


# ARTIFICIAL INTELLIGENCE, MUSIC RECOMMENDATION, AND THE CURATION OF CULTURE

---

A white paper by Georgina Born, Jeremy Morris,  
Fernando Diaz, and Ashton Anderson



UNIVERSITY OF  
TORONTO

 SCHWARTZ REISMAN INSTITUTE  
FOR TECHNOLOGY AND SOCIETY

**CIFAR**

# CONTENTS

## INTRODUCTION

### 1. ASSUMPTIONS BUILT INTO RECOMMENDATION SYSTEMS

Understanding who builds recommendation systems

Unpacking the theory of the listening subject built into recommendation systems

Exploring the theory of music at play in algorithmic recommendation

### 2. AI'S IMPACT ON PLATFORMIZATION AND MUSIC PRODUCTION

### 3. AI'S INFLUENCE ON THE POLITICAL ECONOMY OF MUSIC DISTRIBUTION

### 4. THE EFFECTS OF GLOBAL SERVICES ON LOCAL MUSIC CULTURES AND ISSUES OF DIVERSITY

### 5. AI'S LONG-TERM IMPACT ON CULTURAL CONSUMPTION

### 6. WHO AND WHAT ARE WE OPTIMIZING FOR?

### 7. RECOMMENDATIONS FOR REGULATORY AND POLICY INTERVENTIONS

Acknowledging and regulating curation and publishing

Concentration and vertical integration

Data collection, use, and trade

Transparency and legibility—and their absence

- a. Transparency concerning consumer information and empowerment
- b. Legibility regarding the nature and functioning of, and controls over, AI curation technologies
  - i. Regulating to give consumers controls
  - ii. Regulating to give artists and creative communities controls
- c. Transparency to enable regulatory means of monitoring, auditing, and evaluating platforms

Diversity of content, and of sources

Artists' and creators' interests—remuneration beyond copyright

### 8. DIRECTIONS FOR FUTURE RESEARCH

## WORKS CITED

## ABOUT US

About the contributors

About the organizations

## REPORT AUTHORS:

**Georgina Born**, Faculty of Music, University of Oxford

**Jeremy Morris**, Department of Communication Arts, University of Wisconsin-Madison

**Fernando Diaz**, Fairness, Accountability and Transparency, Microsoft Research

**Ashton Anderson**, Department of Computer Science, University of Toronto; Schwartz Reisman Institute for Technology and Society

Date of publication:  
June 1, 2021

Production by  
Daniel Browne

Proofreading by  
Daniel Browne  
Jovana Jankovic

Photos by  
Seif Eddin Khayat  
Kelly Sikkema  
Nadia Sitova  
Jess Watters

# INTRODUCTION



*The authors would like to acknowledge that this white paper could not have been written without input from the discussions, meeting notes, and post-workshop conversations and emails contributed by those who attended the CIFAR workshop “AI and the Curation of Culture” in Paris, October 2019. The authors would like to thank Blake Durham, Maria Eriksson, David Hesmondhalgh, Nicole Klassen, Larissa Mann, Alison Powell, Jui Ramaprasad, Reginold Royston, Robert Prey, and Sam Way for their written contributions and additions to certain sections of the document. Of course, since the paper emerged from a shared workshop event, the authors would also like to thank all the participants for their contributions to the various workshop sessions and meeting notes that provided a key source text for the writing of this final white paper. This includes Laurent Charlin, Guillaume Heuguet, David Reiley, Nick Seaver, Bob Sturm, Paul Théberge, and Michael Ugwu.<sup>1</sup>*

Our access to the media that we cherish and love—from music to movies to books and more—is undergoing fundamental changes thanks to recent developments in artificial intelligence (AI). From the use of AI in the production process—there are bots that can generate “photos,” and algorithms that have been signed to “record deals”—to AI-driven recommendation engines that deliver us personalized movie or TV recommendations, artificial intelligence is emerging as a key buzzword and technology in the cultural industries. Media and cultural goods like music, books, and video that have traditionally been distributed through radio and television stations, cinemas, public libraries, and retail stores now circulate on digital platforms and services like Spotify, iTunes, Pandora, YouTube, and Netflix, where a variety of AI-driven technologies shape how this cultural content is packaged, presented, and discovered. These innovations may seem like simply a new way to receive the same kind of media we’ve always enjoyed. But much research remains to be done on the cultural and societal impacts of these AI and algorithmically-driven developments as they influence the production, circulation, and consumption of culture.

