, there are indeed 16 fakes, but also 4 authentic artworks. In contrast, the second node labeled “Fake1” contains only fakes. Adding additional variables could improve this, but without any absolute certainty. An ideal model should have a 100% correct classification rate. In practice, this statement means that the closer the model’s proportion of correct classifications is to 100%, the less impact errors have on the results. To assess this impact, we computed the sensitivity (percentage of correct classification for fake artworks) and the specificity (percentage of correct classification for original artworks) for each analysis.

For the Fake1 model, 82.7% of fakes and 99.6% of original artworks were correctly classified. In comparison, the Fake2 model achieved a slightly higher sensitivity of 89.0%, while retaining the same specificity of 99.6%. This improvement in sensitivity in the Fake2 model can be attributed to its broader scope, which includes distinguishing original artworks from all categories of fake artworks in our database. By encompassing a wider range of fake types, the Fake2 approach leverages more comprehensive data, allowing for finer distinctions and ultimately leading to better overall accuracy in identifying fakes.

Interpretation of classification trees

A classification tree algorithm works by progressively dividing data into more specific groups based on key variables. It begins with an initial split and continues to create further divisions based on additional variables (Moisen, 2008). Each division forms a rule that aids in categorizing the data into different classes. The terminal nodes of the tree represent the ultimate classifications. This process generates clear, interpretable rules and helps in understanding the likelihood of each classification being accurate, making it a valuable tool for predictive modeling and decision-making.

The results of such rule interpretation are presented in Tables 1 and 2. To interpret each rule, we need to understand how each predictive variable contributes to the classification outcome. To do this, we analyze the significance of each variable in relation to the classification decision. This involves examining the roles of various predictors, for example, “provenance”, “date”, “title”, and “presence of a signature or seal”, in determining the authenticity of the artwork. Such an interpretation in plain English is important to enable non-specialists to understand and use our results. By understanding the interplay between these factors, we can grasp why the classification tree assigns a high probability to be either original or fake to artworks meeting these specific criteria.

Table 1. Rules and interpretations for the Original vs Fake1 model

Terminal node	Rule	Interpretation
1	With a provenance	Works of art with complete provenance, as stated in accompanying documents on the art market, have a 99% probability of being original
2	Without a provenance With a date With a titled With a signature/ seal	When the documents accompanying an artwork do not contain information on provenance but include details such as the date noted on the support, the artwork’s title, and a signature or seal, the probability of the artwork being original remains high at 89%
3	With a provenance With a date With a titled Without a signature/seal In France	When the documents accompanying an artwork do not contain information on provenance but include details such as the date noted on the support, the artwork’s title, and a signature or seal, yet the owner of the artwork is located in France, the probability of the artwork being original decreases slightly to 75%
4	Without a provenance With a date With a titled Without a signature/seal Not in France	When the documents accompanying an artwork do not contain any information about provenance or a signature but include information about the date of creation and a title, the artwork has an 80% probability of being fake if the owner is in any country other than France
5	Without a provenance With a date Without a titled	Artwork without a provenance, without a title, and with a date have a 100% probability of being a fake
6	Without a provenance Without a date	Artwork without a provenance and without a date noted in accompanying documents have a 97% probability of being a fake

Validation and performance analysis of the classification models

As mentioned earlier, the CART approach resulted in a sensitivity and specificity of 82.7 and 99.6% respectively for the first model, and 89.0 and 99.6% for the second model, and these results are excellent, as they are close to 100%. However, we cannot rule out the possibility that these good performances are due to specific characteristics of our data rather than to the overall quality of the approach. This is why we applied a cross-validation procedure.

The cross-validation results presented in [Tables 3 and 4](#) revealed that our model still performs well when tested on other samples. Notably, the specificity of 99.3% computed on the validation set closely mirrors that of the total sample (99.6%) for both analyses. This indicates that the model reliably identifies original artworks across various datasets, thereby reducing the risk of misclassifying original artworks as fake.

