To See or Not to See: Understanding the Tensions of Algorithmic Curation for Visual Arts

Ramya Srinivasan
Fujitsu Research of America
USA

ABSTRACT

Algorithmic recommendation is one of the most popular applications of machine learning (ML) systems. While the implication of algorithmic recommendation has been studied in the context of high-stakes domains such as finance and healthcare, there has been very little focus in understanding its impacts with respect to the arts domain. Given that ML is increasingly finding place in the arts domain such as in generative arts and content analysis, in this paper, we examine the tensions of algorithmic curation in the context of visual arts. Through case studies, we describe how curatorial algorithms that are oblivious of broader socio-cultural contexts could potentially result in ethical concerns such as over-representation and misattribution, to name a few. Towards addressing some of these concerns, the paper offers design guidelines. Specifically, the paper outlines repair strategies that suggest ways 1) to engage with cultural stakeholders in building visual art curatorial algorithms, 2) to unlearn biases embedded in digital artworks and their meta-data, and 3) emphasize the need to establish regulatory norms specific to the use of ML in visual art curation. Taking cue from the process employed by artwork curators, the paper also describes how authenticity can be prioritized by re-calibrating visual art curatorial algorithms. The paper also suggest ways through which the potential of state-of-the-art ML curatorial algorithms can be re-imagined towards empowering the audience of artworks. We hope the insights presented in the paper spark interdisciplinary discussions and pave way for fostering reformation in algorithmic curation of visual arts.

CCS CONCEPTS

• **Computing methodologies → Machine learning algorithms**; • **Human-centered computing → Human computer interaction (HCI)**; • **Social and professional topics**;

KEYWORDS

algorithmic recommendation, visual arts, case studies, machine learning, curation

ACM Reference Format:
Ramya Srinivasan . 2024. To See or Not to See: Understanding the Tensions of Algorithmic Curation for Visual Arts. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24), June 03–06, 2024, Rio de Janeiro, Brazil.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3630106.3658917

FAccT '24, June 03–06, 2024, Rio de Janeiro, Brazil

ACM ISBN 979-8-4007-0450-5/24/06
https://doi.org/10.1145/3630106.3658917

1 INTRODUCTION

From social media feeds to personalized search results, our daily digital life is replete with algorithmically-recommended content, shaping individual experiences and perceptions [115]. The massive scale at which online data is being produced coupled with technological advancements, in particular breakthroughs in machine learning (ML), has made algorithmic recommendation a very attractive venture. These developments have in turn triggered debates concerning the counter-benefits of algorithmic recommendation: the biases and prejudices recommendation algorithms can embed [22], the historical discrimination algorithms can amplify [9], and the power centralization that algorithms can reinforce [51].

Past studies have investigated issues of concern in algorithmic recommendation in high-stake domains such as healthcare [3], finance, and law [62]. There are also works that have examined social media/news feeds to uncover issues related to algorithmic biases [74] and exacerbation of power imbalance due to algorithmic interventions [99]. In parallel, there have been efforts to address these biases by designing bias-mitigating algorithms [44, 68] or through human-centered approaches [26, 79, 81, 111]. While there are a few studies that examine new age media such as music [1, 10, 73, 95, 102], there is not as much focus on visual art forms.

Visual arts deserves attention for various reasons. Firstly, visual arts encompasses a broad and wide variety of both traditional and modern art forms from ancient paintings to modern photography and video-making. From pre-historical times, visual art has been a form of communication deeply imprinted in human nature, having also the power to facilitate cognitive development in individuals [100, 114]. Researchers even argue that visual arts can provide a glimpse into history and culture —from the prehistoric cave paintings to the modern times, visual arts is regarded as a powerful means of storytelling [108], helping in documenting important social, religious, cultural, political, and other historic events of the past [27, 63].

More importantly, visual arts has been explored widely within the ML community—ML technologies are making inroads into a variety of visual arts applications such as for generating art, content analysis, artist identification, and style modeling [52, 89, 106, 117]. Amidst these advancements, many concerns have surfaced. There are philosophical and epistemological concerns–questions have been raised with regards to creative potential of algorithms [118]. Then, there are a variety of social, cultural, and economic concerns pertaining to the impact of AI art on traditional artists [45, 119]. As authors in [45] note, from reputational damage and economic loss to plagiarism and copyright infringement, traditional artists have had to experience a variety of harms due to the proliferation of generative arts. Thus, understanding the ethical and normative implications of algorithmic curation in the context of visual arts becomes necessary. For these reasons, in this paper we discuss issues related to algorithmic curation of visual arts.

It is to be noted that the term 'curation' is often used in the context of selection, organization, and presentation of artworks, especially in galleries, museums, art exhibitions, and other related venues. Curation often involves expert humans who oversee the selection and presentation process. On the other hand, recommendation is a more broad term that can refer to a wider set of applications (beyond arts) where lay users and algorithms can take the place of expert humans in selecting, organizing, and presenting content. In the context of this paper, we are focusing more on the curation and less on recommendation, specifically to refer to the process of selection, organization, and presentation of artworks across all platforms where artworks can be showcased. Thus, these could include physical platforms such as museums and physical galleries as well as online art galleries. In the rest of the paper algorithmic curation refers to ML based curation.

Motivated by the recent applications of ML algorithms in the visual arts domain, the paper discusses two important but non-exhaustive issues of conflict that arise as a consequence of algorithmic curation. First, we discuss the tensions surrounding the generation and digital reproduction of artworks. With the emergence of various generative AI tools and interfaces, artworks are being created and reproduced for a variety of tasks such as for simulating the styles of ancient artists and for rendering new artforms. Although generative AI might have enhanced accessibility to art, it has also paved the way for adverse downstream effects. Through two visual art case studies concerning artist style simulation, we illustrate how the nuanced elements of artists' style and the traditions of art movememts can be overlooked by generative algorithms. The second aspect of tension that the paper discusses concerns how algorithmic curation which was seen as a means to include global artistic narratives can potentially lead to exclusion of certain context specific narratives– two visual art case studies– one concerning curation of artworks aided by ChatGPT and another related to similar artwork retrieval based on culture are described to illustrate how algorithmic curation can lead to mis-attribution and other ethical concerns.

The paper outlines three potential pathways towards fostering reformations in algorithmic curation—through 'repair', 'recalibration', and 're-imagination'. Repair outlines ways to 1) engage with cultural stakeholders in building visual art recommendation algorithms, 2) unlearn biases embedded in digital artworks and their meta-data, and 3) emphasize the need to establish regulatory norms specific to the use of ML in visual art curation. We discuss how some of these measures can be implemented in practice. Recalibration calls for deliberation and a temporal re-orientation by shifting the focus towards a slow and enduring algorithmic curation whereby aspects such as authenticity are prioritized over speed and efficiency. Finally, we discuss ways through which the potential of state-of-the-art ML curatorial algorithms can be re-imagined towards empowering the audience of artworks.

Specifically, the paper aims to broaden the understanding of the implications of algorithmic curation with respect to visual arts by offering :

- A snapshot of the tensions associated with algorithmic curation for visual arts, taking into account the implications induced by recent ML advancements
- Accompanying case studies, illustrating some ethical concerns that can surface
- Potential pathways towards reforming the current state of algorithmic curation of visual arts–how can existing curatorial algorithms and platforms be *repaired*, why they need to be *recalibrated*, and how the the potential of state-of-the-art ML curatorial algorithms can be *re-imagined* towards empowering the audience of artworks. Design guidelines are discussed, a summary of which is provided in Table 1

2 BACKGROUND

Machine learning (ML) technologies are being used for a variety of visual arts related applications such as for content analysis [89, 106], artist classification [61], style transfer[117], etc. A comprehensive review of ML based art applications can be found in [18]. Furthermore with the emergence of generative algorithms, ML technologies are also being used for supporting existing artistic practices [23, 49] and for establishing novel methods of creating and re-mixing media [58]. In this section, we briefly situate our paper in the specific context of ML based curation of artworks and their associated concerns. The paper complements and extends previous works that have investigated biases in recommendation algorithms— by discussing issues specific to visual arts and by suggesting potential pathways towards addressing the concerns outlined.