To address this need, in the fall of 2019, we held a two-day workshop on “AI, Recommendation, and the Curation of Culture” in conjunction with [CIFAR’s Pan-Canadian Artificial Intelligence Strategy](#) to explore the issues surrounding the increasing use of artificial intelligence and recommendation engines in the cultural industries, using music as a particularly salient case study. The workshop was designed to include a diversity of stakeholders, from users and media scholars to computer

scientists, content creators, and platforms/music distribution companies. We intended to stoke dialogue and discussion amongst these stakeholders in order to help bridge technical and ethical divides between these different communities. Our aim was collectively to think through more socially and ethically informed system development and more technically informed media analyses that are attuned to the power of creative arts at the level of individual identities, local art and music scenes, and regional and (post)national cultural communities.

We began the workshop focused on a few central themes. We were interested, first, in **the long-term impacts of AI-driven technologies on cultural consumption and creation**. For example, AI recommendation systems employ data on consumer behaviour while also directly influencing consumer choices, thereby shaping cultural literacies as well as medium- and long-term trends in consumption and taste. Similarly, algorithmic distribution may influence creators’ incentives about what type of content to produce in order to receive exposure to listeners or reach fans. We intended the workshop to generate discussion on these matters, while also addressing the possibilities of creating and coding algorithms and recommendation systems that work for the good of the various cultures and communities involved.

A second central theme was **the need to weigh against the excessively normative power of existing AI-driven technologies so as to enable them to serve diverse communities and give access to diverse genres of music, culture, and art**. On the one hand, many AI-based recommendation systems incorporate key

<sup>1</sup> We would like to acknowledge the funding provided by CIFAR, through their Pan-Canadian AI Strategy, which supported our event as part of their AI & Society series of workshops. These events aimed to bring together diverse experts from across disciplines, sectors, and borders to address some of the most fundamental challenges posed by AI. Special thanks to Gaga Boskovic, Amy Cook, and Wendy Hallé from CIFAR for supporting and coordinating the workshop. Learn more at [www.cifar.ca/ai](http://www.cifar.ca/ai).

social and cultural variables such as gender, race, ethnicity, and genre as ranking signals that have significant implications for the results provided to consumers. While recommendation algorithms are always socially and culturally inflected, then, relying on overly generalized demographic and identity classifications simplifies more complicated and overlapping categories with the risk of reinforcing rigid, simplistic, and misleading identitarian boundaries. The algorithms can privilege content or relationships that are well-represented and known, already proven and popular over the novel, innovative, and different. When these assumptions get fed back into databases and algorithmic inputs, they perpetuate and deepen essentialist suppositions about identities, practices, and tastes in opaque ways. Even when consumer representations do not explicitly adopt pre-defined categorical variables—as with so-called learned representations—the influence of data biases due either to population demographics or to platform ecosystems and the normative assumptions encoded in objective functions and model architectures can result in similar problems (Broussard 2018, Noble 2018, Benjamin 2019). We intended our workshop to reflect on the ways in which implicit and explicit biases are built into the design of digital platforms and how they are likely to affect the use of such platforms.

We also realized, however, **that the curation of culture has always involved, and always will involve, balancing competing objectives.** Recommendation systems can be optimized for different objectives, including those of consumers (e.g. accessing relevant or stimulating content), creators (e.g. reaching new audiences and fans), or distributors (e.g. gaining revenue). Yet how recommendation systems balance these objectives is opaque. Moreover, while unique access to consumption data allows recommendation systems ostensibly to improve recommendations for creators and consumers, there are incentives to use data to generate first-party cultural content either through in-house creators or generative algorithms, practices that threaten to displace artists and creators—whether or not they have equal or similar access to consumption data. A general principle at stake here is the way that the extraction of personal data has been privatized and corporatized by curation platforms, but as yet without any public interest debate or regulatory intervention in terms of the potential for a rebalancing through accountability and transparency. In the current situation where AI and algorithmically-driven recommendation systems are typically built on closed, proprietary code and data, how can we understand and audit these systems to uncover, and debate, whose interests and objectives are being optimized for, in which ways, and why?

During our two-day workshop, these three themes became the starting point for expansive discussions on the role of AI and algorithmically-driven recommendation systems in the curation of culture. This document presents key debates, ideas, and avenues for future research that emerged from the event. In particular, we report on **six priorities** that emerged from the group, which we use as structuring elements for the body of this document: 1) Assumptions built into recommendation systems; 2) AI's impacts on platformization and music produc-

tion as, in turn, they affect musical forms; 3) AI's influence on the political economy of music distribution, media advertising, and social media; 4) The effects of global services on local musical cultures in relation to issues of diversity; 5) AI's long-term impact on cultural consumption, including the consequences of music services' data collection and curation practices for how users discover and experience music; and 6) Competing logics of algorithmic "optimization." We conclude by drawing out a number of 7) Recommendations for regulatory and policy interventions, and 8) Directions for future research.

