21

CURATION, CONTENT, CREATION

Computer Approaches to the Fine Arts

Javier de la Rosa and Juan-Luis Suárez

Portrait of Edmond de Belamy, an "algorithmically" generated painting featured at New York auction house Christie's, sold for $432,500 in October of 2018. An anonymous bidder acquired the piece over the phone for a price nearly forty-five times the highest estimate. In the process, Christie's became the "first auction house to offer a work of art created by an algorithm," for whatever that is worth.[1] The gold-framed but otherwise dull depiction of a fictional man mimicking the style of nineteenth-century European portraiture set the path for more contenders to enter the increasingly bigger artificial intelligence–fueled art market. A few months later, New York's HG Contemporary gallery in Chelsea presented *Faceless Portraits Transcending Time*, an exhibition where all paintings were created by a computer; in particular the paintings were produced by the "collaboration" between an algorithm named AICAN and its creators,[2] although arguably the machine did most of the work.[3] And while some have long foreseen the apocalypse of art in these new approaches,[4] it remains to be seen for how long these attempts will survive in a market that favors freshness and originality over innovation, as John McCormack put it when referring to the "precious bubble."[5] Certainly, it has become almost impossible to escape the hype of technology in every field imaginable, and the frontiers of what separates human from artificial creation is ever blurrier. However, every technological process, no matter how far-fetched or future changing, has people involved one way or another. Algorithms and artificial intelligence build upon human knowledge, be it as the curators of the data sets used to teach the machines or the specialists who validate newly produced artifacts. To get to the current moment in the fine arts and art history, an invaluable amount of labor has been put in throughout time by many experts, curators, art historians, librarians, scholars, and students. This is not to say that the technological advances can be trivialized; rather, that for them to burst into the art scene the previous painstaking effort is often made invisible or even forgotten. We try to highlight how those endeavors are moving or might move into the digital domain, with a great deal of human ingenuity involved. In this chapter we describe three of the main aspects embroiled in the generation of such AI-crafted—allowing the craftsmanship metaphor—paintings: curation, content, and creation. We illustrate how each step is a rich research field in itself and provide details on specific projects that have been carried out and the ranging variety of tools used. The chapter is also meant to point to the variety of paths opening up for experts of the fine arts in the digital age.

Curation

Data-related projects in fine arts usually start with the construction of a data set, which in the humanities might also be referred to as a corpus or collection, depending on the discipline. It is important to note that this first step, the construction of the data set, represents the most burdensome phase since in most cases, and certainly in cases that deal with history prior to the digital age, researchers need to digitize the materials. If, on the contrary, one works with contemporary art it is more likely that the information has already been digitized. When the primary elements of a data set are images or digitized versions of images, metadata plays a fundamental role. Quality curation of content is an intensive and time-consuming task performed manually by humans, but once a data set has been built, its many secrets can be exploited in multiple forms until exhausted. The formats vary from formal relational databases built with schemas and models of the reality in mind to somewhat chaotic and ad hoc spreadsheets or text files.[6] Regardless, all of them attempt to capture some particular nuances of the relevant subject domain: a slice of reality in computer form from a human's vantage point. The Hispanic Baroque Project is an example of such practices.[7] With the ambitious goal of characterizing the complexity of the First Atlantic Culture, the project set to trace interactions that generated the cultural complexity that was characteristic of the Hispanic Baroque from the mid-sixteenth to early nineteenth centuries, and that allowed for their reproduction in, and transfer of models to, other cultural settings even today.[8] Methodologically, in order to analyze three spheres of Baroque culture—the Baroque constitution, its religious expressions, and its urban aspects—some ideas from complexity theory such as emergence, dynamic stability, or efficiency were applied. These concepts, once arcane to humanistic endeavors, started to make sense precisely at the intersection of various disciplines and methods, and became instrumental for the project once the use of graph theory and network science was leveraged. Thus, the project embarked on the construction of several databases and tools to better achieve its goals. One of these databases still active today is BaroqueArt, a data collection of Hispanic Baroque painters, places, and paintings from 1550 to 1850.[9] Populated through fieldwork data collection and manual data extraction from existing print catalogues, the database contains close to seventeen thousand items from fourteen hundred bibliographical resources, where each item might be annotated by up to two hundred descriptors. Since the scope of the database is Hispanic Baroque art, entries are of varying geographies and artistic styles. The selection follows historical context as well as artistic and scholarly principles such as influence and knowledge transfer. Interestingly, the aforementioned descriptors were put together using a team-based, manual curation approach to create a comprehensive ontology based on rich annotations of the main object of study: the paintings.[10] These descriptors ranged from the abstract (e.g. "virtuality") to the concrete (e.g. the kind of ink used on the physical canvas) and served as a proxy to build a network of paintings.[11] After annotation, a link was established between any pair of paintings if they shared at least five or more annotations, using this cardinality as the weight of the created edge in the graph. The resulting network, divided into periods of twenty-five years, together with the geographical information that tracked the place of production and current location of the paintings, allowed for a rich analysis that proved the importance of symbols in the creation and sustainability of the First Transatlantic Culture.[12] The study introduced a new methodology that aspired to establish the foundations of a digital geography of art and culture. At the same time, it left a database of thousands of paintings with extremely detailed metadata, from religious descriptors such as "virgins" or "saints," to the use of the "civil" descriptor that signaled more secular aspects of Latin American paintings as those nations moved toward independence. In other words, it curated a vast number of artifacts for future filtering, querying, and retrieving that could be utilized to, for example, select all paintings of male portraits in a specific style.