2.1 Algorithmic curation of artworks: Design and Interaction Studies

From virtual and augmented reality systems to recommendation engines, researchers have extensively investigated on how to design interactive creative systems. From ethnographic fieldwork to people-powered collaborative system design, these studies have been informed by computational principles from various fields such as HCI, AI, humanities, and the social sciences.

For example, HCI researchers have offered ethnographic fieldwork based insights for the design studio culture [105]. In [48], a theoretical and methodological framework of art-based modes of inquiry in HCI and the broader STEM fields is offered. Studying the notion of embodied interactions in museums, the authors in [98] note that mode type impacts the number of visitors that interact with the installation, the gestures that people do, and the amount of time that visitors spend observing the data on display and interacting with the system. In [69], the potential of exhibits to engage new audiences in collaborative scientific discussions as part of people-powered research in science museums is studied. The authors in [87] discuss how visitors to museums and cultural heritage sites use and link digital information with physical information to shape others' understandings of cultural heritage. In [75] implications for how design can support learners' expression across culturally relevant themes are identified. An investigation of cross-cultural differences for website design is provided in [6]. To deepen people's experiences with public art, [33] provides designers with the grounds for informed inspiration in ideating such systems. The authors in [54] examine the role of live streams in the preservation of intangible digital heritage.

The richness of artworks and other such cultural heritage poses both opportunities and challenges in the development of tools for its curation and exploration [5]. Curation methods have to not only enhance information discovery and analysis, but also aid the audience by presenting authentic and interesting content. Thus, designing and delivering accessible and enjoyable experiences of cultural artefacts is an important research and development area that calls for a collaborative effort across disciplines.

2.2 Algorithmic curation of artworks: Applications

With the abundance of digital artworks, conventional modes of acquisition patterns are being disrupted. Art institutions are increasingly resorting to ML algorithms for curation [12, 39]. Recent years have witnessed several applications wherein ML technologies have been employed for curation of visual arts. One prominent application concerns personalized visual art recommendation based on user preferences. In this regard, the authors in [113] propose a multimodal visual art recommendation pipeline to customize the results based on user preferences. In another work [112], the authors propose a visual art recommendation system based on textual and visual features, and perform user-centric evaluations for assessing the quality of the recommendations.

Recommendation algorithms for visual arts have also been used to aid in understanding history. For example, a deep learning approach to cluster visual arts is proposed in [17] for historical knowledge discovery. In [37], the authors have designed an interactive web app that allows users to find pairs of semantically related artworks that span different cultures, media, and millennia. In [15], the authors propose a visual art recommendation system that they claim to be generalizable to any recommendation application. With rapid ML advancements, the number of applications related to visual art recommendation is likely to grow.

2.3 Algorithmic curation of artworks: Concerns

Despite the aforementioned progress, critics state that recommendation systems seem to flatten culture into numbers, that they normalize ever-broadening data collection, and profile their users for commercial ends [85, 86]. It is argued that algorithms are nothing more than objects of marketing and that this produces distinctive representations of algorithms that have evolved from technical and educational imagery [84].

Casting modernity as an ongoing contest between visuality and countervisuality, or 'the right to look', Nicholas Mirzoeff in his book *The Right to Look: A Counterhistory of Visuality* argues that the 'right to look' is the claim to a subjectivity that has the autonomy to arrange the relations of the visible and the sayable [64]. It is ironical that although artists are beginning to deconstruct algorithms to visualize and counter-visualize their underpinnings, how algorithms curate content, what they promote, and what gets seen by the viewers largely remains opaque. Researchers argue that AI should not only be understood not only as a tool for artists but also as cultural and political design material [13]. The author in [65] argues that an emerging layer of companies – referred to as the 'infomediaries' – are increasingly responsible for shaping how audiences experience cultural content. The author adds that the cultural content towards which infomediaries point users has less to do with user preferences and more to do with a supposed fit with quality that companies care about.

Recent research has exposed the biases associated with curatorial algorithms, prompting for greater diligence and emphasis on bias-mitigating strategies. However, researchers state that even if there was greater transparency about how algorithmic curation works [7], the dynamic and highly fluid nature of the algorithms which are constantly tuned based on the data collected from users implies that some amount of opaqueness always exists, that there is something 'impenetrable' about their performance [30]. The rapid pervasiveness of algorithms has given impetus to anxieties about the entanglement between the social, cultural, and the algorithmic –artworks have become intriguing in many ways– *"they are invisible yet omnipresent, proprietary yet pervasive, and with assumed socio-political powers that co-produce our lives"* [24].

While recent studies have investigated the implications of algorithmic curation on social media [29, 79, 103] and online music platforms [1, 10, 73, 95, 102, 103], there is not much focus on examining the impacts of algorithmic curation of visual arts. Works such as [33, 54, 69, 75, 105] that investigate the interplay between technology, culture, and creative practices offer valuable insights for designing people powered ML systems that can facilitate collaborative scientific advancements. This paper complements and extends previous works by focusing on the specific use case of algorithmic curation of visual arts. Through case studies, we illustrate the contradicting objectives of algorithmic curation for visual arts, while also suggesting potential pathways towards addressing some of the issues of concern.

3 THE TENSIONS OF ALGORITHMIC CURATION

Although a variety of tensions exist in the context of algorithmic curation of visual arts, here, we examine two points of friction. Through visual art case studies, we describe how generative algorithms used for curating artworks in the styles of ancient artists can be misleading, and how algorithmic curation can decontextualize artistic narratives by not considering the broader social-cultural contexts surrounding the artworks.

3.1 Digital Reproduction vs Irreproducibility

Starting from the times of mechanical reproduction and photography to the more recent ML based reproduction and generation, the discussion around pros and cons of digital reproduction of arts has had a long history. On the one hand, digital reproduction of artworks enables institutions to expand their presence beyond physical boundaries of museums and galleries while also allowing the audience to search for artworks from a large repository of worldwide collections. On the other hand, digitization and generation can pave way for many adverse downstream effects [39].

In the seminal essay 'The Work of Art in the Age of Mechanical Reproduction', Walter Benjamin writes *"...Unmistakably, reproduction as offered by picture magazines and newsreels differs from the image seen by the unarmed eye. Uniqueness and permanence are as* *closely linked in the latter as are transitoriness and reproducibility in the former"* [8]. Although Benjamin's comments are in the context of mechanical reproductions as relevant to the 1930s when the essay was written, the argument concerning the unique, indefinable, and irreproducible aura associated with the visual artworks applies to the current times of digital reproductions [19]. With ML generated art, there are additional layers of complexities [21, 66]. Generative art can be embedded with a variety of biases, including racial bias, gender bias, and cultural appropriation which can stem from the datasets used in training these models [91].

Furthermore, even if unethical datasets are deprecated, models and datasets derived from the source dataset can continue to raise ethical issues [56]. In a recent art exhibition, artist Trevor Paglen and researcher Kate Crawford uncovered the biases associated with the 'person' category of ImageNet, a foundational image dataset [20]. The creators of the dataset subsequently removed the offensive images, but there were already plenty of generative models and datasets based on ImageNet by then. For example, audit on the text to image generation CLIP model (trained on ImageNet) has revealed a variety of gender and racial biases [109]. CLIP and several other generative models have been extensively used to (re)create famous artworks of the past, albeit embedded with stereotypes and biases [91]. Below we describe two visual art case studies involving generative algorithms and describe some ethical concerns that arise in the context of curation.