Before addressing each of these sections, we want to state at the outset that, as we shared our various definitions of and ideas about AI, we recognized that AI is not just a specific technology, or suite of technologies. It's also a set of discourses about technology that encompasses a series of questions: what does AI mean to everyday users? How does the business press write about AI? And how do companies that use AI talk about it to attract users or investors? AI, we observed, is part technology, part discourse, and part hype—what Kate O'Riordan (2017) calls an "unreal object." It is a technological manifestation, a rhetorical tool, and a rich surface for dreams, fantasies, and cultural and commercial speculations. It is a magnet for venture capital, a strategic tool for attracting research funding, and a lure for businesses and organizations that seek to automate, optimize, and outsource complicated tasks. It is at once a promise, a technical reality, and a volatile experiment that deepens interdependencies between humans and machines in ways we do not yet understand but that are nonetheless advancing with accelerating speed.

AI is real for, and felt by, those living in the midst of its outputs: those profiled by facial recognition systems, those assessed by credit-ranking algorithms, those receiving treatment based on machine-assisted medical evaluations, and those independent artists cut out of the market for popular music by corporate platforms. At the same time, AI is an abstract idea that raises fear, embodies hope, and fills markets, governments, and cultural domains with an intoxicating sense of hypermodernity and cutting-edge progress. AI is "unreal" because it is a set of real technologies but also a receptacle for conflicting, competing, and sometimes contradictory imaginaries of what AI can do or will be able to do.

By demystifying terms like "artificial intelligence," "machine learning," and "algorithmic recommendation," we wanted to account for the fact that these terms have meanings that are wildly divergent, depending on which constituency or audience is using them. From the realm of science fiction, we have visions of AI as a means of using computing to get machines to do tasks that normally require human intelligence. From the perspective of engineers building these technologies, AI is the collection of inferential statistics applied as a central component of advanced computational systems. From the perspective of public and private investors, these terms signal state-of-the-art technologies and forecast profitability. From the perspective of musicians, the same terms connote risks of redundancy and threats to livelihood, but also tools for music creation. From




the perspective of social justice workers, the involvement of AI can ameliorate systemic injustices—but more often reinforces them.

Ultimately, there is nothing artificial about artificial intelligence, nor necessarily intelligent in the usual meanings of the term. Algorithms are deeply human and non-human in both their conception and deployment, so to suggest that AI and algorithms are purely technical entities is inadequate. AI is more than a simple technology; in regard to recommendation it consists, rather, of a set of relationships between creators, curators, audiences, commercial entities, engineers, and machines. If we conceive of AI as unreal in this way, we can understand it as a new incarnation of much longer entanglements between culture and technology. Caught between machine learners, algorithmic classification systems, and techno-assisted ways of seeing, hearing, and sensing, the workshop left us wondering: what is AI doing? What do we want it to be doing? What do we fear? What do we anticipate? How does the hype surrounding AI shed light on currents in society? What does it obscure? And how can we formulate and engage in a positive and productive critique of AI and the uses of AI without contributing to its fetishization?

We reflected, finally, on how many of the claims made about AI and its potential benefits and drawbacks for a cultural industry like music were not necessarily new. Social theorists have long theorized the relationship between human agency, social structures, and technologies, but how can we best draw upon social theories that pre-date these systems to understand how AI and algorithmic recommendation systems enable, transform,

or constrain social action? Should we view algorithms as structural properties of social systems that are both the medium and the outcome of the practices that constitute these systems (Giddens 1984)? Or should we leave behind the sociological concept of structure altogether, and see AI and algorithmic recommendation as themselves actors participating in material-semiotic relations with other actors (Latour 1988)? One productive approach is to view these automated systems as institutions—doing in a new way what social institutions have always done: “making routine decisions, solving routine problems, and doing a lot of regular thinking on behalf of individuals” (Douglas 1986, 47).

It is also worth remembering that the arrival of other, earlier, curation technologies brought about similar levels of hype, excitement, and fear. The advent of radio DJs, the iPod’s “shuffle” button, and online digital distribution: each of these developments raised similar questions about the assumptions built into these then-new models of cultural curation, the effects of the technologies on the production of music and the political economy of the music industry, and their impact in mediating local and global music cultures. While the workshop was very interested in what was new about algorithmic recommendation and other uses of AI in music, we also recognized how these new technologies often re-frame, exacerbate, and revive long-standing issues. In other words, we advocate tempering the hype and claims of novelty surrounding AI and recommendation algorithms with a critical historical perspective that locates AI as simply the latest iteration of a much longer entanglement of cultures and technologies (Marvin 1988, Gitelman 2008, Sterne 2012).