Content

One natural step after the analysis of metadata is to study the actual content of the digital images of paintings: that is, what it is that they depict, how this is materialized in the painting, what meanings those depictions carry with them, and what can we learn from them. In this sense, different approaches have been tested over the years. Particularly notorious and numerous are those that seek to identify—by computational means—styles in a painting, either the author's personal signature or the pictorial style within a time period.[13] Although less prolific in the literature, other avenues of research have been explored: extraction of meaning,[14] influence networks,[15] and trends over time,[16] to name a few. It is in this latter category that we can place the Faces Project, which studied how people were depicted across time and what relation those depictions had to present day evolutionary conceptions of perceived attractiveness. The main goal of the project was to assess whether human beauty had always been universally perceived in the same manner or had changed over time and across cultures.[17] This is, however, an almost unprovable issue since beauty is a changing and subjective notion and as such cannot be measured scientifically. If we add that we lack reliable photographic evidence of most periods of human history, we can agree that accessing the past of the notion of beauty is a difficult task. In addition, no representative sample of the respective populations of the past will ever be possible to obtain as we can only make inferences as to the representativeness of the human types depicted in art history across cultures. To overcome these limitations, the project considered instead a definition of perceived attractiveness valid in evolutionary psychology and neuro-aesthetics that is supposed to be universal yet comprehensible and feasible to extract numerically. It was then used as a proxy for beauty on a corpus of over 120,000 paintings from all over the world and time periods. However, for different reasons the collection ended up having a major participation of Western paintings and styles, thus limiting the scope of the results. Before the popularization of WikiArt, the current de facto paintings database,[18] collecting a data set of these characteristics was a challenge, and the project had to turn to a comprehensive but private collection of images made public on the internet.[19] Fortunately, the situation is slowly changing and we can now find cultural institutions releasing their data openly to the public and for free.[20] Paintings are possibly the only artifacts known to have reliably recorded the various aspects of human attributes before the invention of photography, and while there is no guarantee that depicted people are representative of the general population, it still is the best resource there is to track changes in the representation of the human face.[21] It is important to note that the project does not make any claims on current or past standards of beauty. Rather it applies a supposedly objective approximation of it on a faceted corpus of paintings and analyzes the results and their implications. That is to say, if we were to accept a scientific definition of perceived attractiveness, the Faces Project describes how depicted people, periods, and styles would fare against it. In assessing the perceived attractiveness of a given face, two main factors influence the outcome. First, symmetry—measured as the sum of the distances of facial attributes to the center of the face as defined by the perpendicular line that crosses it through the middle horizontal point of the nose and that divides the face in two hemi-faces. Second, averageness—a slightly more elusive concept that is defined as the difference between an individual's facial attributes and those of the average of a group. Effectively, this calculation consists in the difference between an individual's face symmetry value and that of the average composite image of all faces of a group, being a group defined by a time period or a painting style. At the technical level, all of these measurements built upon proven face recognition algorithms that provided exact facial feature coordinates within images: a couple of points for each eye, three for the mouth, and another three for the nose. Moreover, these borrowed techniques from machine learning also gave estimations of gender, age, and the head pose of the recognized faces, allowing for a demographic aggregation and rich comparison of the results.

The project showed that the representation of human faces has not remained constant and that there are substantial differences between the faces depicted between the fifteenth and eighteenth centuries when compared to those of both the thirteenth and twentieth centuries. Especially significant was the decrease in the perceived attractiveness of faces in twentieth-century paintings, as the freedom of artists and the openness of society fostered the representation of different types of human face other than those associated with classical styles.