3.1.1 Case study 1. A generative ML model called 'CycleGAN' [117] claims to model the style of famous artists like Monet and van Gogh. However, in training the model, the influence of art movements (i.e., tendencies or styles in art with a specific common philosophy influenced by various factors such as cultures, geographies, political-dynastical markers, etc. [107]) is not considered. It is to be noted that Monet was primarily influenced by Impressionism, a modern art movement while van Gogh's renditions were influenced by the Post-Impressionism art movement. When images across art movements are combined as was the case in the CycleGAN model, the fact that art movement is a potential confounder is ignored, thereby leading to biased representations. This is an illustration of the Simpson's paradox [76]. It was shown that spontaneous and accurate depiction of light along with its changing quality, an important characteristic of Impressionism was missing in the generated Monet versions, and expressive brushstrokes emphasizing geometric forms were missing in generated images corresponding to Post-Impressionist style of van Gogh [92]. Thus, not only the true cognitive abilities of an artist are undermined, but also the larger cultural practices as reflected in the art movements are mis-represented in the process. Despite this fundamental flaw in modeling, CycleGAN has been widely adapted for various applications, including in a recently exhibited application called 'Electronic Curator' [36] wherein the CycleGAN based algorithm not just generates a new image, but also takes the role of a curator evaluating the quality of the art piece.

3.1.2 Case study 2. As another instance, consider 'The Next Rembrandt Project' [80] wherein paintings of artist Rembrandt (owned by the Rembrandthuis museums) were processed to generate paintings in the artist's style. In training the algorithm, the original images were divided into smaller patches to make it amenable for processing, a step which can greatly degrade the original image resolution and color calibration [78]. Experts argue that the generated images do not reproduce Rembrandt's style, and could contributes towards misattribution [16]. Despite such issues, a similar algorithm was used to 'reproduce' the lost fragments of Rembrandt's 1642 artwork 'Night Watch' (Please see Figure 1). The missing part was restored using a smaller seventeenth-century copy made by artist Gerrit Ludens (Please see Figure 2), on which the ML algorithm applied 'Rembrandt's style' to complete the missing left portion of Figure 1 [78]. The algorithm worked on small patches of both paintings' reproductions, comparing the two styles. Despite the poor quality of image used in (mis)creating the lost fragment, the 'reproduced' painting was featured at the Rijksmuseum museum. From stereotyping artists to misattribution, these case studies illustrate some of the concerns induced by the use of new age digital reproductions/generations in curatorial applications.

3.2 Inclusion vs Exclusion

In the interpretation of visual arts, museum work often involves the establishment of a canon through which order is created 'by giving authority to certain texts, figures, ideas, problems, discursive strategies and historical narratives' [41]. Researchers argue that algorithms lack the intellectual predisposition of artistic intention and in their attempt to form a narrative of contemporary humanism, may overwhelm and choke the development of historical schools of thought [77].

Artworks are embedded in social, cultural, environmental, political, and religious contexts—the artist behind the artwork, the style or art movement that influenced the artwork, the materials that were used in creating the artwork, the commissioning authorities, the purpose behind the artwork characterizes the value and relevance of artworks. Algorithms seldom consider these aspects in curation. Algorithms often determine value by means of indexing and connections [46]. It has been argued that algorithmic curation can tag significant portion of contemporary events in the visual arts as 'historical' by the diffusion of information that is compatible to algorithmic computation [77]. For instance, the 'Art Genome Project' recommends other artworks to users based on personal taste and previous choices, an approach that has been criticized for being over-simplistic and devoid of contexts [39]. In the process, audience may be made subjects to an invisible history as perceived by the algorithm, a history that is decontextualized and misrepresented [57].

3.2.1 Case study 3. Recently, students and faculty in Duke university's Art, Art History and Visual Studies department along with Duke Digital Art History and Visual Culture Research Lab set out to understand the usefulness and limitations of AI as a curator in museums. In their experiment, they gathered a dataset of nearly 14,000 objects in the Nasher Museum's collection by developing a tool to extract publicly accessible information from the museum's databases [71]. This was further transformed into machine-readable data that is understandable by OpenAI's ChatGPT platform. Specifically, the research developed a series of prompts and instructions for ChatGPT that asked it to select artworks for the exhibition. While there were many aspects to admire in the collection displayed by the

Figure 1: 'Night Watch', a 1642 artwork of Dutch artist Rembrandt. It is believed that some portions of the original artwork are missing in the available versions. Image source https://www.wikiart.org [107].

algorithm, there were also notable flaws [31]. For example, the algorithm called one painting a sculpture and mis-titled another, raising questions related to mis-attribution. It was also observed that the algorithm picked a non-descriptive bowl to be displayed alongside a Dali [31], potentially leading to simplification of a complex work of art. Additionally, as the author in [31] writes, the language used in the descriptions was somewhat rigid and repetitive.

3.2.2 Case study 4. As another instance, consider 'MosAIc', an interactive web app that allows users to find pairs of semantically related artworks that span different cultures, media, and millennia [37]. Specifically, visual similarity search is conducted using user supplied filters or 'conditions' to find pairs of similar images that span distinct subsets of the image corpus. Examples of filters or conditions that users can provide include 'culture' and 'medium'. However as the authors themselves acknowledge, the method chooses some items significantly more than others, this is known as the 'hubness problem' [37]. This means certain artworks are shown to be more similar to the query primarily because the algorithm sees it in the nearest neighborhood of the query often. Although the artworks may be distinct in many ways, the algorithm fails to recognize the nuances of artistic traditions, potentially contributing to misattribution and cultural appropriation. Thus, while the original intention of leveraging algorithms was to provide a diverse and authentic representation of world culture and history, such examples raise questions concerning the authenticity and diversity associated with algorithmic curation.

4 DISCUSSION

In what follows, we outline three non-exhaustive potential pathways towards fostering reformations in algorithmic curation of

Figure 2: A copy of 'Night Watch' by artist Gerrit Ludens which includes the missing left portion of the original artwork. Figure source: https://www.nationalgallery.org.uk/paintings/gerrit-lundens-after-rembrandt-the-company-of-captain-banning-cocq-the-nightwatch listed under the license https://creativecommons.org/licenses/by-nc-nd/4.0/

visual arts—through what we refer to as 'repair', 'recalibration', and 're-imagination'. In section 4.1 under repair, we discuss what could to be done to fix existential issues of concern in algorithmic curation. Section 4.2 (recalibration) discusses the need for a temporal reorientation facilitating deliberation instead of rapid scaling of algorithmic curation. Section 4.3 discusses various ways through which the potential of ML algorithms could be leveraged towards empowering the audience of visual arts. Table 1 provides a summary of the design recommendations.

4.1 Repair: Unlearning biases, facilitating regulation, and restoring dialogue

The first repair pathway concerns unlearning the biases embedded in ML models and training datasets. For decades now, ML systems interpret tasks (be it recommendation, classification, or anything else) as something that is objective and definable. For example, algorithmic curation largely perceives artworks as static; in contrast, artworks are dynamic and continuous bodies of knowledge shaped by a variety of social, cultural, religious, and political contexts. Thus, there is a need to to recognize the co-constructed elements that govern artworks as opposed to the existing prejudiced notions of discreteness, immobility, and objectivity that are wrongly associated with artworks [11, 53]. Such a repair process may embed modes of human interaction with technology and with each other in ways that surface values as contingent and ongoing accomplishments, as suggested in [42].

Specifically, techniques that strategically limit the influence of a data point in the training procedure can been seen as a viable option to unlearn existing biases [55]. Formally called 'machine unlearning', these are techniques to make ML models 'forget' certain training data instances (e.g., those that embed biases or infringe

Reformative measure	**Potential Pathways**
Repair:unlearning biases	'machine unlearning' techniques to limit the influence of biased data instances; active learning methods to interactively update identity labels of artworks
Repair: regulation	standardization measures for artwork content quality and moderation; establishment of best practices for art datasets documentation
Repair: restore dialogue	leveraging expert interaction data for evaluation of curation results; employing jury learning for resolving artwork metadata disagreements
Recalibrate	temporal reorientation that prioritizes authenticity and fairness over speed; multi-level multi-interest recommendation algorithms for fine grained analysis of individual and collective interactions embedded in artworks
Reimagine	conscious data contribution, data protection measures, cultural resistance, leveraging visual art to raise awareness, anti-consumption activism

Table 1: Summary of potential pathways towards reforming the current state of algorithmic curation of visual arts

artists' copyright) so that their adverse impact on the outcome can be minimized [35, 96]. Additionally, active learning techniques wherein experts guide the ML model by labeling the most uncertain and biased data instances can be helpful in interactively updating identity labels of artworks, thereby facilitating co-construction [34].