# 1. ASSUMPTIONS BUILT INTO RECOMMENDATION SYSTEMS

---



One of the main ways that AI influences the distribution of cultural goods is through algorithmically-generated recommendations. In online platforms such as Spotify, Netflix, or YouTube, recommendation algorithms automatically suggest pieces of content that a user can choose to consume. Often these recommendations occur in context: the recommended track or album may be different at the startup of the application from when the user reaches the end of an album. Recommendation systems are omnipresent on online platforms so that a user can always make an easy choice. In certain situations, streaming platforms will even “autoplay” algorithmic recommendations; paradoxically, users have their “choice” chosen for them.

Recommendation systems are sociotechnical systems (Bijker, Hughes, and Pinch 1987, Bijker and Law 1992). They are the result of a collective design process, involving many people in many roles—algorithm designers, user data researchers, data scientists and engineers, product managers, and more—who build these systems based on a series of assumptions: about the nature of users, about the problems recommendation engines are apparently solving and the goals they are intended to serve, and about the nature of music itself. Such assumptions are embedded in recommendation systems so that, when the ensuing cultural curation is multiplied across the billions of recommendations presented to users by online platforms, they significantly affect the nature of individual cultural experiences. Moreover, the process is magnified and multiplied cumulatively, across time, and globally, across populations, regions, and cultures. The total magnitude of effects represents an unprecedented degree of automatized intervention in the way people and communities encounter and experience culture. Unpacking the assumptions built into recommendation systems is therefore an urgent task. We break it down into three main elements: 1) Understanding who builds recommendation systems and the social organizations and institutions in which they are embedded; 2) Unpacking the theory of the listening subject built into recommendation engines, one with uncanny resemblance to the neoliberal subject (Rose 1992, Gershon 2011); and 3) Exploring the theory of music at play in algorithmic recommendation,

one with an uncanny likeness to recorded versions of global commercial popular music.

## Understanding who builds recommendation systems

Recommendation systems, like any technology, are shaped by those who create them. The first research direction we propose under this theme is to better understand who is involved in building them, what their demographic profile is, and how this shapes their design practices—for example, through the potentially gendered social imaginaries that recommendation engineers bring to product conceptualization and design (cf. Oudshoorn 2003, Oudshoorn et al. 2004). This suggests the need for research not only on the demographics of leading platforms but also their company structure, the division of organizational roles, team and unit formation, as well as particularly influential figures like user researchers and algorithm designers. To take a much-discussed example: it is well-known that women and people of colour are underrepresented in the technology sector (Broussard 2018). Without enough of these voices “in the room,” the impact of design decisions on specific groups may not be adequately addressed. These underrepresented groups tend also to be poorly represented in the higher echelons of the corporate chain of command at media platform companies. Hence, even when there are diverse voices in the room, they may be present only in positions with relatively little control over decision-making and design. Moreover, recommendation systems are highly technical and require advanced mathematical and computational expertise. But due to academic disciplinary divisions, engineers and researchers with these forms of expertise are unlikely also to possess requisite tools and insights from the social sciences and cultural studies that would equip them to analyze the implications of their design decisions. If appropriate critical cultural expertise is not sufficiently represented in the design community, the potentially negative as well as positive social and cultural implications of recommendation systems are unlikely to be addressed. Understanding the existing demographic and disciplinary makeup of the builders of recommendation systems is, then, one key to unpacking the assumptions built into recommendation systems.

## Unpacking the theory of the listening subject built into recommendation systems

The second element we want to highlight involves unpacking the theory of the listening subject at play in recommendation systems. Those involved in designing recommendation systems rely on models of individual listeners to guide who they are designing for. These models may be informal and intuitive, based on personal experience; qualitative models based on corporate user research; quantitative models based on information retrieval effectiveness “metrics;” or statistical associations between listener behaviour and delayed outcomes like user retention. Such models clearly affect how recommendation systems operate. So the questions arise: how representative are the models of individual users being employed? From what data are these models derived, and interpreted according to which disciplines or paradigms? Are more “valuable” users given more attention than others? Any summation of an evolving user population must leave some people out—who are they? Feedback and data signals play a crucial function in the design process—so if user models evolve, in what ways and how are recommendation system designers updating them? What data is used, and how are metrics and feedback signals interpreted and operationalized? While these questions as yet remain unanswered, one thing is clear: it is that recommendation systems make highly normative assumptions about the listening subjects they model and whom they purport merely to serve.