The Faces Project also set the path for a digital demography in the fine arts as it explored demographic distributions and age pyramids for the first time. As an unexpected side effect of the project, valuable data on the demographics of people in paintings were uncovered, both by century and by painting style. Therefore, it became possible to characterize, for example, the most common types of individual depicted in a specific period of time, including the preferred age and gender in a particular painting style,[22] or the composition, poses, proportions, and spatial locations in which particular types appear.

Creation

The final step in the pipeline is the creation process itself. We now have a sense of how important the proper curation of collections is and how it enables the selection of paintings in different painting styles. Moreover, through the analysis of their content, further filters can be applied (e.g. retaining those paintings that display only one woman in a central role). Although very different methodologically, both aspects combine to produce rich labeled samples from which, theoretically, a machine can learn. However, while learning by example is one of the ways in which art can be artificially produced, all AI-generated artworks are technically under the umbrella of what is called generative art.

Generative art, in which human participation is drastically reduced if not completely eliminated from the creative process, is by no means a novel idea.[23] Although many different methods exist, both disciplinary (mechanical, chemical, biological, mathematical, robotic, etc.) as well as conceptual (emergence, evolution, embodiment, self-organization, etc.), the creation of art by autonomous systems has been historically tied to rule-based and algorithmically determined computer processes. Since its birth in the computer graphics scene in Stuttgart in 1965 with *Generative Computergraphik*, presumably the first exhibition to show generative computer art,[24] dynamic artwork-systems have evolved in parallel with the technology that made them possible. From the mid-1960s onwards, artists played with the idea of generating art by leveraging concepts from Artificial Life (e.g., L-systems, genetic algorithms, cellular automata) and Artificial Intelligence. AARON, a piece of software built over a forty-year period by the abstract painter Harold Cohen, is perhaps one of the first attempts to use artificial intelligence to produce art.[25] Its style and aesthetic dexterity have evolved over many existing versions since its conception: from basic abstract black and white compositions to the inclusion of color and figurative elements. It even debuted in an exhibition at the Tate Gallery in London, United Kingdom, in 1983, where the works are still housed after the passing in 2016 of its creator.

However, what makes AARON fundamentally different from AICAN (Artificial Intelligence Creative Adversarial Network), namely the technique used to generate the series *Faceless Portraits*

Transcending Time, is the school of thought from which the two approaches of artificial intelligence come. AARON uses what is called a symbolic approach, in which the expert of the domain must manually encode a number of rules in the algorithm that govern the creation process. By contrast, AICAN uses a connectionist approach, meaning that instead of hard coding the rules that guide the creative process, the algorithm learns the general ideas from previous examples until it is able to produce new and hitherto unseen artworks. In computer science literature, the latter method is referred to as an artificial neural network, since it is vaguely inspired by the biological neural networks in animal brains. A neural network is more a framework for a specific learning task than an algorithm itself, since no prior task-specific rules are programmed into them.

Typically, neural networks are internally structured in layers connected to each other in such a way that the input layer receives a numerical representation of an object to be classified and the output layer produces values that assign the input to one of several possible classes. Each of these layers have artificial neurons that perform some computation based on an activation function and some pre-defined weights. With some reservations, the process of learning or training is then reduced to processing the input iteratively and adjusting the weights until a certain loss function is minimized in the output layers. When the number of layers increases, we talk about deep neural networks or deep learning.[26]

In the case of image recognition by artificial intelligence, a common kind of neural net is the Convolutional Neural Network (CNN) that works by reducing images successively using a form of matrix multiplication named convolution—technically, a cross-correlation—until the resulting image can be assigned a class from a limited number of options.[27] For example, in a data set comprising exclusively images of paintings by Joan Miró and Van Gogh, a CNN might be able to extract features of the images at an increasing level of abstraction: beginning with the identification of pixels that form edges and borders; progressing to the identification of blobs and shapes; and, finally, distinguishing individual strokes, themes, or styles. Instead of feeding all this information independently, a CNN automatically generates identifying characteristics from the data sets it processes. This results in one of the best and worst qualities of this method: A neural network will always be as good (and as biased) as the data used to train it. As described, CNNs are very good at classifying images, but not very useful at generating them by themselves. One way to adapt them to the task of image creation is to run the classification process in reverse order: If the outermost layers of a CNN encode information from images by moving from the more abstract to the more concrete, it is possible to reverse the process just before the final layers and move from the dense codification of an image (or a random list of numbers for that matter) back to an image again.

