

## **ISSUE STATEMENT**

The effort to utilize machine learning algorithms is becoming widely adopted by visual arts collections as a method to improve access to collections as digitization, born-digital materials, and online open-access efforts continually increase. By spotlighting the high levels of bias and low levels of transparency that often accompany algorithmic object description models, I aim to argue best practices for using machine learning in large visual arts collections.

## ISSUE PAPER

*Algorithmic Description: Using Machine Learning for Arts-Based Collection Metadata*

April 2021

---

### *Introduction*

Information professionals across various types of collection institutions, libraries, archives, and museums, are often faced with a similar issue. The constant and incredible struggle of an ever-increasing backlog of materials to process, which swells as digitization efforts and born-digital materials proliferate. Limitations of time, funds, and labor continue to block accessibility to collection objects, even as acquisitions may continue. Employment of machine learning for processing digital image collections presents itself as an emerging development and one possible solution to this stagnation. Particularly, computer vision, a subcategory of machine learning, is starting to be implemented in visual arts collections to alleviate long accumulations of unprocessed materials. The use of computer vision in arts-based collections automates digital image analysis and processing to increase metadata description in the hope for its subsequent accessibility. Currently, the growing integration of this new trend in arts collections continues, and broader adoption of the technology appears to be widely embraced by many collection stewards.

From this perspective, the connection of LAM collections and machine learning seems naturally appropriate. The AI technology is apt to process and analyze large amounts of data, like the data held by large collection repositories. Yet, as intrigue for this application of machine learning increases, there are significant issues to address. As recent research studies and literature suggest, machine learning algorithms are known to reflect high levels of bias and low levels of transparency. How can information

professionals working in LAMs confront these potentially harmful effects of machine learning during their use of it? Many institutions that could benefit from automated assistance in collections processing are also institutions which hold equity and inclusivity at the core of their mission. Is there a way to tackle both of these priorities so that institutions can implement responsible collection description and earlier access to materials through the use of machine learning algorithms? Through analysis of recent case studies, interview accounts, and literature, this essay proposes visual arts LAMs can mitigate algorithmic bias by promoting transparency of computer vision models, demonstrating caution, and establishing accountability. The following discussion will introduce computer vision, its implementation in visual arts collections, and follow with the considerations of its ethical application.

### *What is Computer Vision?*

Computer vision is a form of artificial intelligence that automates digital image analysis through a trained machine learning algorithm and convolutional neural network (CNN) that is able to capture unlabeled images and form judgments about their features and qualities. Computer vision as supervised machine learning means its competency is dependent on the training dataset it learns from, as it makes conclusions based on what it was trained to be correct. The training dataset requires a lot of data, so that the machine learning algorithm can effectively teach itself about the context, similarities, and differences of visual data. The CNN model is a deep learning algorithm that captures the input image, breaks down its pixels, assigns significance based on its learnable weights and biases, makes predictions, and checks the accuracy of those

predictions.<sup>1</sup> As Adam Greenfield explains, “the first goal of machine learning is to teach an algorithm how to generalize. A sound algorithm is one that is able to derive a useful classifier for something it hasn’t encountered from the things it has been shown.”<sup>2</sup> Once the machine learning algorithm has taught itself to detect patterns, it’s able to recognize shapes, objects, colors, as well as more nuanced features, such as faces and sentiment. The application of computer vision for visual art collections in LAMs may help to identify descriptive attributes of a work such as its subjects, style, composition, genre, creator, context, relationships, and even authenticity, with minimum human involvement.<sup>3</sup> If done carefully and conscientiously, the algorithm’s ability to identify such characteristics of a digital image can allow it to execute analyses and make certain determinations about it that can source new understandings, reveal connections, and benefit its overall metadata for discoverability and access.<sup>4</sup>

Since the early experimentation of computer vision began in 1959 by two neurophysiologists, David Hubel and Torsten Wiesel, its application has reached far outside of the library and information science field.<sup>5</sup> In 1974, when the technology of optical character recognition (OCR) and intelligent character recognition (ICR) were developed, the possibilities of applying computer vision widened to include documentation and recognition of vehicle plates, mobile payments, invoices, etc.<sup>6</sup> In 2001, the first face recognition applications was introduced by Paul Viola and Michael

---

<sup>1</sup> Sudeepti Surapaneni, Sana Syed, and Logan Yoonhyuk Lee, “Exploring Themes and Bias in Art using Machine Learning Image Analysis,” *2020 Systems and Information Engineering Design Symposium (SIEDS)*, (2020): 1-6, doi: 10.1109/SIEDS49339.2020.9106656.

<sup>2</sup> Adam Greenfield, *Radical Technologies: The Design of Everyday Life*, (London: Verso, 2017), 217.

<sup>3</sup> Brendan Ciecko, “6 Ways That Machine Vision Can Help Museums,” *Cuseum* (blog), last modified March 10, 2016, <https://cuseum.com/blog/6-ways-that-machine-vision-can-help-museums>.