Recommender models of listening both presume and create what Robert Prey (2018) calls “algorithmic individuation”: a particular, abstracted idea of the listening subject that amounts to a summation of serial acts of individual behaviour, specifically acts of “choice,” that are then built into the algorithm. Crucially, “along with being subjected to identification practices [by recommendation systems], individuals are also subjectified by them” (Prey 2018, 1088)—that is, listeners’ subjectivities are iteratively moulded through interaction with the systems. In recommendation systems, then, identity categories “are not determined at the outset, but rather performed into being through the user’s actions” (ibid., 1088). Probabilities prevail: hence, “A visitor to a website might be identified as a ‘Caucasian’ man with a confidence measure of 79% (Cheney-Lippold 2017, 34) and this measure may rise or fall based on a subsequent purchase on the site. From this perspective, the ‘data subject’ is essentially a process *in development*” (Prey 2018, 1088). For John Cheney-Lippold, such processes create a “cybernetic relationship to identification,” in which essentialist notions of identity are replaced by “pliable behavioural models” (Cheney-Lippold 2011, 168). Similarly, Nick Seaver’s (2019) anthropological study of the music informatics researchers designing recommendation systems explores how certain models of the listening subject—in their case, assumptions about how contemporary listeners have access to more music than they could ever hope to listen to, and therefore need music services to help them manage it—became the basis for the design paradigm behind their recommendation systems.

Precisely how the “algorithmic individuation” is modelled varies across platforms and is constantly being redesigned internally. The “raw data” of a listener comes from a variety of sources. On-platform data is gathered from aggressive telemetry applications. Simple play counts (i.e. counting which specific songs have been played or skipped) are the core of many recommendation system algorithms. But on-platform data also includes timestamps of when songs were played or skipped, contextual information (e.g. was the recommendation requested on the home screen, or on an artist’s profile page?), and even granular interaction data (e.g. scrolling). Off-platform data can be provided by users during registration (e.g. demographic data) or procured from data partnerships with other companies. Both on-platform and off-platform data can be used to derive abstracted concepts. Demographic labels can also be imputed using AI. Songs can be augmented with metadata (e.g. Pandora’s Music Genome Project) which are then aggregated and integrated with metadata collected on individual listeners. These data representing the listener are weighed and leveraged in context-specific recommendation decisions. Moreover, behavioural data can be used to infer whether a specific recommendation was satisfying to the listener. The “data subject” presumed and performed into existence by music streaming services therefore responds to every interaction a listener has with a musical item, which are recorded in real-time, creating a dynamic record of the user’s evolving, non-essential musical tastes and responses. Rather than reified intersectional identity attributes (gender, race, ethnicity, class, age, and so on) being the grounding for how recommendation systems hail the listener, “one’s ‘algorithmic identity’ is in a constant state of modulation” (Prey 2018, 1088). Hence, in addition to the malleability of categorical assignments, the categories themselves are diffuse and change in response to user behaviour, especially in so-called representation learning systems such as deep neural networks.

The kinds of person envisaged and performatively cultivated by recommendation systems are also illuminated by writers who have addressed how neoliberal governmentality shapes our subjectivities (Rose 1992, Gershon 2011). For these authors, neoliberalism issues an injunction regarding the cultivation of individual selves: “Become whole, become what you want, become yourself: the individual is to become... an entrepreneur of itself, seeking to maximize its own powers, its own happiness, its own quality of life, though enhancing its autonomy and then instrumentalizing its autonomous choices in the service of its lifestyle. The self is to style its life through acts of choice, and when it cannot conduct its life according to this norm of choice it is to seek expert assistance” (Rose 1992, 11). AI’s escalating role in delivering exactly the “expert assistance” foreseen here, in many areas of life, is perhaps uncontroversial—and in recommendation systems it comes to fruition in the curation of culture and music. In this vein, the normative listening subject is presumed to be someone faced with an overwhelming world historical archive of music online, someone who, when faced with those “acts of choice” that style her life, feels a need for the assistance of the recommendation algorithm to navigate music and make her relationship



with music manageable. But she is also someone whose musical tastes and interests are presumed to evolve according to a universal logic derived from an analysis of the aggregated behaviour of millions of listeners. Any alternative idea that musical taste or musical identity might evolve or be enlivened by moving in unpredictable directions—by jumps, or breaks, or through unruly processes of branching, drifting, negation, inversion, difference, or dissension, or through camp or surreal associations—is unconscionable in this modelling.

Seaver's anthropological study of music informatics researchers confirms these assumptions about the listening subject built into recommendation systems. His interviewees repeatedly noted "what had come to seem obvious: contemporary music listeners have access to more music than they could ever hope to listen to, and they need something to help them manage it. In conversation, it was easy to forget that on-demand, large-catalog subscription services like Spotify had only existed for a few years... The resulting problem was what psychologists call 'the paradox of choice' (Schwartz 2004) or 'choice overload'" (Seaver 2019, 4-5).