The second aspect of repair concerns facilitating regulation. Most of the proposed ML regulatory frameworks relate to high stake domains such as finance and healthcare. Barring recent works such as [90] wherein the authors provided a set of checklists for responsible development of art datasets and [45] wherein the authors provide recommendations for protection of artists' rights, there is not much work related to regulatory principles for the arts domains. As noted in [2], existing curatorial platforms are not bound by any filtering quality or neutrality requirements. Lack of appropriate metrics to quantify content quality and neutrality could be a potential hindrance in determining content quality. Thus establishing standardization and documentation best practices for algorithmic curation of visual art becomes essential to enable fairness, reliability, and trust. Such regulatory measures need to be drafted in consultation with legal and cultural experts and could lay the foundation in enforcing specific requirements with regards to art work provenance, copyright, authenticity, use, and distribution.

The third repair pathway concerns restoring dialogue. Algorithmic curation platforms largely focus on content producers, curators, and consumers. However, consulting a wider set of stakeholders (e.g., art historians, legal experts, cultural policy makers, artists, archivists, etc.) will not only usher in diverse perspectives but also facilitate dialogue [21]. There is a pressing need to restore such dialogue, especially given the amplification of historical and societal injustice, misattribution, and revisionism brought about by algorithmic curation.

Even beyond the arts domain, the need for going back to the concerned people and sites under study is something that HCI researchers have strongly emphasized [28]. Dialogues can enable a restorative effect by serving as a means of deescalating ideological confrontations between different cultural groups [2]. Thus, restoring dialogue between all the concerned stakeholders and using these informed opinions in ML algorithm design/evaluation can be a promising repair pathway.

One way of restoring dialogue is by asking art experts to interact with algorithmic curation results, and subsequently leveraging the interaction data to repair the algorithm. For example, interactions from meeting observations, interviews, documentation, and online interaction data were leveraged to show how non-technical art experts can explain and repair sociotechnical breakdowns – when their expectations for similarity between art images and artists differ from the similarity relations produced by the algorithm [25, 82]. Another way of restoring dialogue could be to consult domain experts for verifying the annotations provided by lay workers. For example, jury learning techniques could be leveraged in resolving annotation disagreements in a cost effective manner by determining when expert advice has to be sought [32].

4.2 Recalibrate: Encouraging deliberation and temporal reorientation

In the seminal work *'Thinking Fast and Slow'* [47], Daniel Kahneman remarks that good decision-making demands a tradeoff between the efficiency of fast paced intuitive thinking and the accuracy of slow and deliberative thinking. A significant portion of technological acceleration can be attributed to speed–the fact that algorithms save time is one of the main factors for their appeal and adoption. However, as Kahneman and several other scholars argue, using computational tools to speed up (or slow down) certain decisions is not a 'neutral' adjustment–how quickly or slowly algorithms make decisions is often intertwined with latent objectives that are in turn governed by individual, organizational, social, or political motives [38, 93].

The following example serves as an illustration to the aforementioned point. At a public exhibition organized in London in January 2023, Oxford researchers launched a project called 'The Algorithmic Pedestal' that highlighted the differences between algorithmic and human curation. Specifically, drawing on the Metropolitan Museum of Art's open access collection, both a human artist and an algorithm used in Instagram were invited to select images for this exhibit. Sorting through thousands of images, twenty to thirty images from each (human artist and the Instagram algorithm) were chosen to be displayed in a particular layout and order [40]. The human artist involved, namely, Fabienne Hess, chose to select images that corresponded to the concept of 'loss'. According to Hess, loss is a uniquely—and universally—human experience. The images that were displayed were part of her collection called "Dataset of Loss", which she created over the course of three years as a resistance to the dominant algorithmic ways of seeing.

As mentioned in [40], Hess's curatorial process was driven by the human experiences of time, curiosity, and patience as she spent years physically exploring collections in an embodied fashion, learning about each object's stories and photographing them during site visits [40]. On the other hand, the rationale behind Instagram algorithm's selection was largely opaque due to its blackbox nature. It is argued that many of the algorithms used by social media platforms like Instagram are not motivated by the goal of artistic display, instead they are optimized to maximize social connections and selling ads [72]. Curatorial platforms that jump to judgment before considering relevant contexts can decontextualize art history leading to many adverse downstream effects. The speed at which decisions are made can privilege certain values over others (e.g., monetary benefit over artistic appeal).

Instances such as the above reiterate the need for deliberation in the entire process [70]. Although deliberation implies more time, making algorithmic curation less attractive and efficient, it can decrease the probability of misattribution, over-representation, and other systemic biases. There is thus a need to recalibrate curatorial algorithms to weigh in factors such as authenticity and reliability—choosing to use existing systems without attending to questions about their temporalities implicitly prioritizes speed and efficiency over decision-making accuracy [93].

It is important to note that cultural phenomena are inherently multi-scale spanning time and space, entailing varying levels of interactions. In the illustration described above, artist Hess's curation was a slow and detailed process. It was not merely labeling what the images depicted, but instead a process that asked questions about origin and content while resisting being categorized and labelled [40]. Thus, deliberation would require curatorial algorithms to be able to scale down to the micro-levels of individual artistic renditions and up to the macro-levels of collective societal interactions that shape the artworks. Towards understanding and capturing such multi-level interactions in curation, a class of techniques called 'multi-interest recommendation algorithms' that aim to learn multiple aspects using the interaction data between users and items could serve as a starting point [116]. These algorithms could be leveraged to learn distinct aspects of artworks while optimizing for correlation between user interest and artwork authenticity. Additionally, hierarchical ML methods such as graph neural networks and agent based modeling methods that take into account scaling and universal dynamics in both social (e.g., in learning the influence of artists' lineage and art movements) and physical properties (e.g., in understanding the art material) could be used in conjunction with multi-interest recommendation algorithms for enabling fine grained analysis in visual art curatorial algorithms [14, 50, 97].

4.3 Re-imagine: Exploring opportunities to resist status quo and enabling reconceptualization

Audience and users of curatorial platforms can help in enabling reformation in the state of curatorial algorithms. By audience and users, we are referring to a) human curators involved in galleries and museums b) art connoisseurs and private art collectors as well as c) casual viewers.

In 2015, Swedish design student Johanna Burai began a project to change Google Image search results. After being annoyed with the racial bias in the search results – while, for example, searching for depictions of hands – she decided to test and correct the algorithm [101]. Although her idea might not have been revolutionary, the audit experiment elucidated the possibilities of users to resist algorithmic power [104].

It becomes important to acknowledge the mutual co-construction of algorithms and their users, and re-imagine alternative uses of algorithms, including how algorithms can help in auditing some of its own products. For example, the authors in [101] reframe the public and scholarly debates on algorithmic power, drawing on media as practice to advance a framework for studying algorithms with a focus on user agency. In the work 'Manifesto for the Broken Machine' [88], Prof. Sarah Sharma investigates how feminism can help in reframing the technological as a mode of resistance, rather than a form of control. Prof. Sharma writes that technologies such as ML often neglect systemic forms of social injustice and conveniently discard and replace "malfunctioning parts (nonconforming subjects)". She reckons *"But the idea of our contemporary social-political-economic system as an already-broken machine full of the incompatibly queer, raced, classed, and sexed broken-down machines is politically exciting for feminism"*. Taking cue from Prof. Sharma's concept of brokenness, the artwork 'This Recommendation System is Broken', provides a critical reflection on discriminatory practices induced by most common algorithmic recommendation systems, uncovering epistemologies of not knowing: incomplete, inaccurate, and unidentified digital accounts [94].