<sup>4</sup> Babak Saleh, Kanako Abe, Ravneet Singh Arora, and Ahmed Elgammal, “Toward Automated Discovery of Artistic Influence,” *Multimedia Tools and Applications* 75 (2016): 3565–3591, <https://doi.org/10.1007/s11042-014-2193-x>.

<sup>5</sup> D. H. Hubel and T.N. Wiesel, “Receptive Fields of Single Neurones in the Cat’s Striate Cortex.” *The Journal of Physiology* 148 (1959), doi: 10.1113/jphysiol.1959.sp006308.

<sup>6</sup> “Computer Vision,” IBM, accessed December 2020, <https://www.ibm.com/topics/computer-vision>.

Jones,<sup>7</sup> and throughout the 2000s, the procedure of how visual datasets are annotated emerged with the development of new CNN models and image data sets.

### *Applications*

How well different CNNs perform on art and art historical images has been recently researched by numerous scholars, including Sudeepti Surapaneni, Sana Syed, Logan Yoonhyuk Lee (University of Virginia School of Data Science), Sean Yang, et. al. (University of Washington), and Adrian Lecoutre, Benjamin Negrevergne, Florian Yger. These three studies identify the most extensively used CNNs in image classification tasks as ResNet 50, ResNet 101, Inception-Resnet-V2, and AlexNet, with the latter two displaying the highest performance. Frequently used datasets in these research studies are ImageNet (over 14 million images), The Metropolitan Museum of Art's online collection (375,000 images), WikiArt (140,000 images), and Artsy (27,000 images). In the commercial world, some of the most prominent computer vision tools are being developed and sold by Microsoft, IBM, and Google.

In mid-October 2020, Microsoft announced the recent computer vision developments of their AI product, Azure Cognitive Services, stating, "our Computer Vision image captioning capability now describes pictures as well as humans do."<sup>8</sup> The bold statement was followed by their claims of improved content discoverability, text extraction, image, and video analysis in the effort to advance visual data processing. How exactly Microsoft's product works requires some uncovering and specialized

---

<sup>7</sup> P. Viola and M. Jones, "Rapid Object Detection Using a Boosted Cascade of Simple Features," *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2001): 511-518, doi: 10.1109/CVPR.2001.990517.

<sup>8</sup> John Roach, "What's That? Microsoft's Latest Breakthrough, Now in Azure AI, Describes Images as Well as People Do," *The AI Blog, Microsoft*, October 14, 2020, <https://blogs.microsoft.com/ai/azure-image-captioning/>.

knowledge, as their most public-facing explanation is rather imprecise; vaguely asserting their algorithm “pulls from a rich ontology of more than 10,000 concepts and objects to generate value” without any immediate further details regarding its dataset’s source or process.<sup>9</sup>

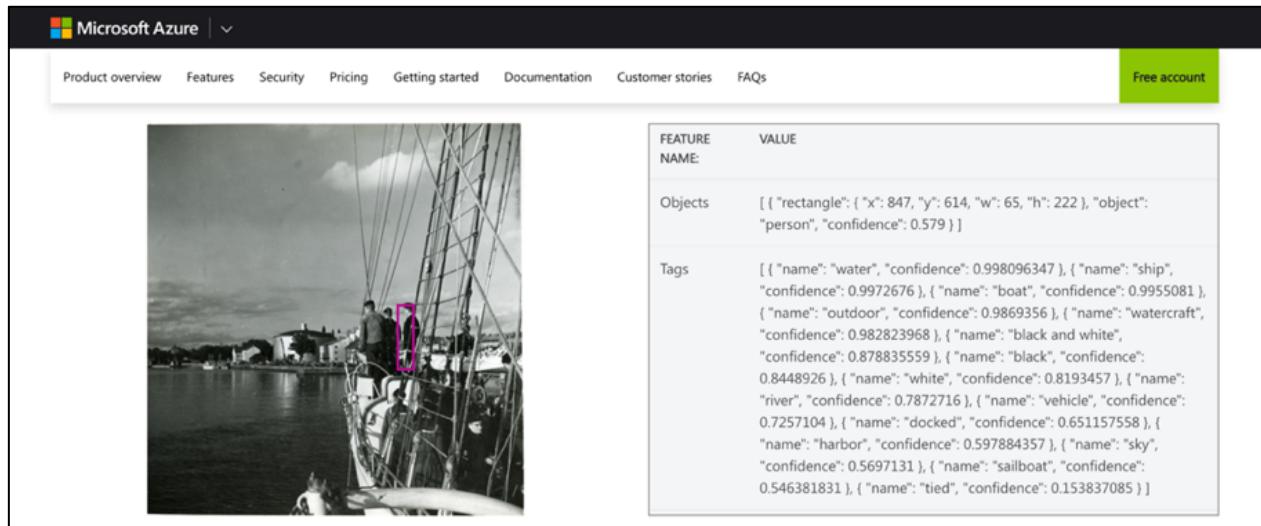


Figure 1. Microsoft Azure results for *Övningsfartyget Gladan anlöper Visby hamn*. 1947. Photograph. Accessed through the DigitaltMuseum. <https://digitaltmuseum.se/011014890070>.