This is indeed the idea behind Generative Adversarial Networks (GANs) the core mechanism inside AICAN, with the addition of a clever trick: every time a new image is generated from a random input vector by the generative network, another network (the discriminant) evaluates (classifies) the resulting image as valid or not depending on how real the image would look to a human. These two networks enter in a contest with each other in a zero-sum game. Every time the generator presents a new, more appealing candidate, the discriminator has also become better at the process of evaluation. For each misclassification, information about the difference between the generated image and a valid image is fed back to both networks. After many iterations of the training process, the images that pass the filter imposed by the discriminator are surprisingly natural looking, even to the point of deceiving the human eye.[28] These GANs learn approximations of the meaning of visual features in their own way. Sometimes this resembles the human way of understanding things, while on other occasions it is basically senseless.[29] AICAN is a variation of a GAN where the discriminator tries to assess whether generated images are of artistic value. Certainly, the approach seems to have some validity since both the creators of *The Portrait of Edmond de Belamy* and the series *Faceless Portraits Transcending Time* have received acknowledgment for their work either in monetary form or as public recognition,[30] either of which might meet the ambitions of the works' creators.

Conclusion

In this chapter, we have shown how the curation of metadata is far from a burden imposed by library and catalogue practices, but rather is a process can generate fruitful and impactful research. The proper curation of data sets might eventually be leveraged to enrich the analysis of the content of catalogued images. This, in turn, could be used to build generative models that produce pieces of artistic value or, at least, to create new data and, therefore, new research questions for the digital fine arts. This process raises some interesting questions.

The first issues that arises is the question of authorship. It is not entirely clear to whom authorial credit should be attributed in the case of AI generated works. Should it be ascribed to the numerous training samples? To the designer of the GAN? To the person who simply uses the trained model to produce new artworks? To the computer artifact itself? To all of them? This latter case has important implications and might require re-examination of our theories of the self,[31] of artistic identity, and of art market practices. The digital copyright wars of the 2000s could become insignificant when compared to upcoming battles against algorithmically generated content. Although writing in different contexts, some scholars have stated that knowing the author of a work "changes its meaning by changing its context [. . . ,] certain kinds of meaning are conferred by its membership and position in [. . .] the *oeuvre*."[32] This might have the effect of deeming computational creations as mere "computer-output" instead of paintings by virtue of the simple fact that they are computer-generated.[33] This perception often arises after discovering that the artwork in question is computationally generated, a point that affects its authenticity and hinders its aesthetic approval: The computer could generate new pieces endlessly with no human supervision or mediation.[34] Furthermore, while every generated piece is presumably unique, the ontological issue is not reasonably addressed. As David Cope has argued, "the fact that human composers die, [. . .] has consequences for aesthetic valuation: someone's *oeuvre* is valued in part because it is a unique set of works, now closed."[35] And for him this poses the question of whether we can identify the artwork, the uniqueness, behind the countless unique images that an artist's program can generate.[36] Do we have a theoretical framework to appreciate or measure these contributions? Do we need one or is it permissible to dismiss it, since humans cannot conceive of the infinite artist?

The further question arises as to the ontology of art itself. Does art need to reflect the inimitable experience of the artist? Does it need to be a manifestation of thought or emotion articulated in a particular medium? Does it need to manifest human creativity? Part of the praise for AICAN derives from the fact that it was considered, or labeled, as being more creative than regular generative adversarial networks. However, due to its unpredictable nature, the concept of creativity continues to be as hard to define as it has always been,[37] which makes it a problematic attribute to quantify and thus questionable to assign to machine outputs of any kind.[38] This chapter has sought to highlight the fact that pictorial art, either human or machine created, cannot fully escape either its predecessors (in the form of art's histories) or assessments (of aesthetic merit and ontology) agreed upon by experts in fine arts communities. The latter is of the utmost importance if the creators of AI generated works seek recognition of some sort. Nevertheless, the infinite—computer—artist has come to stay and we should be grateful if it stimulates the imagination of its creators: human artists.

Lastly, there is the persistent concern in all kinds of representational and statistical learning: intention. Computer generated art lacks a force that drives the creation of artworks. Ranging from monetary to altruistic motives, technical mastery to expressive desire, machines are not only unable to provide reasons for their decisions, but the processes that mediate creation remain opaque and deprived of interpretive or communication powers.[39] And while these issues have deep consequences in other fields (for example, if misused in policing), in the context of art it gives the impression of emptiness or purposelessness. Perhaps, a new kind of art will emerge, one that can only be appreciated or deciphered by other machines, a true machine art, made for and by machines, with no humans in the loop.