In terms of sensitivity, the results are 5–7% lower in the validation sets than in the original sample. These slight variations in sensitivity were expected, especially considering that fake artworks are nearly ten times fewer in number than original artworks in our dataset. This imbalance can make it more challenging for the model to learn and accurately classify the minority class (fakes), as it is trained on a dataset dominated by originals. Consequently, this can lead to a slight reduction in the model’s sensitivity, as the model may become better at identifying originals than fakes.

Overall, the cross-validation results confirm that our model is both accurate and generalizable. These findings support the model’s predictive reliability.

Table 2. Rules and interpretations for the Original vs Fake2 model

Terminal node	Rule	Interpretation
1	With a provenance With a title	Works of art with complete provenance and a title, as stated in accompanying documents on the art market, have a 99% probability of being original
2	With a provenance Without a title	Artwork with provenance, as stated in accompanying documents on the art market, but without a title, have a 100% probability of being fake or presents high doubts about its authenticity. Such works could be hypothetical pieces attributed to the artist, indicating a high risk regarding their authenticity
3	Without a provenance With a date With a title	When the documents accompanying an artwork do not contain information on provenance but include details such as the date noted on the support, the artwork's title, and a signature or seal, the probability of the artwork being original decreases to 85%
4	With a signature/seal Without a provenance With a date With a title Without a signature/ seal In France	Artwork without a provenance and signature, with a date and title, whose owner is in France, have their probability of being original decreased slightly to 75%
5	Without a provenance With a date With a title Without a signature/ seal Not in France	Artworks without a provenance and signature, with a date and title, whose owner is not in France have an 81% probability of being a fake
6	Without a provenance With a date Without a title	Artworks without a provenance, without a title, and with a date have a 100% probability of being a fake
7	Without a provenance Without a date	Artworks without a provenance and without a date have a 99% probability of being a fake

Table 3. Cross-validation results for the “Original vs Fake1” model. We provide the sensitivity and specificity for both the training and the validation sets

	Minimum	Q1	Median	Mean	Q3	Maximum
Training set: Sensitivity	62.5	79.3	82.4	85.2	85.5	96.8
Training set: Specificity	98.3	99.3	99.4	99.5	99.6	100.0
Validation set: Sensitivity	46.7	73.3	77.9	77.5	82.5	96.2
Validation set; Specificity	96.0	99.2	99.3	99.3	99.5	100.0

Table 4. Cross-validation results for the “Original vs Fake2” model. We provide the sensitivity and specificity for both the training and the validation sets

	Minimum	Q1	Median	Mean	Q3	Maximum
Training set: Sensitivity	73.5	84.2	86.8	86.7	89.4	98.0
Training set: Specificity	97.9	99.2	99.4	99.4	99.6	100.0
Validation set: Sensitivity	60.8	79.4	82.9	82.8	86.3	96.9
Validation set; Specificity	95.4	99.1	99.3	99.3	99.5	100.0

Evaluation of additional works

To assess the model's predictive validity, we tested it on two artworks that were not included in the original training set but were subsequently authenticated (or rejected) by Vallotton experts. The model correctly predicted both outcomes, assigning a 99% probability of being a fake to the fake piece and a 99.8% probability of authenticity to the original artwork. These results closely aligned with expert evaluations based on provenance and stylistic analysis, reinforcing the model's reliability and practical relevance.

Discussion

This study aimed to rationalize and enhance the detection process of fake artworks, particularly for non-professional users who may lack the specialized expertise typically required in traditional authentication settings. With this goal in mind, we set out to develop an approach that could help democratize access to authentication methods, thereby enabling a broader range of stakeholders, including collectors, dealers, institutions, and researchers, to assess authenticity risks with greater confidence and objectivity. Specifically, the study examined how statistical analysis of observed artwork characteristics can distinguish original artworks from fakes. To this end, we tested two core hypotheses:

- H1.* Artworks can be grouped into distinct, homogeneous clusters based on publicly available features.
- H2.* Some of these clusters are predominantly composed of original artworks, while others are primarily made up of fakes.

Our findings support both hypotheses. Using the CART algorithm, we identified clear and consistent clusters based on typical features found in art market documentation, such as provenance, title, date, and signature, thus validating [H1](#). Moreover, the model revealed that certain clusters were largely composed of either original or fake artworks, offering strong support for [H2](#).