In the art exhibition 'Recontres Internationales Paris-Berlin', the piece 'What the Robot Saw', which is a live, continuously-generated robot film curated, analyzed, and edited using computer vision, neural networks, and contrarian search algorithms, highlighted some of the least viewed videos on YouTube, featuring first person narratives by some of the people that commercial ranking algorithms ignore [4]. What the Robots saw illustrates the complex relationship between the curatorial ML algorithms and the humans, depicting the awkward contemporary collision of performed selves and screen-centric perceptions [4]. Figure 3 is an image curated by What the Robot Saw, demonstrating gender stereotypes embedded in commercial recognition algorithms that are used for categorizing human faces.

Both the artworks 'This Recommendation System is Broken' and 'What the Robot Saw' are illustrations of how visual arts can be leveraged to uncover the drawbacks of algorithmic curation. Researchers argue that the field of visual arts organically possesses the ability to protest, challenge, inspire, and be instrumental in humanizing technological advances [59]. Guided by humans-in-the-loop, the potential of ML algorithms can be thus used to co-curate visual arts such that it helps in promoting cultural neutrality [43], facilitating situational interpretation and community participation [67].

Collective proactive people participation can play a pivotal role in re-conceptualizing algorithmic curation. As studies have pointed out, there is a need to shift the focus from people's wants and deficits towards a deep understanding of people's assets and capacities [110]. Leveraging Swidler's theory of culture-in-action, the

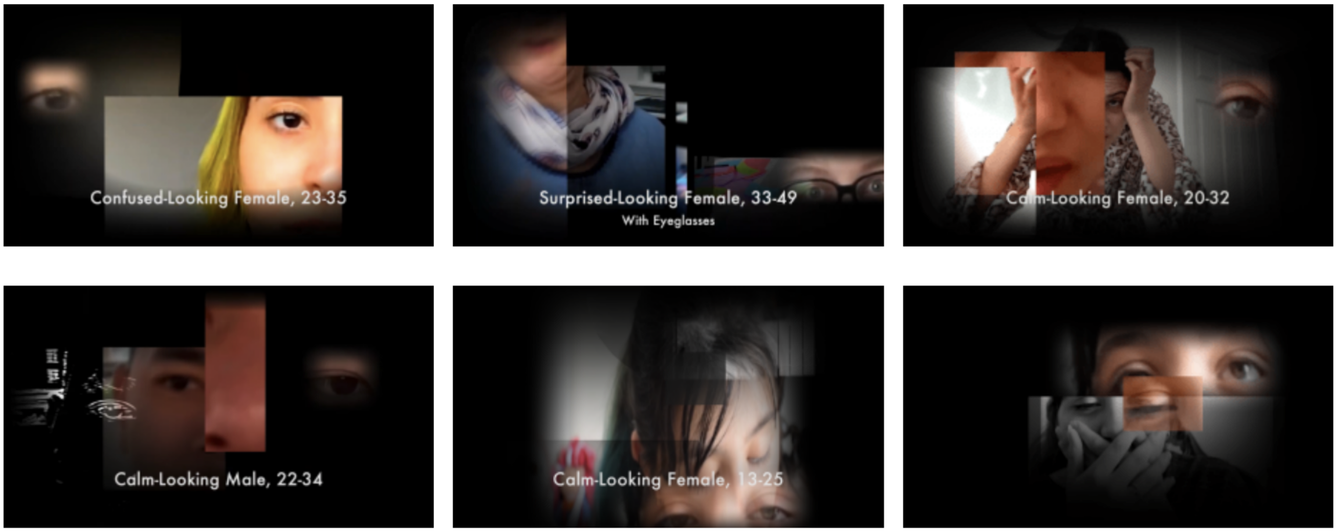

Figure 3: ©: Amy Alexander —An illustration from 'What the Robot Saw'. Personal videos from among those with low subscribers and view counts are showcased by What the Robot Saw project. The section headers are derived from Amazon Rekognition highlighting how commercial algorithms choose to categorize human faces. Figure source: https://amy-alexander.com/live-performance/what-the-robot-saw/ Artist ©: Amy Alexander. Permission obtained from the artist for including the image in the paper

authors in [110] propose an analytical lens for unpacking the complex relationship between people's capacities, goals, and structural limitations.

There could be many channels through which people can reclaim some of their power back from the algorithms that currently determine what they see. Human-machine co-curation platforms that provide ample control to users in the curation process can be beneficial in this regard. For instance, in [60], the authors proposed training machine learning models that allow curators to document past curatorial practices of the newspaper's photo librarians, retrace the editorial selection from photo-assignments and propose archival paths supported by several historical documents. Instead of relying on unique identifiers as seen in ontology-based curatorial practices, the proposed method promotes dialogue between archives by relying on metadata and image recognition.

Additionally, cultural resistance and anti-consumption activism can also reduce power centralization [83]. This could entail conscious data contribution whereby users are conscious of what artworks they endorse via likes, shares, subscriptions, etc., data protection measures that safeguard the interests of artists [45], and by displaying activism against forces that undermine cultural values. All these measures can be helpful in resisting status quo and help in re-imagining the purpose and principles governing algorithmic curation of visual arts.

5 LIMITATIONS

The list of concerns outlined or the strategies suggested in the paper are not exhaustive by themselves, and represents a subset of the broader set of concerns and reformatory pathways. The paper offered a set of potential pathways towards reforming algorithmic curation of visual arts without necessarily delving deep into the challenges involved in their implementation and deployment. For example, development of regulations for visual art content moderation requires establishment of international standards which in turn may necessitate policy intervention and consensus across government bodies and cultural organizations worldwide. Additionally, restoring dialogues with relevant stakeholders like artists would perhaps require the creation of inclusive forums through which all stakeholders can freely interact and exchange opinions/ideas without feeling intimidated or sidelined. In a future work, the goal is to study a broader set of ethical issues and interview artists and other stakeholders to incorporate first-person accounts in shaping the reformation of algorithmic curation of visual arts.

6 CONCLUSIONS

The field of visual arts has an important bearing on society through culture, history, philosophy, and economy. The growing use of ML technologies in visual arts such as for content analysis and curation makes it imperative to understand the ethical implications of such applications. In this paper, we examined the tensions of algorithmic curation of visual arts, taking into account recent ML advancements. Through visual art accompanying case studies, we highlighted how these tensions can lead to a variety of ethical concerns such as overrepresentation and misattribution. The paper suggested repair strategies to engage with cultural stakeholders in building visual art curatorial algorithms, to unlearn biases embedded in digital artworks and their meta-data, and emphasized the need to establish regulatory norms specific to the use of ML in visual art curation. The paper described the need for prioritizing authenticity through a recalibration of visual art curatorial algorithms, and suggested ways through which the potential of state-of-the-art ML curatorial algorithms can be re-imagined towards empowering the audience of artworks. We hope the insights presented can spark interdisciplinary discussions and shed light on promising future directions.

Ethical Considerations: The author has adhered to the institutional ethical guidelines in conducting the research. The author does not foresee any adverse consequences from this work.

Author Positionality: The author's background spans computer vision, machine learning, and applied ethics, with a special focus on understanding the ethical impacts of creative ML pipelines. The author has experience working with both traditional and generative visual artists. The author has also extensively interacted with the broader visual art community members such as art historians, visual art connoisseurs, and art journalists.

REFERENCES

[1] Ferraro A, Serra X, Bauer C. Scholer F, and Thomas P (eds.). 2021. Break the Loop: Gender Imbalance in Music Recommenders. *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval* (2021).

[2] Tarek Abdelzaher, Heng Ji, Jinyang Li, Chaoqi Yang, John Dellaverson, Lixia Zhang, Chao Xu, and Boleslaw K. Szymanski. 2020. The Paradox of Information Access: Growing Isolation in the Age of Sharing. *ArXiv* (2020).

[3] Deena Abul-Fottouh, Melodie YunjuSong, and Anatoliy Gruzd. 2020. Examining algorithmic biases in YouTube's recommendations of vaccine videos. *International Journal of Medical Informatics* (2020).

[4] Amy Alexander. 2022. What the Robot Saw. *Recontres Internationales Paris-Berlin* (2022).