Notes

1. Paris-based arts collective *Obvious* was the organization behind the production of the piece. Christie's, "The First Piece of AI-Generated Art to Come to Auction," last modified December 12, accessed April 1, 2019, www.christies.com/features/A-collaboration-between-two-artists-one-human-one-a-machine-9332-1.aspx.
2. AICAN stands for Artificial Intelligence Creative Adversarial Network. See Ahmed Elgammal et al., "CAN: Creative Adversarial Networks, Generating 'Art' by Learning About Styles and Deviating from Style Norms" (June 21, 2017), https://arxiv.org/abs/1706.07068.
3. "AI and Machine Learning Invade a New York Art Gallery," *The Atlantic*, last modified March 6, accessed April 1, 2019, www.theatlantic.com/technology/archive/2019/03/ai-created-art-invades-chelsea-gallery-scene/584134/.
4. See, for example, Anthony O'Hear, "Art and Technology: An Old Tension," *Royal Institute of Philosophy Supplements* 38 (March 1, 1995): 143–58, doi:10.1017/S1358246100007335; http://journals.cambridge.org/abstract_S1358246100007335.
5. Jon McCormack et al., "Ten Questions Concerning Generative Computer Art," *Leonardo* 47, no. 2 (April 1, 2014): 135–41, doi:10.1162/LEON_a_00533; www.mitpressjournals.org/doi/abs/10.1162/LEON_a_00533.
6. Stephen Ramsay, "Databases," in *A Companion to Digital Humanities* (Malden, MA and Chichester: Wiley-Blackwell, 2004).
7. "The Hispanic Baroque Project [Archived]," accessed April 1, 2019, http://projects.cultureplex.ca/www.hispanicbaroque.ca/.
8. Juan-Luis Suárez and Estefanía Olid-Peña, "Hispanic Baroque: A Model for the Study of Cultural Complexity in the Atlantic World," *South Atlantic Review* 72, no. 1 (January 1, 2007): 31–47, www.jstor.org/stable/27784678.
9. "BaroqueArt," accessed April 1, 2019, http://baroqueart.cultureplex.ca/.
10. "BaroqueArt Descriptors Ontology: An Ontology to Formalize the Way an Artwork Is Described," accessed April 1, 2019, http://ontologies.cultureplex.ca/baroqueart/spec/index.html.
11. Juan-Luis Suarez, Fernando Sancho Caparrini, and Javier de la Rosa Perez, "The Art-Space of a Global Community: The Network of Baroque Paintings in Hispanic-America," *IEEE*, October 2011.
12. Juan-Luis Suarez et al., "Towards a Digital Geography of Hispanic Baroque Art," *Literary and Linguistic Computing* 28, no. 4 (December 1, 2013): 718–35, doi:10.1093/llc/fqt050.
13. See, for example, Babak Saleh and Ahmed Elgammal, "Large-Scale Classification of Fine-Art Paintings: Learning the Right Metric on the Right Feature" (2016), doi:10.11588/dah.2016.2.23376; https://search.datacite.org/works/10.11588/dah.2016.2.23376; Higor Y.D. Sigaki, Matjaž Perc, and Haroldo V. Ribeiro, "History of Art Paintings Through the Lens of Entropy and Complexity," *Proceedings of the National Academy of Sciences of the United States of America* 115, no. 37 (September 11, 2018): E8594, doi:10.1073/pnas.1800083115, www.ncbi.nlm.nih.gov/pubmed/30150384; Daan Wynen, Cordelia Schmid, and Julien Mairal, "Unsupervised Learning of Artistic Styles With Archetypal Style Analysis" (May 28, 2018), https://arxiv.org/abs/1805.11155; Ahmed Elgammal et al., "The Shape of Art History in the Eyes of the Machine" (January 23, 2018), https://arxiv.org/abs/1801.07729; Laure Thompson and David Mimno, "Computational Cut-Ups: The Influence of Dada," *The Journal of Modern Periodical Studies* 8, no. 2 (November 1, 2018): 179–95, doi:10.5325/jmodeperistud.8.2.0179, www.jstor.org/stable/10.5325/jmodeperistud.8.2.0179; Eva Cetinic and Sonja Grgic, "Automated Painter Recognition Based on Image Feature Extraction," Croatian Society Electronics in Marine—ELMAR,

September 2013. For an introduction to the topic, see Lev Manovich, "Data Science and Digital Art History," *International Journal for Digital Art History*, no. 1 (2015).