IBM has also developed computer vision products and services, such as their Watson Visual Recognition product and available code patterns for classifying works of art; although, with much more accessible statements about their methods and with greater emphasis on education. IBM maintains their Trusted AI campaign, “AI Explainability 360”, which offers open-source information to guide users in understanding machine learning and elucidate how their models determine and assign labels. “Black box machine learning models that cannot be understood by people are achieving impressive accuracy on various tasks. However, as machine learning is increasingly used to inform high stakes decisions, explainability and interpretability of the models is becoming essential.” Although IBM’s commitment to transparency and

<sup>9</sup> “Computer Vision,” Microsoft, accessed December 2020, <https://azure.microsoft.com/en-us/services/cognitive-services/computer-vision/#features>.

education in AI is a worthy effort, the greatest amount of attention and success has been attributed to the computer vision products developed by a chief competitor in AI innovation, Google.

The concept of computer vision and art has been popularized by the Google Arts and Culture Lab, which has gained remarkable attention for their AI projects using visual collections from approximately 1,000 institutions around the world.<sup>10</sup> Google's computer vision product, Google Vision, has been experimented with by several prominent visual art collections, including by The Met,<sup>11</sup> The Getty,<sup>12</sup> MoMA,<sup>13<sup>14</sup></sup> LACMA<sup>15</sup>, and a number of others, such as the Harvard Art Museums, Cleveland

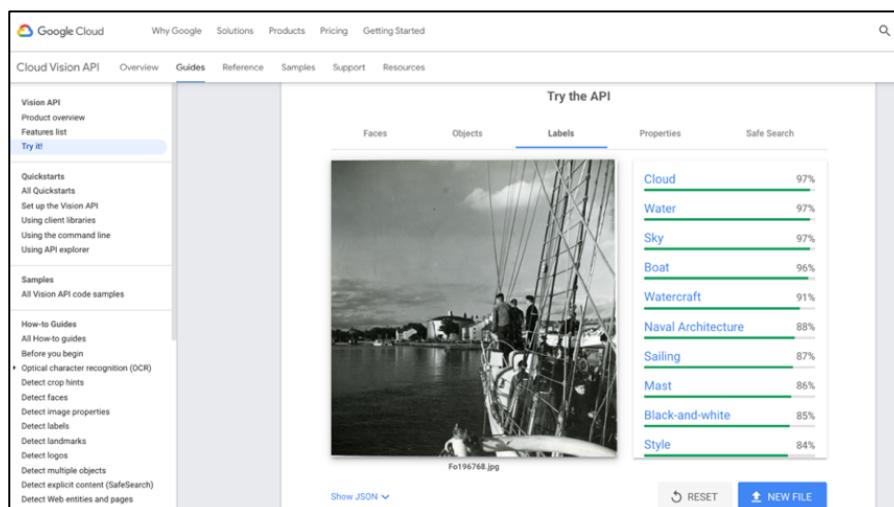


Figure 2. Google's Cloud Vision API results for *Övningsfartyget Gladan anlöper Visby hamn*. 1947. Photograph. Accessed through the DigitaltMuseum. <https://digitaltmuseum.se/011014890070>.

<sup>10</sup> Amit Sood, "Every Piece of Art You've Ever Wanted to See—Up Close and Searchable," Google Arts & Culture, uploaded June 19, 2016, accessed December 2020, <https://www.youtube.com/watch?v=CjB6DQGalU0>.

<sup>11</sup> Sarah Robinson, "When Art Meets Big Data: Analyzing 200,000 Items from The Met Collection in BigQuery," *Google Cloud* (blog), last modified August 7, 2017, <https://cloud.google.com/blog/products/gcp/when-art-meets-big-data-analyzing-200000-items-from-the-met-collection-in-bigquery>.

<sup>12</sup> Nathaniel Deines, "Does It Snow in L.A.? What Computer Vision Saw in Ed Ruscha's Sunset Boulevard," *Getty Iris* (blog), *Getty Museum*, October 7, 2020, <http://blogs.getty.edu/iris/does-it-wq2wsnow-in-la/>.

<sup>13</sup> "MoMA & Machine Learning," *Experiments with Google* (blog), *Google Arts & Culture*, March 2018, <https://experiments.withgoogle.com/moma>.

<sup>14</sup> "Identifying Art Through Machine Learning: A Project with Google Arts & Culture Lab" MoMA, accessed December 2020, <https://www.moma.org/calendar/exhibitions/history/identifying-art>.

<sup>15</sup> Sarah Pham, email message to author, March 26, 2021.