[5] Liliana Ardissono, George E. Raptis, and Noemi Mauro. 2022. Special Issue on AI and HCI Methods and Techniques for Cultural Heritage Curation, Exploration and Fruition. *Applied Sciences, 12(19)* (2022).

[6] Amanda Baughan, Nigini Oliveira, Tal August, Naomi Yamashita, and Katharina Reinecke. 2021. Do Cross-Cultural Differences in Visual Attention Patterns Affect Search Efficiency on Websites? *CHI Conference on Human Factors in Computing Systems* (2021).

[7] D Beer. 2009. Power through the algorithm? Participatory web cultures and the technological unconscious. *New Media and Society 11: 985–1002* (2009).

[8] Walter Benjamin. 1969. The Work of Art in the Age of Mechanical Reproduction. *Illuminations, edited by Hannah Arendt, translated by Harry Zohn, from the 1935 essay New York: Schocken Books, 1969* (1969).

[9] Saumya Bhadani. 2021. Biases in Recommendation System. *ACM RecSys* (2021).

[10] Marc Bourreau and Germain Gaudin. 2021. Streaming platform and strategic recommendation bias. *Journal of Economics and Management Strategy* (2021).

[11] Judith Butler. 2011. Bodies that matter: On the discursive limits of sex. *Routledge* (2011).

[12] Susan Cairns and Danny Birchall. 2013. Curating the Digital World: Past Preconceptions, Present Problems, Possible Futures. *Annual conference of Museums and the Web* (2013).

[13] Baptiste Caramiaux and Sarah Fdili Alaoui. 2022. Explorers of Unknown Planets, Practices and Politics of Artificial Intelligence in Visual Arts. *Proceedings of the ACM on Human-Computer Interaction Volume 6, Issue CSCW2 Article No.: 477, pp 1–24* (2022).

[14] Simon Carrignon, Tom Brughmans, and Iza Romanowska. 2020. Tableware trade in the Roman East: Exploring cultural and economic transmission with agent-based modelling and approximate Bayesian computation. *PLOS ONE* (2020).

[15] Manuel Cartagena, Patricio Cerda, Pablo Messina, Felipe Del Río, and Denis Parra. 2020. CuratorNet: Visually-aware Recommendation of Art Image. *Proceedings of the Fourth Workshop on Recommendation in Complex Environments* (2020).

[16] Sarah Cascone. 2016. New Rembrandt Artwork Created via 168,263 'Painting Fragments'. *Artnet News: https://news.artnet.com/art-world/new-rembrandt-through-data-analysis-466517.* (2016).

[17] Giovanna Castellano and Gennaro Vessio. 2022. A Deep Learning Approach to Clustering Visual Arts. *International Journal of Computer Vision* (2022).

[18] Eva Cetinic and James She. 2022. Understanding and Creating Art with AI: Review and Outlook. *ACM Transactions on Multimedia Computing, Communications, and Applications, Volume 18, Issue 2* (2022).

[19] Noah Charney and Zocalo Public Square. 2016. Has the Incredible Accuracy of Art Reproduction Ruined the Way We Experience Masterpieces? *Smithsonian Magazine: Arts and Culture August Issue* (2016).

[20] Kate Crawford and Trevor Paglen. 2019. Excavating AI: The politics of training sets for machine learning. (2019).

[21] Antonio Daniele and Yi-Zhe Song. 2019. AI+Art= Human. *AAAI AI Ethics and Society* (2019).

[22] Abhisek Dash, Abhijnan Chakraborty, Saptarshi Ghosh, Animesh Mukherjee, and Krishna P. Gummadi. 2021. When the Umpire is also a Player: Bias in Private Label Product Recommendations on E-commerce Marketplaces. *ACM FAccT* (2021).

[23] Nicholas Davis. 2019. Creative Sketching Partner: A Co-Creative Sketching Tool to Inspire Design Creativity. *International Conference on Computational Creativity* (2019).

[24] Patricia de Vries. 2019. Algorithmic Anxiety in Contemporary Art: A Kierkegaardian Inquiry intothe Imaginary of Possibility. *Published by Institute of Network Cultures, Amsterdam* (2019).

[25] Vicente Dominguez, Ivania Donoso-Guzmán, Pablo Messina, and Denis Parra. 2020. Algorithmic and HCI Aspects for Explaining Recommendations of Artistic Images. *ACM Transactions on Interactive Intelligent Systems Volume 10 Issue 4 Article 30* (2020).

[26] Motahhare Eslami, Aimee Rickman, Kristen Vaccaro, Amirhossein Aleyasen, Andy Vuong Karrie Karahalios, Kevin Hamilton, and Christian Sandvig. 2015. I always assumed that I wasn't really that close to [her]: Reasoning about invisible algorithms in the news feed. *CHI Conference on Human Factors in Computing Systems* (2015).

[27] Jennifer Farrell. 2017. Art as Influence and Response: A First Look at World War I and the Visual Arts. *Met Museum Blogs-World War I and the Visual arts* (2017).

[28] Sarah Fox and Daniela K. Rosner. 2016. Continuing the Dialogue: Bringing Research Accounts Back into the Field. *CHI Conference on Human Factors in Computing Systems* (2016).

[29] Anna Gausen, Wayne Luk, and Ce Guo. 2022. Using Agent-Based Modelling to Evaluate the Impact of Algorithmic Curation on Social Media. *Journal of Data and Information Quality* (2022).

[30] T Gillespie. 2014. The relevance of algorithms. *Gillespie T, Boczkowski P and Foot K (eds) Media Technologies: Essays on Communication, Materiality, and Society. Cambridge, MA: MIT Press, pp.167–193* (2014).

[31] Brian Gordon. Retrived April 2024. Meet the new curator at teh Duke art musuem. It isn't human. *https://www.newsobserver.com/entertainment/arts-culture/article279310784.html* (Retrived April 2024).

[32] Mitchell L. Gordon, Michelle S. Lam, Joon Sung Park, Kayur Patel, Jeffrey T. Hancock, Tatsunori Hashimoto, and Michael S. Bernstein. 2022. Jury Learning: Integrating Dissenting Voices into Machine Learning Models. *ACM CHI* (2022).

[33] Tim Gorichanaz. 2020. Engaging with Public Art: An Exploration of the Design Space. *CHI Conference on Human Factors in Computing Systems* (2020).

[34] Denis Gudovskiy, Alec Hodgkinson, Takuya Yamaguchi, and Sotaro Tsukizawa. 2020. Deep Active Learning for Biased Datasets via Fisher Kernel Self-Supervision. *IEEE CVPR* (2020).

[35] Varun Gupta, Christopher Jung, Seth Neel, Aaron Roth, Saeed Sharifi-Malvajerdi, and Chris Waites. 2021. Adaptive Machine Unlearning. *NeurIPS* (2021).

[36] Eran Hadas and Eyal Gruss. 2022. The Electronic Curator or How to Ride Your CycleGAN. *https://iscma.scm.cityu.edu.hk* (2022).

[37] Mark Hamilton, Stephanie Fu, Mindren Lu, Johnny Bui, Darius Bopp, Zhenbang Chen, Felix Tran, Margaret Wang, Marina Rogers, Lei Zhang, Chris Hoder, and William T. Freeman. 2021. MosAIc: Finding Artistic Connections across Culture with Conditional Image Retrieval. *Proceedings of Machine Learning Research 1:1–23* (2021).

[38] Alex Hanna and Tina Park. 2020. Against Scale: Provocations and Resistances to Scale Thinking. *CSCW Workshop Reconsidering Scale and Scaling in CSCW Research* (2020).

[39] Manique Hendricks. 2018. The Algorithm as Curator: In Search of a Non-narrated Collection Presentation. *Stedelijk Studies Journal* (2018).

[40] Laura Herman. Retrieved August 2023. The Algorithmic Pedestal. *https://www.algorithmicpedestal.info/about* (Retrieved August 2023).

[41] E. Hooper-Greenhill. 2000. Museums and the interpretation of visual culture. *London and New York: Routledge* (2000).