A key strength of our approach lies in its transparency and ease of use. Although CART is not the most advanced decision tree algorithm, its ability to handle diverse variable types and model non-linear relationships makes it particularly well suited to our dataset. Furthermore, given that the art market is often difficult to understand and can sometimes resemble a “black box” system, it is advantageous to use a well-defined statistical algorithm. Additionally, our approach relies solely on routinely available documentation, making it a practical and widely usable tool for collectors, dealers, institutions, and researchers. In our practical illustration, we have used all 65 available characteristics from the works as potential variables for the CART model, but it is perfectly possible to use only a much smaller set of variables. This could make it possible to test very specific hypotheses concerning certain types of counterfeiting. For example, we assume that a work considered stolen may lead certain actors to produce a counterfeit that reappears in its place on the market, but this counterfeit should then have characteristics very close to those of the stolen painting to have a chance of being credible. Conversely, in the notes taken by Vallotton during his lifetime, there is sometimes mention of a series of paintings, although neither the number nor the exact subjects are known. In this case, counterfeits that were quite different from one another, in terms of subjects and sizes for example, could have been produced.

Despite its promise, our study has certain limitations. Most importantly, the model was trained exclusively on works related to Félix Vallotton's artistic practice. While the resulting decision rules proved robust within this context, the relative weight and interaction of such features may differ substantially across artists. A work of art is inherently the product of individual creativity, shaped by the artist's personality, historical context, and cultural environment. Consequently, correlations among the characteristics of original artworks are artist-specific and cannot be mechanically extrapolated to other artists. Because fake artworks

are created in reference to the stylistic, documentary, and market conditions surrounding a particular artist, they are likewise shaped by artist-specific patterns. The indicators that distinguish original artworks from fakes in one artistic context may therefore not retain the same predictive value in another. As a result, both the characteristics of original artworks and the strategies adopted by counterfeiters are contingent upon the specific artistic environment, limiting direct transferability of the model across different artists.

Accordingly, the model outcomes summarized in [Tables 1 and 2](#) are exclusively applicable to Vallotton's works. That said, the model incorporates several predictive variables that are, in principle, common across artists, such as "presence of a signature" or "date". However, the strength and statistical significance of these variables, as well as their interactions with other features, may vary from one artist to another.

This variability is further illustrated by prior empirical research on price formation and financial returns in the art market ([Renneboog and Spaenjers, 2013](#); [Cinefra et al., 2019](#)). [Cinefra et al. \(2019\)](#) found that the variables "signed" and "dated" were not statistically significant in their analysis of Joan Miró's works. They explain this result by noting that the vast majority of Miró's paintings are both signed and dated, leaving insufficient variation to detect a measurable effect. By contrast, [Renneboog and Spaenjers \(2013\)](#), analyzing more than one million auction transactions involving paintings and works on paper by multiple artists, demonstrated that these same features can be highly significant in broader and more heterogeneous datasets. Their hedonic regression results indicate that the presence of a signature is associated with an average price premium of up to 31%, while the inclusion of a date increases value by nearly 19%. The divergence between these findings underlines that the impact of seemingly universal features such as signature and date is highly context-dependent and artist-specific. This contrast reinforces the need for caution when attempting to generalize predictive models across different artistic settings. In addition, [Cinefra et al. \(2019\)](#) note that the presence of a signature and a date generally contributes to an artwork's perceived authenticity. They further suggest that the variables authenticated, provenance, and catalogue raisonné can be regarded as indicators of a work's authenticity. Consistent with these observations, our findings likewise identify "provenance" as a central factor in distinguishing original artworks from fakes.