14. Byunghwee Lee et al., "Heterogeneity in Chromatic Distance in Images and Characterization of Massive Painting Data Set," *PLoS One* 13, no. 9 (September 1, 2018): e0204430, doi:10.1371/journal.pone.0204430, www.ncbi.nlm.nih.gov/pubmed/30252919.
15. Babak Saleh et al., "Toward Automated Discovery of Artistic Influence," *Multimedia Tools and Applications* 75, no. 7 (April 2016): 3565–91, doi:10.1007/s11042-014-2193-x, https://search.proquest.com/docview/1776306844.
16. Leonardo Impett and Franco Moretti, *Totentanz: Operationalizing Aby Warburg's 'Pathosformeln,'* Stanford Literary Lab, Stanford, CA, 2017.
17. Javier De la Rosa and Juan-Luis Suárez, "A Quantitative Approach to Beauty: Perceived Attractiveness of Human Faces in World Painting" (2015), doi:10.11588/dah.2015.1.21640, https://search.datacite.org/works/10.11588/dah.2015.1.21640. The limitations of the data set skewed the analysis and results toward Western paintings and conceptions of beauty.
18. "WikiArt: Visual Art Encyclopedia," accessed April 1, 2019, www.wikiart.org/.
19. Although the original site was shut down and a new domain is now available with limited functionality "Ciudad De La Pintura (City of Paintings)," accessed April 1, 2019, www.ciudadpintura.com/.
20. See, for example, the released of The MET's own collection of paintings as free and open data sets and their joint initiative with Microsoft and the MIT to support an external API for programmatic access to their data, "Open Access at the MET," accessed April 1, 2019, www.metmuseum.org/about-the-met/policies-and-documents/open-access.
21. There is some controversy about this point. On the one hand, during certain periods of history only those in a position of power could command the creation of paintings, possibly biasing the representativeness of the depicted population. It is also known that sometimes painters exerted some creative licenses adapting the faces of their models to whatever standards of beauty were reigning at the time (see Álvaro Pascual-Chenel, "Juegos De Imagen Y Apariencia: Simulación, Disimulación Y Propaganda Política Durante El Reinado De Carlos II," *El Universo Simbólico Del Poder En El Siglo De Oro* (January 1, 2012): 4.). On the other hand, even when only random people were depicted on a figurative fashion, no claim can be made on how that would relate to the general population's beauty. The Faces Project does not try to make the claim that only beautiful people were painted nor that the study can be blindly interpolated to all of human paintings, rather it just establishes a framework to assess how those that were indeed painted for posterity would be evaluated from a biologically evolutionary definition of perceived attractiveness today. Moreover, the study never states that our current standard of beauty is better or worse in any aspect to those of the past.
22. Ironically, most modern methods to estimate age and gender from pictures are now based on neural networks.
23. Alan Dorin et al., "A Framework for Understanding Generative Art," *Digital Creativity* 23, no. 3–4 (December 2012): 239–59, doi:10.1080/14626268.2012.709940.
24. Margaret A. Boden and Ernest A. Edmonds, "What Is Generative Art?" *Digital Creativity* 20, no. 1–2 (June 1, 2009): 21–46, doi:10.1080/14626260902867915, www.tandfonline.com/doi/abs/10.1080/14626260902867915.
25. See Harold Cohen, "The Further Exploits of AARON, Painter," *Stanford Humanities Review* 4, no. 2 (1995): 141–58; Harold Cohen, "A Million Millennial Medicis," in *Explorations in Art and Technology* (London: Praxis, 2002), 91–104.
26. For a good introduction see Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep Learning* (Cambridge, MA and London, UK: MIT Press, 2016).
27. Ibid., Ch. 20.
28. Tero Karras, Samuli Laine, and Timo Aila, "A Style-Based Generator Architecture for Generative Adversarial Networks" (December 12, 2018), https://arxiv.org/abs/1812.04948.
29. Shan Carter et al., "Exploring Neural Networks with Activation Atlases," *Distill* (March 6, 2019), doi:10.23915/distill.00015, https://distill.pub/2019/activation-atlas/.
30. "Faceless Portraits Transcending Time," HG Contemporary New York, accessed April 1, 2019, www.hgcontemporary.com/exhibitions/faceless-portraits-transcending-time.
31. To put it in perspective, the problem is the same that autonomous vehicles face if they run over a pedestrian. Who would be to blame in that case?
32. Harold Love, *Attributing Authorship: An Introduction* (Cambridge, MA: Cambridge University Press, 2002) 46.