[42] Lara Houston, Steven J. Jackson, Daniela K. Rosner, Syed Ishtiaque Ahmed, Meg Young, and Laewoo Kang. 2016. Values in Repair. *CHI Conference on Human Factors in Computing Systems* (2016).

[43] Han-Yin Huang and Cynthia C. S. Liem. 2022. Social Inclusion in Curated Contexts: Insights from Musuem Practices. *ACM FAccT* (2022).

[44] C. Ilvento, M. Jagadeesan, and S. Chawla. 2020. Multi-category Fairness in Sponsored Search Auctions. *ACM FAccT* (2020).

[45] Harry Jiang, Timnit Gebru, Jessica Cheng, Lauren Brown, Mehtab Khan, Johnathan Flowers, Abhishek Gupta, Deja Workman, and Alex Hanna. 2023. AI Art and its impact on artists. *AIES* (2023).

[46] David Joselit. 2013. After Art. *Princeton and Oxford: Princeton University Press* (2013).

[47] Daniel Kahneman. 2013. Thinking Fast and Slow. *(1st pbk. ed ed.). Farrar, Straus and Giroux, New York. OCLC: ocn834531418* (2013).

[48] Laewoo Kang, Steven Jackson, and Trevor Pinch. 2022. The Electronicists: Techno-aesthetic Encounters for Nonlinear and Art-based Inquiry in HCI. *CHI Conference on Human Factors in Computing Systems* (2022).

[49] Thomas Kerdreux, Louis Thiry, and Erwan Kerdreux. 2020. Interactive Neural Style Transfer with Artists. *International Conference on Computational Creativity* (2020).

[50] Edward D. Lee, Bryan C. Daniels, Christopher R. Myers, David C. Krakauer, and Jessica C. Flack. 2020. Scaling theory of armed-conflict avalanches. *Physical Review E* (2020).

[51] Liu Leqi, Dylan Hadfield-Menell, and Zachary C. Lipton. 2021. When Curation Becomes Creation: Algorithms, microcontent, and the vanishing distinction between platforms and creators. *Queue, Volume 19, Issue 3* (2021).

[52] Bingchen Liu, Kunpeng Song, Yizhe Zhu, and Ahmed Elgammal. 2020. Sketch-to-Art: Synthesizing Stylized Art Images from Sketches. *Asian Conference on Computer Vision* (2020).

[53] Christina Lu, Jackie Kay, and Kevin R. McKee. 2022. Subverting machines, fluctuating identities: Re-learning human categorization. *ACM FAccT* (2022).

[54] Zhicong Lu, Michelle Annett, Mingming Fan, and Daniel Wigdor. 2019. I feel it is my responsibility to stream: Streaming and Engaging with Intangible Cultural Heritage through Livestreaming. *CHI Conference on Human Factors in Computing Systems* (2019).

[55] Christopher A. Choquette-Choo Hengrui Jia Adelin Travers Baiwu Zhang David Lie Nicolas Papernot Lucas Bourtoule, Varun Chandrasekaran. 2021. Machine Unlearning. *IEEE S and P* (2021).

[56] Alexandra Sasha Luccioni, Frances Corry, Hamsini Sridharan, Mike Ananny, Jason Schultz, and Kate Crawford. 2022. A Framework for Deprecating Datasets: Standardizing Documentation, Identification, and Communication. *ACM FAccT* (2022).

[57] Mykola Makhortykh, Aleksandra Urman, and Roberto Ulloa. 2021. Hey, Google, is this what the Holocaust looked like? Auditing algorithmic curation of visual historical content on Web search engines. *First Monday, Volume 26, Number 10* (2021).

[58] Marian Mazzone and Ahmed Elgammal. 2019. Art, Creativity, and the Potential of Artificial Intelligence. *Arts* (2019).

[59] Nicholas Mcguigan and Alessandro Ghio. 2019. Art, accounting and technology: unravelling the paradoxical in-between. *Meditari Accountancy Research* (2019).

[60] C. Meghini, V. B. Lenzi, Daniele Metilli, and Filippo Benedetti. 2019. Introducing narratives in Europeana: A case study. *International Journal of Applied Mathematics and Computer Science* (2019).

[61] T. Mensink and J. Van Gemert. 2014. The rijksmuseum challenge: Museum-centered visual recognition. *Proceedings of International Conference on Multimedia Retrieval* (2014).

[62] Silvia Milano, Mariarosaria Taddeo, and Luciano Floridi. 2020. Recommender systems and their ethical challenges. *AI and Society* (2020).

[63] Joyce Miller. 2003. Using the Visual Arts in Religious Education: An Analysis and Critical Evaluation. *British Journal of Religious Education* (2003).

[64] Nicholas Mirzoeff. 2011. The Right to Look: A Counterhistory of Visuality. *Duke University Press* (2011).

[65] Jeremy Wade Morris. 2015. Curation by code: Infomediaries and the data mining of taste. *European Journal of Cultural Studies* (2015).

[66] Michael Muller, Lydia B Chilton, Anna Kantosalo, Charles Patrick Martin, and Greg Walsh. 2021. GenAICHI: Generative AI and HCI. *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (2021).

[67] Reese Muntean, Alissa N. Antle, Brendan Matkin, Kate Hennessy, Susan Rowley, and Jordan Wilson. 2017. Designing Cultural Values into Interaction. *CHI Conference on Human Factors in Computing Systems* (2017).

[68] Preetam Nandy, Cyrus DiCiccio, Divya Venugopalan, Heloise Logan, Kinjal Basu, and Noureddine El Karoui. 2022. Achieving Fairness via Post-Processing in Web-Scale Recommender Systems. *ACM FAccT* (2022).

[69] Mmachi God'sglory Obiorah, James K. L. Hammerman, Becky Rother, Will Granger, Haley Margaret West, Michael Horn, and Laura Trouille. 2021. U!Scientist: Designing for People-Powered Research in Museums. *CHI Conference on Human Factors in Computing Systems* (2021).

[70] William Odom, Richard Banks, Abigail Durrant, David Kirk, and James Pierce. 2012. Slow technology: critical reflection and future directions. *Proceedings of Designing Interactive Systems* (2012).

[71] Nasher Musuem of Art at Duke University. Retrived April 2024. Act as if you are a curator at the Nasher Musuem of Art. *https://nasher.duke.edu/exhibitions/act-as-if-you-are-a-curator-an-ai-generated-exhibition/* (Retrived April 2024).

[72] University of Oxford. Retrieved August 2023. New Exhibit highlights differences between algorithmic and human curation. *https://www.ox.ac.uk/news/2022-12-08-new-exhibit-highlights-differences-between-algorithmic-and-human-curation* (Retrieved August 2023).

[73] Niko Pajkovic. 2021. Algorithms and taste-making: Exposing the Netflix System's Operational Logics. *Convergence: International Journal of New Media Technologies* (2021).

[74] V. Papa and T. Photiadis. 2021. Algorithmic Curation and Users' Civic Attitudes: A Study on Facebook News Feed Results. *Information* (2021).

[75] William Christopher Payne, Yoav Bergner, Mary Etta West, Carlie Charp, R. Benjamin Benjamin Shapiro, Danielle Albers Szafir, Edd V. Taylor, and Kayla DesPortes. 2021. danceON: Culturally Responsive Creative Computing. *CHI Conference on Human Factors in Computing Systems* (2021).

[76] Judea Pearl. 2009. Causality: Models, Reasoning and Inference, 2nd Edition. *Cambridge University Press* (2009).

[77] Michael Pepi. 2011. Iconology in the Age of the Algorithm. *ARTWRIT* (2011).

[78] Ana Peraica. 2021. The Work of Art in the Age of Neural Reproduction: Works of Rembrandt van Rijn in Convolutional Neural and Generative Adversarial Networks. *Art Documentation: Journal of the Art Libraries Society of North America, volume 40, issue 2* (2021).

[79] Emilee Rader and Rebecca Gray. 2015. Understanding User Beliefs About Algorithmic Curation in the Facebook News Feed. *CHI Conference on Human Factors in Computing Systems* (2015).