Provenance presents particular challenges for statistical measurement. Ideally, it should consist of a comprehensive and continuous chronology of ownership from the moment of creation to the present, including documented transfers, modes of sale, and successive locations of the work. [Cinefra et al. \(2019\)](#), for instance, defined provenance as "the number of owners the artwork had had until the time of the auction". In our study, we adopted a more detailed approach by collecting all available provenance-related information. To this end, we structured provenance-related characteristics into three complementary variables. "Without provenance" was assigned to paintings lacking any ownership chronology, while "with provenance" indicated the presence of such documentation. The variable "recent provenance" refers to paintings for which the recorded ownership history does not begin close to the time of the work's creation but only several decades later; in other words, the early stages of the work's circulation are undocumented. Paintings acquired directly from the artist but lacking a complete subsequent ownership chain were coded as "first provenance".

The results further indicate that, among the variables within this group, only the presence or absence of a complete and well-documented provenance was retained by the algorithm as an independent predictor with sufficient classification power to reliably distinguish authentic works from fakes; the other variables were not statistically significant. Artworks supported by a complete provenance showed a 99% probability of being original. Its predictive strength exceeded that of other predictive variables, as it was the only indicator able to classify autonomously rather than as part of a cluster of associated characteristics (see [Table 1](#)). As expected, the lack of provenance, especially when combined with other missing key indicators, was associated with a substantial decline in authenticity probability, dropping to between 15 and 0% (see [Tables 1 and 2](#)).

Although stylistic, material, and documentary features may vary substantially across artistic practices, the existence of a coherent ownership history remains a structurally important and widely recognized component of due-diligence procedures and authenticity assessments in the art market. For this reason, provenance may be considered one of the few indicators whose predictive relevance could, in principle, allow for cautious generalization beyond the specific case examined. However, such generalization must remain qualified, as certain artists, historical periods, or market conditions may systematically generate incomplete or fragmented ownership records, thereby limiting its statistical significance and reducing its predictive power in specific contexts.

From a practical perspective, our model could be implemented by organizations devolved with authenticating authority and possessing archives of fakes, such as artists' foundations, research institutes, and museums. Since these institutions generally have exclusive access to the necessary data and are best positioned to guarantee users' trust, it is appropriate that they serve as the primary custodians of these detection methods. Accordingly, these institutions could directly benefit from the model as a complementary technical tool. In cases of uncertainty, statistical outputs may support stylistic judgments and reduce the need for extensive justification in relatively straightforward assessments. Likewise, art market intermediaries, such as auction houses, galleries, and legal advisers, could encourage clients to conduct preliminary evaluations or carry out these assessments internally before seeking formal expert opinions. Overall, this structured pre-assessment approach could enhance efficiency, improve decision-making, and strengthen due-diligence practices for both experts and market participants.

Applying this method to other artists would require the creation of new databases, which would involve initial costs. These expenses, however, could be progressively amortized through moderate user fees that remain lower than the cost of traditional expertise, while the development of supporting software would remain comparatively inexpensive. Appropriate governance frameworks would nevertheless be essential to ensure data accuracy, ideally under the supervision of artists' foundations or similarly specialized institutions.

Beyond private initiatives, public organizations may also play a role in implementing such systems. Interpol's ID-Art application for tracing stolen works illustrates the potential of comparable tools. However, the coordinator of Interpol's Works of Art Unit noted that developing a database dedicated to counterfeit artworks would require judicially confirmed cases, which Interpol itself is unable to compile despite recognizing the necessity of such a resource ([The Art Law Foundation, 2021](#)).

In this context, the proposed model represents a pragmatic compromise: although its database includes non-judicially confirmed counterfeits, its outputs are statistical risk indicators. Grounded in expert knowledge and validated quantitatively, these results could function as a strong preventive tool, encouraging users to seek professional advice rather than act fast. While such databases would initially operate on an artist-by-artist basis, state authorities could stimulate their development through foundations or museums and deploy artist-specific applications, similar to those used for stolen art.

Future applications of our approach will depend heavily on the quality and scope of available data. Ideally, data on original artworks should be sourced from catalogues raisonnés, monographs, and archival documents. However, the reliability of such sources is contingent on the reputation of the experts behind them, as no formal standards currently govern their production. Equally essential is access to verified data on fakes identified by recognized experts; however, individuals and institutions are often reluctant to acknowledge possessing such data, particularly when they themselves have been deceived by counterfeiters. Moreover, to ensure statistical reliability, datasets must be sufficiently large, as smaller sample sizes introduce greater variance and reduce predictive consistency. Ultimately, beyond the quality of the data, the classification algorithm also plays a key role in determining the relevance of the results. The research presented in this article is a first step toward evaluating the value of statistical classification for the authentication of works of art.