What is certain is, once again, that recommendation systems are highly normative about the listening subjects projected and then designed into the system, whether that normativity relates to the conception of the listening "individual" being catered for, or to assumed segmentations or aggregations of listener populations. Recommender models engaged in algorithmic individuation not only reduce the complex, multifaceted nature of human subjectivities but they downplay and effectively deny the social and the embodied dimensions of music listening. It is the profoundly normative set of base assumptions about the listening subject outlined here on which we urgently need more research, with the results fed back into the cultures of algorithmic design so as to mitigate and transform the existing normative models, and, thereby, their performative effects on listening. We urgently need research probing the distance between such models and the diversity of users' actual listening practices and curatorial desires, particularly those not captured by the models. Building critical self-reflection on these issues into the cultures of design of recommendation systems is, then, a crucial step towards creating more musically, socially, and culturally aware and responsive algorithms and recommendation systems. In naming these issues, we seek to stimulate debate about them with the engineering community, with the intention of achieving a greater diversification of how listening, and the evolution of musical experience, are understood and modelled.

### **Exploring the theory of music at play in algorithmic recommendation**

The third element of the assumptions built into music recommendation systems that we want to highlight is how such systems embody a theory of music: that is, a set of underlying assumptions about what music is. At base this turns on music being conceived as a form of numerical information, a signal. As Seaver puts it: "In the computer, an audio file

is a long list of numbers for telling speakers how to vibrate... The fundamental task of computer audition is to reduce that series of numbers, typically 44,100 of them a second in CDs or MP3s, to a much smaller set that meaningfully represents the content. These smaller sets are called 'feature representations' (in that they represent not the audio itself, but relevant features of it) and they serve as the input for higher-level algorithms" (Seaver 2015, 1). Pursuing the question of what music is, Seaver expands: "To treat music as a signal means to treat it as a series of numbers laden with pattern... At one scale, we find pitches—an A vibrating at 440 Hz. At another, we find tempo—120 beats per minute. Melody, rhythm, and meter repeat on their own timescales, and at a higher scale we find song structure: verse and chorus repeating a few times over the course of a few minutes". Music becomes, then, formal patterns latent in "the data stream, just waiting to be mathematically recognized" (2015, 2).

What we hope to highlight here is how this theory—or ontology—of music has the effect of absencing or externalizing (Callon 1998) crucial features of music and musical experience long cherished by listeners, musicians, and participants in many of the world's diverse music cultures, and long analyzed by humanist scholars from musicology, ethnomusicology, and music sociology: music's embeddedness in and dependence upon the body, the socialities of performance as they are animated among performers and audiences, material sound-producing and sound-transmitting devices such as musical instruments and loudspeakers, and particular venues, locations, and sonic environments.

As such, machine learning-based classifications of audio features may be highly efficient, but they do not necessarily reflect an intelligent analysis of music, nor do they entail any real understanding of the nature of what music is taken to be in many cultures around the world. As yet, recommendation systems do not even process how human auditory systems always mediate acoustic signals, creating subjective experiences of music that are, again, diverse and profoundly encultured. Moreover, computer models of audio features are usually based on very particular qualities drawn from recordings of global commercial popular music, rather than reflecting the breadth of musical sounds and cultures worldwide. We suggest, then, that another area to which existing professional cultures of algorithmic design need to become sensitive is what computational representations of musical sound currently leave out.

Ultimately, the theory (ontology) of music immanent in recommendation systems depends on very particular qualities drawn from global commercial popular music that are then reified by the systems and universalized as though they are essential features of all musics. Central facets of this process are: the substitution of sound recordings for music's existence *per se*, when recording is itself a radical transformation of music and a practice alien to many of the world's music cultures (Deo 2022); the elevation of machine-readable music scores as a core information source, as though, again, they are a



universal feature of the world's musics, when scores are an irrelevance to oral/aural, non-notated musics; the elevation of such features as pitch contour, tempo, metre, and beat as universally relevant and primary expressive musical features as opposed to all the other aesthetic and expressive dimensions central to musics worldwide; and the assumption that the 3 to 5 minute "track" is a universal formal unit of music. Such problematic universalizations of what are very specific and limited aspects of music may in part be explained, once again, by a lack of interdisciplinary collaboration. Notably, the

music informatics research underpinning recommendation system design appears not to have developed in close dialogue with those fields that specialize in the qualitative analysis of music—ethnomusicology, musicology, music sociology, music analysis. The transformation of the science behind music recommendation that would stem from such dialogue—and thus change the theory of music built into recommendation—is urgently needed, but as yet there are few signs of these developments, and especially of them making a difference to recommendation design.



## 2. AI'S IMPACTS ON PLATFORMIZATION AND MUSIC PRODUCTION

---



AI's impact is also being felt directly by musicians creating music with artificially intelligent technologies, as well as by those creating music for platforms that employ AI and recommendation algorithms as part of their back-end infrastructure for making content discoverable. How should we conceptualize the role of AI both as a set of production tools and as an essential part of the infrastructure of music discovery and distribution? And how can we best understand these sometimes exciting and often anxiety-inducing changes within a longer history of sociotechnical and institutional entanglements?