33. David Cope, *Computer Models of Musical Creativity* (Cambridge, MA: MIT Press, 2005).
34. Boden and Edmonds, "What Is Generative Art?" 24.
35. Ibid., 42.
36. Ibid., 22.
37. See these two opposite views in the classic Margaret A. Boden, *The Creative Mind: Myths and Mechanisms*, 2nd ed. (London: Routledge, 2004), doi:10.4324/9780203508527, www.taylorfrancis.com/books/9781134379576; and the approach taken by some of the creators of AICAN Ahmed Elgammal and Babak Saleh, "Quantifying Creativity in Art Networks" (June 1, 2015), https://arxiv.org/abs/1506.00711.
38. Boden, *The Creative Mind*, Ch. 4.
39. Mike Ananny and Kate Crawford, "Seeing Without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability," *New Media & Society* 20, no. 3 (March 2018): 973–89, doi:10.1177/1461444816676645, https://journals.sagepub.com/doi/full/10.1177/1461444816676645.

Bibliography

Ananny, Mike, and Kate Crawford. "Seeing Without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability." *New Media & Society* 20, no. 3 (March 2018): 973–89. doi:10.1177/1461444816676645.https://journals.sagepub.com/doi/full/10.1177/1461444816676645

Boden, Margaret A. *The Creative Mind: Myths and Mechanisms*. 2nd ed. London: Routledge, 2004. doi:10.4324/9780203508527. www.taylorfrancis.com/books/9781134379576.

Boden, Margaret A., and Ernest A. Edmonds. "What Is Generative Art?" *Digital Creativity* 20, no. 1–2 (June 1, 2009): 21–46. doi:10.1080/14626260902867915. www.tandfonline.com/doi/abs/10.1080/14626260902867915

Bogost, Ian. "AI and Machine Learning Invade a New York Art Gallery." *The Atlantic*. Accessed April 1, 2019. www.theatlantic.com/technology/archive/2019/03/ai-created-art-invades-chelsea-gallery-scene/584134/

Byunghwee Lee, Daniel Kim, Seunghye Sun, Hawoong Jeong, and Juyong Park. "Heterogeneity in Chromatic Distance in Images and Characterization of Massive Painting Data Set." *PLoS One* 13, no. 9 (September 1, 2018): e0204430. doi:10.1371/journal.pone.0204430. www.ncbi.nlm.nih.gov/pubmed/30252919

Carter, Shan, Zan Armstrong, Ludwig Schubert, Ian Johnson, and Chris Olah. "Exploring Neural Networks with Activation Atlases." *Distill* (March 6, 2019). doi:10.23915/distill.00015. https://distill.pub/2019/activation-atlas/

Cetinic, Eva, and Sonja Grgic. "Automated Painter Recognition Based on Image Feature Extraction." Croatian Society Electronics in Marine—ELMAR, September 2013.

Christie's. "The First Piece of AI-Generated Art to Come to Auction." Accessed April 1, 2019. www.christies.com/features/A-collaboration-between-two-artists-one-human-one-a-machine-9332-1.aspx

"Ciudad De La Pintura (City of Paintings)." Accessed April 1, 2019. www.ciudadpintura.com/

Cohen, Harold. "The Further Exploits of AARON, Painter." *Stanford Humanities Review* 4, no. 2 (1995): 141–58.

———. "A Million Millennial Medicis." In *Explorations in Art and Technology*, 91–104. London: Praxis, 2002.

Cope, David. *Computer Models of Musical Creativity*. Cambridge, MA: MIT Press, 2005.

The CulturePlex Lab. "BaroqueArt." Accessed April 1, 2019. http://baroqueart.cultureplex.ca/.

———. "BaroqueArt Descriptors Ontology: An Ontology to Formalize the Way an Artwork Is Described." Accessed April 1, 2019. http://ontologies.cultureplex.ca/baroqueart/spec/index.html.

De la Rosa, Javier, and Juan-Luis Suárez. "A Quantitative Approach to Beauty. Perceived Attractiveness of Human Faces in World Painting." 2015. doi:10.11588/dah.2015.1.21640. https://search.datacite.org/works/10.11588/dah.2015.1.21640

Dorin, Alan, Jonathan McCabe, Jon McCormack, Gordon Monro, and Mitchell Whitelaw. "A Framework for Understanding Generative Art." *Digital Creativity* 23, no. 3–4 (December 2012): 239–59. doi:10.1080/14626268.2012.709940

Elgammal, Ahmed, and AICAN. "Faceless Portraits Transcending Time." HG Contemporary New York. Accessed April 1, 2019. www.hgcontemporary.com/exhibitions/faceless-portraits-transcending-time

Elgammal, Ahmed, Bingchen Liu, Mohamed Elhoseiny, and Marian Mazzone. "CAN: Creative Adversarial Networks, Generating 'Art' by Learning About Styles and Deviating from Style Norms" (June 21, 2017). https://arxiv.org/abs/1706.07068