Museum of Art, The Barnes Foundation, and Auckland Art Gallery,<sup>16</sup> establishing it as the most common computer vision tool for visual art collections. For institutions such as these, machine learning—and more specifically, computer vision—has become an aspect of collection management worth supporting and sharing.<sup>17</sup>

Likewise, in academia, we see machine learning and computer vision being heavily researched, worked towards, and applied across a variety of disciplines. In Art History, recent studies by Eva Cetinic, Tomislav Lipic, and Sonja Grgic, "Learning the Principles of Art History with Convolutional Neural Networks" introduces CNN models to predict the visual features of Heinrich Wölfflin's five key visual principles: linear/painterly, planar/recessional, closed form/open form, multiplicity/unity, absolute clarity/relative clarity.<sup>18</sup> Ahmed Elgammal, et. al. similarly looks at the way machine classification of art styles can relate to art historian's approaches to analyzing style, in their article "The Shape

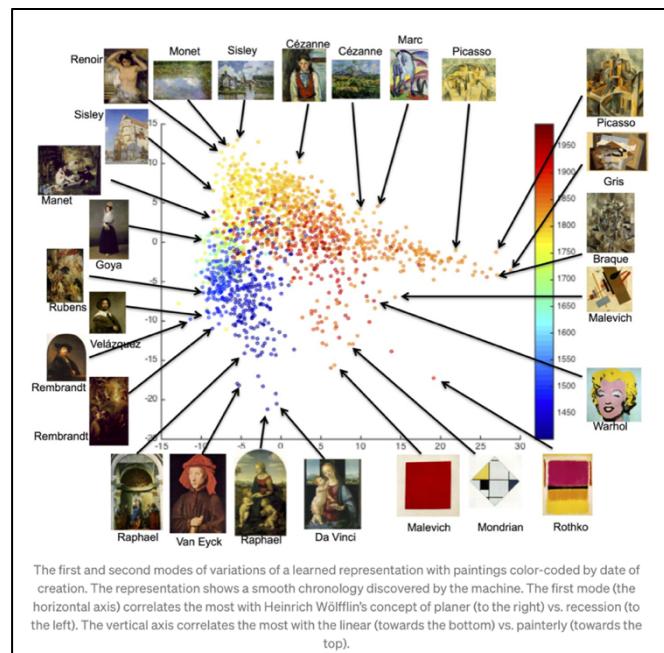


Figure 3. Ahmed Elgammal, Bingchen Liu, Diana Kim, Mohamed Elhoseiny, and Marian Mazzone. 2018. "The Shape of Art History in the Eyes of the Machine". *Proceedings of the AAAI Conference on Artificial Intelligence* 32 no. 1.

<sup>16</sup> Brendan Ciecko, "AI Sees What? The Good, the Bad, and the Ugly of Machine Vision for Museum Collections," *The Museum Review* 5, no. 1 (January 2020), [http://articles.themuseumreview.org/tmr\\_vol5no1\\_ciecko](http://articles.themuseumreview.org/tmr_vol5no1_ciecko).

<sup>17</sup> Melissa Gill and Nathaniel Deines, interview by Jessica Craig, April 2, 2021.

<sup>18</sup> Eva Cetinic, Tomislav Lipic, and Sonja Grgic, "Learning the Principles of Art History with Convolutional Neural Networks," *Pattern Recognition Letters*, 129 (2020): 56-62. <https://doi.org/10.1016/j.patrec.2019.11.008>.

of Art History in the Eyes of the Machine.”<sup>19</sup> Zhu et. al.’s “Machine: The New Art Connoisseur” and Saleh, et. al., “Toward automated discovery of artistic influence” examine art styles and influence through the lens of machine learning conclusions.<sup>20</sup>

In library and information science (LIS), the topic has emerged to become a dominant area of focus among informatics-based scholars concerned with big data, algorithms, and coded biases, such as Safiya Noble, Thomas Padilla, and Ryan Cordell. Brendan Ciecko describes machine learning have increasingly been established in LAMs in many ways, including analyzing patron use, promoting outreach, and, of course, collection processing.<sup>21</sup> Generally, the ongoing dissemination of case studies has resulted in inspiration within the broader LIS field that stimulates curiosity and interest in computer vision’s potential effects for other related collection environments. While its wide implementation is still emerging, current interest is present and rising. A recent OCLC research study conducted by Thomas Padilla from March to August 2019 found that libraries expressed a high level of interest in machine learning and algorithmic solutions for a range of reasons, including increasing efficient collection description, discoverability, and access, along with the prospect of freeing up employee time to meet other shifting demands.<sup>22</sup> While not without concern, however, the idea that a catalog record produced by AI is better than having no record at all was well received.

---

<sup>19</sup> Ahmed Elgammal, Bingchen Liu, Diana Kim, Mohamed Elhoseiny, and Marian Mazzone, “The Shape of Art History in the Eyes of the Machine,” *Proceedings of the AAAI Conference on Artificial Intelligence* 32 no. 1 (2018) <https://ojs.aaai.org/index.php/AAAI/article/view/11894>.

<sup>20</sup> Yucheng Zhu, Yanrong Ji, Yueying Zhang, Linxin Xu, Aven Le Zhou, and Ellick Chan, “Machine: The New Art Connoisseur” *Cornell University arXiv preprint* (2019), <https://arxiv.org/pdf/1911.10091.pdf>.

<sup>21</sup> Brendan Ciecko, “Examining the Impact of Artificial Intelligence in Museums,” *MW17: MW 2017*, last modified February 1, 2017, <https://mw17.mwconf.org/paper/exploring-artificial-intelligence-in-museums/>.