From the perspective of those making music for new platform-based distribution environments, the platforms themselves impact on how musicians craft and create their music. As platforms like Spotify become the primary means of distribution and circulation for music and other content, they begin to shape the ways in which music is produced and readied for them, either through the explicit policies, rules, and guidelines they impose (e.g. songs only accrue royalties after 30 seconds, “albums” can have a maximum 100 songs on them, etc.), or through more hidden acts of infrastructural and algorithmic politics (e.g. some genres and types of songs find more success on certain platforms than others). Cultural commodities become “contingent” in this new environment: producers of cultural content must work to adapt their content for the platforms on which they hope to appear, as well as adjust their business models and plans to match the ways in which the platforms work (Nieborg and Poell 2018, 7).

Whether it is musicians tailoring content to be playlist-friendly on streaming services, or artists tweaking the metadata of their songs so that they surface more readily in searches, or songwriters writing different formats of songs (i.e. many short songs) to receive more plays and higher royalties, there are a number of emerging examples of artists adjusting their music and sounds to match the logics of the streaming platforms on which they appear. The platforms' logics are powerfully performative: they bring about the formats and sounds they favour, while claiming merely to reflect the “nature” of current musical and cultural realities. We need research, then, on the repercussions for musical creativity and music production as musicians continue to uncover these logics and produce under their powerful influence.

For artists creating music by employing recent AI technologies, a whole series of other questions arise. Musicians who have used AI to produce tracks have tended to describe their own involvement in the process not so much as “composing,” but more as “editing” or even “directing.” They set parameters, the machine spits something out, they set new parameters, the machine learns, and so on. The music is composed through constant iteration, like the back-and-forth exchange between a writer and an editor. As Taryn Southern, who released an AI-generated album a few years ago, has described it:

The way I liken the songwriting process in Amper, or with actually many of these tools, is you become more of an editor or a director rather than a composer. You're basically given a ton of raw material and then it's your job to shape that material, to cut that material into something interesting, to transform it. But you're still using the source material from the AI composition tool. (Switched on Pop Podcast, 'AI Music,' 10 min 34 sec)

Of course, the parameters being manipulated are originally set by the AI program itself. Taryn Southern has also talked about how she was drawn to the cinematic soundscapes made by the AI tool Amper, explaining that this was probably because the people who founded Amper came from film composer Hans Zimmer's team, which may make sense of why she found the soundtrack library richer than the rock/pop libraries.

When we look at AI and music production, however, it's important to expand the lens beyond individual artists employing it in their music production practice. In debates about AI, the focus tends to be on the output; but an equally important concern is what is put “into” the AI—the music that constitutes the training set for AI (Drott 2020, Sturm et al. 2019). But other steps in the production chain also matter. For example, AI is increasingly being used as an aid in the A&R process. The recently launched SNAFU Records calls itself “the first full-service record label built on AI-music discovery” of new artists (i.e. A&R). By automating the A&R process, SNAFU claims to be able to identify “undervalued” independent artists earlier than human A&R staff are able to. Once all of these artists are identified by the AI, the most promising 15-20 artists are



then analyzed by the label's in-house team. While we could frame this as an example of the “externalization and automation of taste” (Arielli 2018), we see again that human experts are not so much being removed, but rather being recast as editors or directors in the process—yet without having access to the forms of reasoning about the value of a certain artist, sound, or track being deployed by the machine learner, so that this critically important assessment of value remains opaque. The selection of artists delivered by AI-based A&R processes is, then, entirely dependent on how the system identifies and

determines variables like “undervalued”—how it models what is taken to be musically innovative and of value. Along with examining the theory of music and the theory of the listening subject inscribed in the AI algorithms driving recommendation systems, as noted above, it is therefore clear that we also need to interrogate the theory of value built into these systems, and the way that AI, as well as the increasing pressures to be seen and heard on these crowded platforms, are affecting not just users but the very production of music itself.

**“Along with examining the theory of music and the theory of the listening subject inscribed in the AI algorithms driving recommendation systems [...] we also need to interrogate the theory of value built into these systems, and the way that AI, as well as the increasing pressures to be seen and heard on these crowded platforms, are affecting not just users but the very production of music itself.”**



### 3. AI'S INFLUENCE ON THE POLITICAL ECONOMY OF MUSIC DISTRIBUTION

---



The impact of AI is not just felt at the level of individual artists or platforms. As AI is increasingly integrated into a variety of facets of the music industries, it also begins to re-shape the political economy of music distribution. This intensifies the well-known phenomenon of industry concentration as it is associated with what many now refer to as “platform capitalism” (Srnicek 2017). Concentration of curation and recommendation should concern us since it increases the powers of the platform conglomerates *vis à vis* both creators and consumers; it also impacts negatively on the possibility of unearthing diverse content, thereby disadvantaging local, national, and alternative cultural communities. Since the corporate platforms’ motivations and strategies for promoting certain musics over others are opaque, artists and users do not have the knowledge or tools to understand how AI-influenced recommendation algorithms function and why they recommend what they do. This opacity and the lack of content diversity are synergistic, and the effects are corrosive for democracy and the cultural infrastructure on which it depends. As a recent Canadian government report for UNESCO puts it: “Exposure to [a] diversity of content will contribute to a healthy public discourse, greater social [and cultural] inclusion and a better understanding between countries, cultures and communities” (GCR 2019, 2).