This is why we chose to apply the CART algorithm, which is known for being fast and for producing results that are easy to visualize and interpret. However, for future research, we could use a Bayesian classification algorithm, such as BART, for example. The implementation cost will be higher, but on the other hand, it will be possible to obtain more robust results that are likely to adapt more easily to a wide variety of artists and situations.

Looking ahead, standardizing the methodology would require the development of a generic software capable of processing any structured dataset. Such a tool could facilitate semi-automated assessments across a wide range of artists by integrating all observable characteristics into a unified database and applying consistent analytical procedures. This would make the model more adaptable and scalable, increasing its utility for broader applications.

Conclusion

This study demonstrates that a statistical approach based on routinely available descriptive and historical data can effectively distinguish between original and fake artworks. By confirming both hypotheses, we demonstrate that artworks can be meaningfully grouped into distinct clusters, some clearly associated with original artworks and others with fakes.

Our contribution lies in two main areas: high classification accuracy and practical usability. The model generates clear, data-driven decision rules based on standard documentation, making it especially relevant in the fluid and often opaque art market. It offers a valuable preliminary tool for assessing authenticity, one that complements, rather than replaces, expert judgment. While the elements used here are familiar to provenance researchers, they are rarely deployed within a systematic, data-driven framework. This study formalizes that process, demonstrating how structured analysis can extract new insights from conventional documentation.

Although the current model is specific to Félix Vallotton, the underlying methodology is flexible. With access to high-quality, well-structured data, it could be adapted to other artists and contexts. Future development may include the creation of generic software capable of analyzing standardized datasets for broader application. Access could therefore be partially restricted, so that users would obtain only probabilistic outputs without visibility into the underlying predictive characteristics or datasets. In other words, the technology can be made fully accessible while keeping sensitive data on fake artworks confidential.

The findings reveal statistical associations among documentary and contextual characteristics. These results highlight structural features commonly associated with fake artworks without offering a practical framework for improving them. The model thus helps reduce informational asymmetries for non-specialist market participants and strengthens risk awareness prior to expert consultation. In a field historically marked by opacity, responsible dissemination supports methodological research and contributes to broader fraud prevention efforts.

In a market where the risk of fakes remains ever-present, our findings highlight the potential of structured, data-informed methods to enhance transparency and build trust. By integrating statistical rigor into authenticity assessments, such approaches offer a valuable safeguard for professionals and institutions operating in today's global art market.

Acknowledgments

The authors would like to express their sincere gratitude to the Félix Vallotton Foundation for its invaluable assistance and for facilitating access to the case materials.

1. In this article, the terms *fake* and *counterfeit* are used broadly, even though they have subtle differences in meaning. *Fake* is a more general and widely understood term, referring to objects that are not genuine, regardless of whether they were intended to deceive. *Counterfeit* is often used in a legal context to describe items deliberately made to deceive by imitating an original. For clarity and readability, we use *fake* as the primary term, while *counterfeit* is occasionally used to add nuance.