Elgammal, Ahmed, Marian Mazzone, Bingchen Liu, Diana Kim, and Mohamed Elhoseiny. "The Shape of Art History in the Eyes of the Machine" (January 23, 2018). https://arxiv.org/abs/1801.07729.
Elgammal, Ahmed, and Babak Saleh. "Quantifying Creativity in Art Networks." (June 1, 2015). https://arxiv.org/abs/1506.00711.
The Hispanic Baroque Project. "The Hispanic Baroque Project [Archived]." Accessed April 1, 2019. http://projects.cultureplex.ca/www.hispanicbaroque.ca/.
Impett, Leonardo, and Franco Moretti. *Totentanz: Operationalizing Aby Warburg's "Pathosformeln."* Stanford, CA: Stanford Literary Lab, 2017.
Karras, Tero, Samuli Laine, and Timo Aila. "A Style-Based Generator Architecture for Generative Adversarial Networks" (December 12, 2018). https://arxiv.org/abs/1812.04948.
Love, Harold. *Attributing Authorship: An Introduction*. Cambridge: Cambridge University Press, 2002.
Manovich, Lev. "Data Science and Digital Art History." *International Journal for Digital Art History* no. 1 (2015).
McCormack, Jon, Oliver Bown, Alan Dorin, Jonathan McCabe, Gordon Monro, and Mitchell Whitelaw. "Ten Questions Concerning Generative Computer Art." *Leonardo* 47, no. 2 (April 1, 2014): 135–41. doi:10.1162/LEON_a_00533. www.mitpressjournals.org/doi/abs/10.1162/LEON_a_00533.
O'Hear, Anthony. "Art and Technology: An Old Tension." *Royal Institute of Philosophy Supplements* 38 (March 1, 1995): 143–58. doi:10.1017/S1358246100007335. http://journals.cambridge.org/abstract_S1358246100007335.
Ramsay, Stephen. "Databases." In *A Companion to Digital Humanities*, edited by Susan Schreibman, Ray Siemens, and John Unsworth. Malden, MA and Chichester: Wiley-Blackwell, 2004.
Saleh, Babak, Kanako Abe, Ravneet Arora, and Ahmed Elgammal. "Toward Automated Discovery of Artistic Influence." *Multimedia Tools and Applications* 75, no. 7 (April 2016): 3565–91. doi:10.1007/s11042-014-2193-x. https://search.proquest.com/docview/1776306844.
Saleh, Babak, and Ahmed Elgammal. "Large-Scale Classification of Fine-Art Paintings: Learning the Right Metric on the Right Feature." 2016. doi:10.11588/dah.2016.2.23376. https://search.datacite.org/works/10.11588/dah.2016.2.23376.
Sigaki, Higor Y. D., Matjaž Perc, and Haroldo V. Ribeiro. "History of Art Paintings Through the Lens of Entropy and Complexity." *Proceedings of the National Academy of Sciences of the United States of America* 115, no. 37 (September 11, 2018): E8594. doi:10.1073/pnas.1800083115. www.ncbi.nlm.nih.gov/pubmed/30150384
Suárez, Juan-Luis, and Estefanía Olid-Peña. "Hispanic Baroque: A Model for the Study of Cultural Complexity in the Atlantic World." *South Atlantic Review* 72, no. 1 (January 1, 2007): 31–47. www.jstor.org/stable/27784678.
Suarez, Juan-Luis, Fernando Sancho-Caparrini, Elika Ortega, Javier de la Rosa, Natalia Caldas, and David Brown. "Towards a Digital Geography of Hispanic Baroque Art." *Literary and Linguistic Computing* 28, no. 4 (December 1, 2013): 718–35. doi:10.1093/llc/fqt050.
Suarez, Juan-Luis, Fernando Sancho-Caparrini, and Javier de la Rosa Perez. "The Art-Space of a Global Community: The Network of Baroque Paintings in Hispanic-America." *IEEE*, October 2011.
Thompson, Laure, and David Mimno. "Computational Cut-Ups: The Influence of Dada." *The Journal of Modern Periodical Studies* 8, no. 2 (November 1, 2018): 179–95. doi:10.5325/jmodeperistud.8.2.0179. www.jstor.org/stable/10.5325/jmodeperistud.8.2.0179.
Wynen, Daan, Cordelia Schmid, and Julien Mairal. "Unsupervised Learning of Artistic Styles with Archetypal Style Analysis" (May 28, 2018). https://arxiv.org/abs/1805.11155.