<sup>22</sup> Thomas Padilla, “Responsible Operations: Data Science, Machine Learning, and AI in Libraries,” *OCLC Research*, (December 2019): 6-22, <https://doi.org/10.25333/xk7z-9g97>.

This prospect of advancing collection accessibility through fast description may seem familiar to those in the archival field, to recall the popular dialog that Mark Greene and Dennis Meissner sparked with their controversial 2005 article in the *American Archivist*, “More Product, Less Process.” We’ve since learned that higher-level description for earlier access is not always in the best interest of the users, workers, or subjects in the collection. But if machine-created description can effectively describe materials at the item-level, it may be a solution that finally solves the long-time MPLP debate among archivists; as the AI-based catalog might reply, “why not both?”<sup>23</sup> However, an answer to such a question requires careful and critical thought about the ranging ethical implications of introducing AI to image-based collections.

While progress and efficiency through computer vision may hold some promise for visual art collection description and access, there are concerns worth addressing regarding its financial expense, impact on labor, and overall effectiveness; although, the two most significant issues discussed here are the probability of algorithmic biases and lack of transparency. Critical inquiry of machine learning and computer vision has found bias and obscurity of algorithmic models to be forefront concerns in immediate need of examination. The missions of LAMs are often centered on principles dedicated to ensuring equitable access and use of their collections. Can machine learning, as a technology known to be biased and opaque, contribute to such an objective? Based on the continual movement towards its adoption in LAMs, there seems to be an affirmative position regarding this, thus a more useful question may be how LAMs can manage these issues to mitigate harm for their collections and users.

---

<sup>23</sup> Mark Greene and Dennis Meissner, “More Product, Less Process: Revamping Traditional Archival Processing,” *The American Archivist* 68, no. 2 (September 1, 2005): 208–263, <https://doi.org/10.17723/aarc.68.2.c741823776k65863>.

## *Algorithmic Bias*

Computation and technology are often mistaken as neutral and more objective than human-derived thought. However, scholars such as Cathy O’Neil, Joy Buolamwini, Safiya Noble, Kate Crawford, have discredited this misconception with extensive research and examination. Multiple investigative reports have found machine learning algorithms to reflect the beliefs and values of their developers and, consequently, be just as capable of partiality and subjectivity as humans themselves. Whether intended or not, O’Neil describes, “Models are opinions embedded in mathematics.”<sup>24</sup> The viewpoint here is not asserting algorithmic bias is problematic just because it can create inaccurate annotations (eg. labelling a static image with house, tree, bicycle, etc.), but rather a deeper, much more severe problem with greater consequences. Biased algorithms are able to perpetuate prejudice, bigotry, and racial profiling, just as humans do. This is clearly explained in Noble’s book, *Algorithms of Oppression: How Search Engines Reinforce Racism*, where there are multiple examples of Google’s algorithms performing racist search outputs against Black girls and people of color.

As machine learning and computer vision gradually become commonplace in collections, this truth should not be disregarded. The visual collections held by LAMs often hold cultural, social, religious, and political significance, acting as emblems of practice, tradition, and ideology for communities unique in nature and origin. The responsibility of LAMS to care for such objects extends beyond their physical or material state, but also to their digital representation and access. If computer vision as a machine learning technology is inherently biased just as humans are, then the initial

---

<sup>24</sup> Cathy O’Neil, *Weapons of Math Destruction*, (Broadway Books, 2016).

utilization of it is more likely to be successful if the efforts are spent on managing bias, rather than hopelessly trying to eliminate it. To reduce algorithmic bias for LAMs using computer vision, demonstrating caution and accountability are fundamental first principles of the practice.

The need for caution by LAMs begins at their first point of interaction with computer vision since the operation of it demands one requirement at the start: massive amounts of data to train the algorithm. A crucial note is that datasets are just as prone to bias as the algorithms are, because it is where the bias is derived (after being derived from its creator). If a training dataset is too small or homogenous, what the algorithm can learn will be restricted to the narrow contents of the dataset. The machine learning algorithm only has the ability to perceive what it was taught to learn based on the dataset, so the dataset's contents must be highly varied and diverse. The alternative would lead to a highly biased output. The result of biased datasets can be detrimental, as Ryan Cordell states, "The biases, limitations, and oversights of those datasets will produce flawed research that does not represent the communities libraries seek to serve."<sup>25</sup> To ensure equitable service, the importance of caution should be instilled right away, whether the model is being created or adopted. For computer vision products that allow users to create their own algorithm model (like the ones developed by Microsoft and IBM), attentiveness will be required when choosing and introducing the original training dataset; whereas with products which instead presents their established model for adoption (like Google Vision), the attention will need to be

---

<sup>25</sup> Ryan Cordell, "Machine Learning + Libraries A Report on the State of the Field," *Library of Congress*, (July 14, 2020): i-86, <https://labs.loc.gov/static/labs/work/reports/Cordell-LOC-ML-report.pdf?loclr=blogsig.13>.

directed towards investigating the model's pre-existing sources and predetermined make-up.