References

- Art Law Foundation (2021), "Exclusive members' event: due diligence in the art market", *Interpol's new app to trace and identify stolen cultural property*, available at: <https://artlawfoundation.com/fda-events/exclusive-members-event-due-diligence-in-the-art-market-interpols-new-app-to-trace-and-identify-stolen-cultural-property/> (accessed 29 December 2022).
- Bazley, T. (2010), *Crimes of the Art World*, Praeger, Santa Barbara, CA.
- Blockeel, H. and Struyf, J. (2002), "Efficient algorithms for decision tree crossvalidation", *Journal of Machine Learning Research*, Vol. 3 No. 4, pp. 621-650, doi: [10.48550/arXiv.cs/0110036](https://doi.org/10.48550/arXiv.cs/0110036).
- Borra, S. and Di Ciaccio, A. (2010), "Measuring the prediction error: a comparison of cross-validation, bootstrap and covariance penalty methods", *Computational Statistics and Data Analysis*, Vol. 54 No. 12, pp. 2976-2989, doi: [10.1016/j.csda.2010.03.004](https://doi.org/10.1016/j.csda.2010.03.004).
- Breiman, L., Friedman, J.H., Olshen, R.A. and Stone, C.J. (1984), *Classification and Regression Trees*, Chapman and Hall/CRC Press, Boca Raton, FL.
- Charney, N. (2009), *Art and Crime: Exploring the Dark Side of the Art World*, Praeger, Santa Barbara, CA.
- Chen, L., Chen, J., Zou, Q., Huang, K. and Li, Q. (2017), "Multi-view feature combination for ancient paintings chronological classification", *Journal on Computing and Cultural Heritage*, Vol. 10 No. 2, pp. 7-15, doi: [10.1145/3003435](https://doi.org/10.1145/3003435).
- Cinefra, J., Garay, U., Mibelli, C. and Pérez, E. (2019), "The determinants of art prices: analysis of Joan Miró", *Academia. Revista Latinoamericana de Administración*, Vol. 32 No. 3, pp. 373-391, doi: [10.1108/ARLA-06-2018-0121](https://doi.org/10.1108/ARLA-06-2018-0121).
- Craddock, P.T. (2009), *Scientific Investigation of Copies, Fakes and Forgeries*, Elsevier/Butterworth-Heinemann, Oxford.
- Daniya, T., Geetha, M. and Suresh Kumar, K. (2020), "Classification and regression trees with Gini index", *Advances in Mathematics: Scientific Journal*, Vol. 9 No. 10, pp. 8237-8247, doi: [10.37418/amsj.9.10.53](https://doi.org/10.37418/amsj.9.10.53).
- Ducrey, M. and Poletti, K. (2005), *Félix Vallotton (1865-1925): L'oeuvre Peinte*, Continents, Lausanne, 5.
- Ebitz, D. (1988), "Connoisseurship as practice", *Artibus et Historiae*, Vol. 9 No. 18, pp. 207-212, doi: [10.2307/1483344](https://doi.org/10.2307/1483344).
- Elgammal, A., Kang, Y. and Leeuw, M.D. (2018), "Picasso, Matisse, or a fake? Automated analysis of drawings at the stroke level for attribution and authentication", *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, pp. 42-50, New Orleans, LA, USA, 2-7 February 2018, doi: [10.1609/aaai.v32i1.11313](https://doi.org/10.1609/aaai.v32i1.11313).
- Gini, C. (1921), "Measurement of inequality of incomes", *The Economic Journal*, Vol. 31 No. 121, pp. 124-126, doi: [10.2307/2223319](https://doi.org/10.2307/2223319).
- Loll, C. (2016), "Art, anxiety and the online marketplace", in Tompkins, A. (Ed.), *Art Crime and its Prevention: A Handbook for Collectors and Art Professionals*, Lund Humphries, London, pp. 60-72.
- Lombardi, T.E. (2005), "The classification of style in fine-art painting", Pace University, NY, PhD Thesis.
- Moisen, G.G. (2008), "Classification and regression trees", in Jorgensen, S.E. and Fath, B.D. (Ed. s), *Encyclopedia of Ecology*, Elsevier, Oxford, pp. 582-588.
- Montani, I. (2015), "Exploring transparent approaches to the authentication of signatures on artwork", University of Lausanne, Switzerland, PhD Thesis.