[80] The Next Rembrandt. Retrieved 2022. The Next Rembrandt Project. *https://www.nextrembrandt.com* (Retrieved 2022).

[81] Jonathan Roth, Guillaume Saint-Jacques, and Yinyin Yu. 2022. An Outcome Test of Discrimination for Ranked Lists. *ACM FAccT* (2022).

[82] S. E. Sachs. 2020. The algorithm at work? Explanation and repair in the enactment of similarity in art data. *Information, Communication and Society, Volume 23, Issue 11* (2020).

[83] Jennifer A. Sandlin and Jennifer L. Milam. 2008. Mixing Pop (Culture) and Politics": Cultural Resistance, Culture Jamming, and Anti-Consumption Activism as Critical Public Pedagogy. *Curriculum Inquiry Vol. 38, No. 3* (2008).

[84] Christian Sandvig. 2015. Seeing the Sort: The Aesthetic and Industrial Defense of The Algorithm. *Journal of the New Media Caucus* (2015).

[85] Nick Seaver. 2019. Captivating algorithms: Recommender systems as traps. *Journal of Material Culture, Volume 24, Issue 4, pp 421–436* (2019).

[86] Nick Seaver. 2022. Computing Taste: Algorithms and the Makers of Music Recommendation. *University of Chicago Press* (2022).

[87] Ben Rydal Shapiro and Rogers Hall. 2018. Personal Curation in a Museum. *PACM on Human-Computer Interaction, Vol. 2, No. CSCW, Article 158* (2018).

[88] Sarah Sharma. 2020. A Manifesto for the Broken Machine. *Camera Obscura (Durham, NC), vol. 35, no. 2, Duke University Press pp 170-179* (2020).

[89] Xi Shen, Alexei A Efros, and Mathieu Aubry. 2019. Discovering Visual Patterns in Art Collections with Spatially-consistent Feature Learning. *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)* (2019).

[90] Ramya Srinivasan, Emily Denton, Jordan Famularo, Negar Rostamzadeh, Fernando Diaz, and Beth Coleman. 2021. Artsheets for Art Datasets. *NeurIPS Datasets and Benchmarks Track* (2021).

[91] Ramya Srinivasan and Kanji Uchino. 2021. Biases in Generative Art—A Causal Look from the Lens of Art History. *ACM FAccT* (2021).

[92] Ramya Srinivasan and Kanji Uchino. 2021. Quantifying Confounding bias in Generative Art–A Case Study. *IEEE Systems, Man, and Cybernetics Conference* (2021).

[93] Daniel Susser. 2022. Decision Time: Normative Dimensions of Algorithmic Speed. *ACM FAccT* (2022).

[94] Giulia Taurino. Retrieved 2022. Experiments in Machine Unlearning: On Algorithmic Curation and Museums Marginalia. *https://www.thedigitalreview.com/issue01/taurino-machine/begin.html* (Retrieved 2022).

[95] Alex Taylor. 2021. Are streaming algorithms really damaging film? *https://www.bbc.com/news/entertainment-arts-56085924* (2021).

[96] Phi Le Nguyen Alan Wee-Chung Liew Hongzhi Yin Quoc Viet Hung Nguyen Thanh Tam Nguyen, Thanh Trung Huynh. 2022. A Survey of Machine Unlearning. *IEEE S and P* (2022).

[97] Yu Tian, Jianxin Chang, Yanan Niu, Yang Song, and Chenliang Li. 2022. When Multi-Level Meets Multi-Interest: A Multi-Grained Neural Model for Sequential Recommendation. *SIGIR* (2022).

[98] Milka Trajkova, A'aeshah Alhakamy, Francesco Cafaro, Rashmi Mallappa, and Sreekanth R. Kankara. 2021. Move Your Body: Engaging Museum Visitors with Human-Data Interaction. *CHI Conference on Human Factors in Computing Systems* (2021).

[99] Daniel Trielli and Nicholas Diakopoulos. 2019. Search as News Curator: The Role of Google in Shaping Attention to News Information. *CHI Conference on Human Factors in Computing Systems* (2019).

[100] Christopher W. Tyler and Lora T. Likova. 2012. The Role of the Visual Arts in Enhancing the Learning Process. *Frontiers in Human NeuroScience* (2012).

[101] Julia Velkova and Anne Kaun. 2021. Algorithmic resistance: media practices and the politics of repair. *Information, Communication and Society, 24:4, 523-540* (2021).

[102] Quentin Villermet, Jérémie Poiroux, Manuel Moussallam, Thomas Louail, and Camille Roth. 2021. Follow the guides: disentangling human and algorithmic curation in online music consumption. *ACM RecSys* (2021).

[103] Quentin Villermet, Jérémie Poiroux, Manuel Moussallam, Thomas Louail, and Camille Roth. 2021. Follow the guides: disentangling human and algorithmic curation in online music consumption. *Proceedings of the 15th ACM Conference on Recommender Systems* (2021).

[104] Nicholas Vincent, Hanlin Li, Nicole Tilly, Stevie Chancellor, and Brent Hecht. 2021. Data Leverage: A Framework for Empowering the Public in its Relationship with Technology Companies. *ACM FAccT* (2021).

[105] D. Vyas, G. van der Veer, and A. Nijholt. 2013. Creative practices in the design studio culture: collaboration and communication. *Cogn Tech Work 15, 415–443* (2013).

[106] Xi Wang, Zoya Bylinskii, Aaron Hertzmann, and Robert Pepperell. 2020. Toward quantifying ambiguities in artistic images. *ACM Transactions on Applied Perception (TAP)* (2020).

[107] Wikiart. 2020. Visual Art Encyclopedia. *https://www.wikiart.org* (2020).

[108] Wendy R. Williams. 2019. Attending to the visual aspects of visual storytelling: using art and design concepts to interpret and compose narratives with images. *Journal of Visual Literacy* (2019).

[109] Robert Wolfe and Aylin Caliskan. 2022. Markedness in Visual Semantic AI. (2022).

[110] Marisol Wong-Villacres, Carl DiSalvo, Neha Kumar, and Betsy DiSalvo. 2020. Culture in Action: Unpacking Capacities to Inform Assets-Based Design. *CHI Conference on Human Factors in Computing Systems* (2020).

[111] Allison Woodruff, Sarah E. Fox, Steven Rousso-Schindler, and Jeff Warshaw. 2018. A Qualitative Exploration of Perceptions of Algorithmic Fairness. *CHI Conference on Human Factors in Computing Systems* (2018).

[112] Bereket A. Yilma and Luis A. Leiva. 2023. The Elements of Visual Art Recommendation Learning Latent Semantic Representations of Paintings. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23), April 23–28, 2023, Hamburg, Germany* (2023).

[113] Bereket A. Yilma and Luis A. Leiva. 2023. Together Yet Apart: Multimodal Representation Learning for Personalised Visual Art Recommendation. *UMAP '23: Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization* (2023).

[114] Chao Zhang, Cheng Yao, Jiayi Wu, Weijia Lin, Lijuan Liu, Ge Yan, and Fangtian Ying. 2022. StoryDrawer: A Child–AI Collaborative Drawing System to Support Children's Creative Visual Storytelling. *CHI Conference on Human Factors in Computing Systems* (2022).

[115] Qian Zhang, Jie Lu, and Yaochu Jin. 2020. Artificial intelligence in recommender systems. *Springer Complex and Intelligent Systems* (2020).

[116] Shengyu Zhang, Lingxiao Yang, Dong Yao, Yujie Lu, Fuli Feng, Zhou Zhao, Tat seng Chua, and Fei Wu. 2022. Re4: Learning to Re-contrast, Re-attend, Re-construct for Multi-interest Recommendation. *WWW* (2022).

[117] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. *ICCV* (2017).

[118] Joanna Zylinska. 2022. AI Art: Machine Visions and Warped Dreams. *Open Humanities Press* (2022).

[119] Joanna Zylinska. 2023. Art in the age of artificial intelligence. *Science, Vol 381, Issue 6654, pp. 139-140* (2023).