Examining the authority of established computer vision models could expose insights into what Joy Buolamwini calls, "the coded gaze," which she describes as "a view that posits any technology created by humans will reflect individual or collective values, priorities and if unchecked, prejudices"<sup>26</sup> and situates it as "the embedded views that are propagated by those who have the power to code systems."<sup>27</sup> The coded gaze is a necessary consideration for visual arts collections, which likely hold objects depicting people and faces within their content. The concept of the "looking gaze" is familiar to those in the visual arts, as the feminist theory of the male gaze has dominated art theory, criticism, and movements since Laura Mulvey first coined the phrase in 1975. The coded gaze may also be related to bell hooks' oppositional gaze, which reestablishes the feminist theory as a strategy for Black women to engage critically with mass media representation.<sup>28</sup> Comparably, Buolamwini argues the coded gaze is behind technology applications and models that discriminate against historically marginalized populations; Black women and women of color in particular.<sup>29</sup> There is also evidence of computer vision applications misgendering the subjects of artworks. In results of Surapaneni, Syed, and Lee's research, "the model mislabeled 296 males as females and 363 females as males... a Native American male with long hair, wearing a traditional

---

<sup>26</sup> Joy Buolamwini, "Gender Shades: Intersectional Phenotypic and Demographic Evaluation of Face Datasets and Gender Classifiers." PhD diss., Massachusetts Institute of Technology, 2017.

<sup>27</sup> Joy Buolamwini "InCoding – In the Beginning Was the Coded Gaze," *MIT Media Lab* (blog) *Medium*, May 16, 2016, <https://medium.com/mit-media-lab/incoding-in-the-beginning-4e2a5c51a45d>.

<sup>28</sup> bell hooks, "The Oppositional Gaze: Black Female Spectators," *The Feminism and Visual Culture Reader* (2003): 94-105.

<sup>29</sup> Safiya Noble, *Algorithms of Oppression*, (New York: New York University Press, 2018).

regalia including a breechcloth which resembles a dress was incorrectly classified as a female given those specific features.”<sup>30</sup>

Even while Padilla’s research highlights optimistic impacts of machine learning in library collections, he discusses the importance of evaluating positive impacts against the negatives, stating positives “must be weighed relative to a broader field of misuse spanning applications that lack the ability to recognize the faces of people of color, that discriminate based on color, and that foster a capacity for discrimination based on sexuality.”<sup>31</sup> Therefore, careful attention to the inter-workings of the computer vision application prior to its adoption is essential to ensure proper use. If adopted, caution is no longer a sufficient sole priority, but should be accompanied by transparency and accountability.

The need for transparency and accountability is another aspect of managing bias in algorithmic models. Unfortunately, as O’Neil’s research suggests, transparent models are rare. “Opaque and invisible models are the rule, and clear ones very much the exception.”<sup>32</sup> With the lack of transparency, managing bias and accountability remains difficult and problematic. In the collaborative article by Sarah Myers West, Meredith Whittaker, and Kate Crawford, their *Recommendations for Addressing Bias and Discrimination in AI Systems*, states “Remedying bias in AI systems is almost impossible when these systems are opaque. Transparency is essential and begins with tracking and publicizing where AI systems are used, and for what purpose.”<sup>33</sup> For LAM institutions

---

<sup>30</sup> Sudeepti Surapaneni, Sana Syed, and Logan Yoonhyuk Lee, “Exploring Themes and Bias in Art using Machine Learning Image Analysis,” *2020 Systems and Information Engineering Design Symposium (SIEDS)*, (2020): 1-6, doi: 10.1109/SIEDS49339.2020.9106656.

<sup>31</sup> Padilla, “Responsible Operations,” 9.

<sup>32</sup> O’Neil, *Weapons of Math Destruction*.

<sup>33</sup> Sarah Myers West, Meredith Whittaker, and Kate Crawford, “Discriminating Systems: Gender, Race and Power in AI,” *AI Now Institute* (April 2019): 4, <https://ainowinstitute.org/discriminatingsystems.html>.

whose services are often based on inclusive use and participation, there is the potential for new and more transparent methods of machine learning.

### *Recommendations*

When using computer vision for collections description in a LAM context, there are several ways transparency and accountability can be practically demonstrated. First, inform users that computer vision is used for generating collection metadata and description. As a basic level of transparency, users should be cognizant about the use of computer vision and machine learning for object description. Additionally, providing further information about how computer vision operates in the collection is beneficial for their potential inquiry into the process. Second, information regarding the origin of the training dataset used for the collection should be made easily accessible. Providing readily accessible information about how the computer vision algorithm was trained based on a dataset adds another layer of transparency. Offering educational materials, contact information, and further resources on the explainability of AI in the collection will support their informed use. Third, allow users to control the parameters of the computer vision algorithm applied to their search when exploring a collection.<sup>34</sup> The ability for users to set limitations on the function of AI in their search should be offered. Assumptions about the advantages of AI-generated description should not be forcibly applied to users who may not want to use it in their research discovery process. While this may reduce their searching capabilities, at minimum, it will reveal how much the collection relies on AI for their collection metadata, which allows users the choice to compare and contrast the impact of AI on collection descriptions. Fourth, apply a