- Neuhaus, N.M. (2014), “Art authentication: protection of art experts from a Swiss perspective”, *Art, Antiquity and Law*, Vol. 19 No. 1, pp. 59-74.
- Picard, R.R. and Cook, R.D. (1984), “Cross-validation of regression models”, *Journal of the American Statistical Association*, Vol. 79 No. 387, pp. 575-583, doi: [10.2307/2288403](https://doi.org/10.2307/2288403).
- R Core Team (2021), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, available at: <https://www.R-project.org/> (accessed 8 September 2023).
- Reeves, V.K. (2011), “The rights and risks of experts in French and American courts”, *IFAR Journal*, Vol. 12 No. 4, p. 18.
- Reiss, J. and Sprenger, J. (2020), “Scientific objectivity”, in Zalta, E.N. (Ed.), *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, available at: <https://plato.stanford.edu/archives/win2020/entries/scientific-objectivity/> (accessed 2 July 2024).
- Renneboog, L. and Spaenjers, C. (2013), “Buying beauty: on prices and returns in the art market”, *Management Science*, Vol. 59 No. 1, pp. 36-53, doi: [10.2139/ssrn.1352363](https://doi.org/10.2139/ssrn.1352363).
- Sablatnig, R., Kammerer, P. and Zolda, E. (1998), “Hierarchical classification of paintings using face- and brush stroke models”, *Proceedings of the 14th International Conference on Pattern Recognition*, Vol. 1, pp. 20-24, Brisbane, Queensland, Australia, 20-20 August 1998, doi: [10.1109/ICPR.1998.711107](https://doi.org/10.1109/ICPR.1998.711107).
- Saleh, B., Abe, K., Arora, R.S. and Elgammal, A. (2014), “Toward automated discovery of artistic influence”, *Multimedia Tools and Applications*, Vol. 75 No. 7, pp. 3565-3591, doi: [10.48550/arXiv.1408.3218](https://doi.org/10.48550/arXiv.1408.3218).
- Strobl, C. (2013), “Data mining”, in Little, T.D. (Ed.), *The Oxford Handbook of Quantitative Methods in Psychology*, Oxford University Press, NY, pp. 678-700.
- Teegen, W.R. (2002), “Late 19th–early 20th c. replicas, imitations, and falsifications (?) of brooches from the Pyrmont spring: a light and scanning electron microscopic approach”, *Journal of Roman Archaeology*, Vol. 39 No. 2, pp. 287-289, Supplementary Series.
- Wilkinson, A. (2013), “The giveaway”, *The New Yorker*, available at: <https://www.newyorker.com/magazine/2013/08/26/the-giveaway> (accessed 28 February 2019).
- Yeide, N.H., Walsh, A. and Akinsha, K. (2001), *The AAM Guide to Provenance Research*, American Association of Museums, Washington DC.

Further reading

- Chowdhury, S., Lin, Y., Liaw, B. and Kerby, L. (2022), “Evaluation of tree-based regression over multiple linear regression for non-normally distributed data in battery performance”, *Proceedings of the 2022 International Conference on Intelligent Data Science Technologies and Applications (IDSTA)*, pp. 17-25, San Antonio, TX, USA, 05-07 September 2022, doi: [10.1109/IDSTA55301.2022.9923169](https://doi.org/10.1109/IDSTA55301.2022.9923169).
- Hayes, T., Usami, S., Jacobucci, R. and McArdle, J.J. (2015), “Using classification and regression trees (CART) and random forests to analyze attrition: results from two simulations”, *Psychology and Aging*, Vol. 30 No. 4, pp. 911-929, doi: [10.1037/pag0000046](https://doi.org/10.1037/pag0000046).
- Johnson, R.A. and Wichern, D.W. (2002), *Applied Multivariate Statistical Analysis*, Pearson Education, Great Britain.
- Miglio, R. and Soffritti, G. (2004), “The comparison between classification trees through proximity measures”, *Computational Statistics and Data Analysis*, Vol. 45 No. 3, pp. 577-593, doi: [10.1016/S0167-9473\(03\)00063-X](https://doi.org/10.1016/S0167-9473(03)00063-X).
- Sharp, J. (2020), “Subjectivity”, in Kobayashi, A. (Ed.), *International Encyclopedia of Human Geography*, Elsevier, Oxford, pp. 117-120.

Corresponding author

André Berchtold can be contacted at: andre.berchtold@unil.ch

For instructions on how to order reprints of this article, please visit our website:

www.emeraldgroupublishing.com/licensing/reprints.htm

Or contact us for further details: permissions@emeraldinsight.com