---

<sup>34</sup> Padilla, "Responsible Operations," 10.

tracking feature to object description that allows users to see the alterations of an object's metadata over time. This form of transparency is useful whether the change in the description was caused by a computer or a human, however especially with an algorithmic-induced change, since the algorithmic model should be constantly learning and updating its outputs.<sup>35</sup> This capability contributes to O'Neil's notion of a trustworthy model, "Whatever they learn, they can feed back into the model, refining it... They maintain a constant back-and-forth with whatever world they're trying to understand or predict."<sup>36</sup> This is especially true with the beginning stages of computer vision description in collections since it will likely improve with time. Tracking and publishing those changes keeps the model transparent and its implementation accountable. And fifth, catalog records should list the percentage of certainty assigned to auto-generated metadata descriptions. Revealing the level of certainty given to an auto-tag will benefit how users interpret the description. "Attempts to use algorithmic methods to describe collections must embrace the reality that, like human descriptions of collections, machine descriptions come with varying measures of certainty."<sup>37</sup> Uncertainty in description is inevitable and knowledge of it is crucial for transparency. The ability for error is present whether the description is assigned by a human or an algorithm, however, the reality of it becomes more apparent when the varying levels of accuracy and certainty are acknowledged.

The practical formation of any or all of these five points will contribute to the transparency and accountability of machine learning and computer vision in LAM visual collections. Each point recognizes the possibility of algorithmic bias but

---

<sup>35</sup> Padilla, "Responsible Operations," 10.

<sup>36</sup> O'Neil, *Weapons of Math Destruction*.

<sup>37</sup> Padilla, "Responsible Operations," 13.

addresses it by providing ways for users to diminish its impact. Further investigation is required regarding the true effectiveness of these guidelines; however, the result will likely be impactful and set a precedent regarding how LAM collections implement AI for both their collections and user's advantage.

### *Conclusion*

The attention surrounding machine learning and computer vision has burgeoned in the last decade, encouraging several visual arts LAMs to adopt it for their collection analysis and description. Their success with the emerging technology maintains a great amount of potential for broader applications that could benefit efficiency, progress, and workloads in and across LAM institutions. However, the issues of algorithmic bias and transparency should be taken into account by LAMs, both prior to its adoption and afterward. Demonstration of caution and accountability should be evident throughout the process to minimize harmful effects on the collections, workers, and users. With the practical application of such principles and priorities, machine learning and computer vision hold great promise for visual collections in libraries, archives, and museums; but the red flags need not be overlooked to ensure this.

## Bibliography

Buolamwini, Joy. "Gender Shades: Intersectional Phenotypic and Demographic Evaluation of Face Datasets and Gender Classifiers." PhD diss., Massachusetts Institute of Technology, 2017.

Buolamwini, Joy. "InCoding – In the Beginning Was the Coded Gaze." *MIT Media Lab* (blog) *Medium*, May 16, 2016. <https://medium.com/mit-media-lab/incoding-in-the-beginning-4e2a5c51a45d>.

Cetinic, Eva, Tomislav Lipic, and Sonja Grgic. "Learning the Principles of Art History with Convolutional Neural Networks." *Pattern Recognition Letters*, 129 (2020): 56-62. <https://doi.org/10.1016/j.patrec.2019.11.008>.

Ciecko, Brendan. "AI Sees What? The Good, the Bad, and the Ugly of Machine Vision for Museum Collections," *The Museum Review* 5, no. 1 (January 2020), [http://articles.themuseumreview.org/tmr\\_vol5no1\\_ciecko](http://articles.themuseumreview.org/tmr_vol5no1_ciecko).

Ciecko, Brendan. "Examining the Impact of Artificial Intelligence in Museums." *MW17: MW 2017*. Last modified February 1, 2017. <https://mw17.mwconf.org/paper/exploring-artificial-intelligence-in-museums/>.

Ciecko, Brendan. "6 Ways That Machine Vision Can Help Museums." *Cuseum* (blog). Last modified March 10, 2016. <https://cuseum.com/blog/6-ways-that-machine-vision-can-help-museums>.

Cordell, Ryan. "Machine Learning + Libraries A Report on the State of the Field." *Library of Congress*, (July 14, 2020): i-86. <https://labs.loc.gov/static/labs/work/reports/Cordell-LOC-ML-report.pdf?loclr=blogsig>.

Deines, Nathaniel. "Does It Snow in L.A.? What Computer Vision Saw in Ed Ruscha's Sunset Boulevard." *Getty Iris* (blog), *Getty Museum*, October 7, 2020. <http://blogs.getty.edu/iris/does-it-wq2wsnow-in-la/>.

Google. "Google Arts & Culture." Accessed December 2020. <https://artsandculture.google.com>.

Google Arts & Culture. "MoMA & Machine Learning." *Experiments with Google* (blog). March 2018. <https://experiments.withgoogle.com/moma>.

Greene, Mark, and Dennis Meissner. "More Product, Less Process: Revamping Traditional Archival Processing." *The American Archivist* 68, no. 2 (September 1, 2005): 208–263. <https://doi.org/10.17723/aarc.68.2.c741823776k65863>.

Greenfield, Adam. *Radical Technologies: The Design of Everyday Life*. London: Verso, 2017.

hooks, bell. "The Oppositional Gaze: Black Female Spectators." *The Feminism and Visual Culture Reader* (2003): 94-105.

Hubel, D. H., T.N. Wiesel. "Receptive Fields of Single Neurones in the Cat's Striate Cortex." *The Journal of Physiology* 148 (1959). doi: 10.1113/jphysiol.1959.sp006308.

IBM. "Computer Vision." Accessed December 2020.  
<https://www.ibm.com/topics/computer-vision>.

IBM Research Trusted AI. "AI Explainability 360." Accessed December 2020.  
<http://aix360.mybluemix.net/resources#overview>.

Microsoft. "Computer Vision." Accessed December 2020.  
<https://azure.microsoft.com/en-us/services/cognitive-services/computer-vision/#features>.

MoMA. "Identifying Art Through Machine Learning: A Project with Google Arts & Culture Lab." Accessed December 2020.  
<https://www.moma.org/calendar/exhibitions/history/identifying-art>.

Moriarty, Adam. "A Crisis of Capacity: How can Museums use Machine Learning, the Gig Economy and the Power of the Crowd to Tackle Our Backlogs." *MW19:MW 2019*. Last modified January 15, 2019. <https://mw19.mwconf.org/paper/a-crisis-of-capacity-how-can-museums-use-machine-learning-the-gig-economy-and-the-power-of-the-crowd-to-tackle-our-backlogs/>.

Ngo, Ton, and Winnie Tsang. "Classify Art Using TensorFlow." IBM (blog). October 6, 2017. <https://developer.ibm.com/patterns/classify-art-using-tensorflow-model/>.

Noble, Safiya. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press, 2018.

O'Neil, Cathy. *Weapons of math destruction: How Big Data Increases Inequality and Threatens Democracy*. Broadway Books, 2016.

"Övningsfartyget Gladan anlöper Visby hamn." DigitaltMuseum. Accessed December 2020. <https://digitaltmuseum.se/011014890070>.

Padilla, Thomas. "Responsible Operations: Data Science, Machine Learning, and AI in Libraries." *OCALC Research* (December 2019): 6-22.  
<https://doi.org/10.25333/xk7z-9g97>.

Roach, John. "What's That? Microsoft's Latest Breakthrough, Now in Azure AI, Describes Images as Well as People Do." *The AI Blog, Microsoft*, October 14, 2020. <https://blogs.microsoft.com/ai/azure-image-captioning/>.

Robinson, Sarah. "When Art Meets Big Data: Analyzing 200,000 Items from The Met Collection in BigQuery." *Google Cloud* (blog). August 7, 2017.

<https://cloud.google.com/blog/products/gcp/when-art-meets-big-data-analyzing-200000-items-from-the-met-collection-in-bigquery>.

Saleh, Babak, Kanako Abe, Ravneet Singh Arora, and Ahmed Elgammal. "Toward Automated Discovery of Artistic Influence." *Multimedia Tools and Applications* 75 (April 2016): 3565–3591. <https://doi.org/10.1007/s11042-014-2193-x>.

SAS. "Machine Learning: What It Is and Why It Matters." Accessed December 2020. [https://www.sas.com/en\\_us/insights/analytics/machine-learning.html](https://www.sas.com/en_us/insights/analytics/machine-learning.html).

Sood, Amit. "Every Piece of Art You've Ever Wanted to See--Up Close and Searchable." *Google Arts & Culture*. Uploaded June 29, 2016. Accessed December 2020. Video, 15:00. <https://www.youtube.com/watch?v=CjB6DQGalU0>.

Srinivasan, Ramesh. "Automation Is Likely to Eliminate Nearly Half Our Jobs in the Next 25 Years. Here's What to Do." *Los Angeles Times*, October 29, 2019. <https://www.latimes.com/opinion/story/2019-10-29/opinion-automation-is-likely-to-eliminate-40-of-jobs-in-the-next-25-years-heres-what-we-can-do-a>.

Surapaneni, Sudeepti, Sana Syed, and Logan Yoonhyuk Lee. "Exploring Themes and Bias in Art using Machine Learning Image Analysis." *2020 Systems and Information Engineering Design Symposium* (2020). <https://doi.org/10.1109/SIEDS49339.2020.9106656>.

Viola, P., M. Jones. "Rapid Object Detection Using a Boosted Cascade of Simple Features," *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2001): 511-518. <https://doi.org/10.1109/CVPR.2001.990517>.

West, Sarah Myers, Meredith Whittaker, and Kate Crawford. "Discriminating Systems: Gender, Race and Power in AI." *AI Now Institute* (April 2019). <https://ainowinstitute.org/discriminatingsystems.html>.

Zhu, Yucheng, Yanrong Ji, Yueying Zhang, Linxin Xu, Aven Le Zhou, and Ellick Chan. "Machine: The New Art Connoisseur." *Cornell University arXiv preprint* (2019). <https://arxiv.org/pdf/1911.10091.pdf>.