The new platform models for music rely on a heavy investment in datafication: the tendency to convert everything into data, which is then taken to be a marker and a generator of value. Although the music industries have long relied on data (e.g. Billboard charts, Soundscan technologies, Columbia House Record Club, file-sharing data, etc.), the datafication of music is vastly amplified in an era where every play, skip, pause, and “like” can be captured and fed back as inputs for future curatorial decisions and choices. Streaming platforms become data brokers, connecting a vast networked market of advertisers, brands, and listeners. And while purchasing a subscription (at least on Spotify) may buy listeners out of the need to listen to ads, both free and subscription models are built primarily around the collection of data and the extraction of insights from that data. The business of music has, then, morphed to become geared as much around analyzing listening data and crafting discovery algorithms as it is around finding and nurturing emerging and established musicians.

In addition, the confluence of the music and computing industries achieved by the major platforms has resulted in greater precarity for music companies, as they begin to rely on similar funding models to those driving the tech industries. Many Silicon Valley companies exist only to be bought: the company itself is the product, and the data the company collects—whether it’s a music company or a dating app—is an index of that company’s (re)sale value. In an unprecedented shift, in platform capitalism the service provided by a company is tertiary to the company’s datafied commercial value. For many music industry companies, this means that their material interest in music has largely come to be geared around how they can represent themselves to investors via their data holdings—and much less to do with the profitability of their sale of music. In sum, this financialization of the music industries prioritizes data over music, fuelling both the commercial drive to gather, broker, and trade data and the imperative to overcome whatever regulatory barriers with regard to data security and privacy may exist as a result of governments’ attempts to intervene in and oversee these markets. And the wider effect of this political economy of the new industry is to prioritize a deep yet highly selective, quantified, and abstracted knowledge of user habits and preferences over the cultivation of the well-being and flourishing of musicians and audiences.

The implementation of AI and machine learning in music services affects different stakeholders differently. For users, who have now become the owners of information about their music preferences as well as fan labourers (Durham 2018), AI provides a trade-off of “mass personalization” in exchange for the extraction of their personal data. For designers, who labour within firms using machine learning, AI externalizes the collective aspects of music and offers an atomized approach focused on individuals rather than the social relationships at the heart of the erstwhile experience of music. For copyright owners, machine learning and AI can be used to exercise even greater controls over data and distribution. In turn, for musicians trying to succeed on these platforms, who are also users of the platforms, there is a trade-off between potentially gaining added access to user data (although this is far from assured) and concerns about the relatively weak agency they are likely to have over the distribution of and remuneration for their music. In perhaps the most extreme development, AI



has begun to displace musicians' labour when used by companies to generate music (often for television or film), in this way transforming what were previously intermediaries and service providers into "creators" and copyright owners. Finally, advertisers are beneficiaries of the increased amount of user data that AI and machine learning provide, although the process is a dehumanizing one in which listeners are abstracted, reduced, and translated into data points, fuelling the construction of individual listening profiles.

Ultimately, the new political economy of music, like that of many other tech-influenced industries, relies on an amplified kind of data colonialism (Couldry and Mejias 2018) in which audience behaviours and preferences are extracted, traded, and sold to advertisers and other intermediaries, as well as back to users in the form of "socially"-curated playlists. Identifying how these multiple changes alter the political economy and power structures of the former music industries—by identifying historical continuities or shifts in the industry's power relations—is a compelling and a major area for future research.



## 4. THE EFFECTS OF GLOBAL SERVICES ON LOCAL MUSICAL CULTURES AND ISSUES OF DIVERSITY



Cultural diversity, we propose, is itself a supreme human value. As the political theorist Bhikhu Parekh has argued: “No culture embodies all that is valuable in human life and develops the full range of human possibilities. Different cultures thus... complement each other, [and] expand each other’s horizon of thought.” By giving access to other cultures, cultural diversity can enable people to “step out of their culture,” to “see [its] contingency... and relate to it freely rather than as a fate... Since [it] fosters such vital preconditions of human freedom as self-knowledge [and] self-criticism, it is an objective good” (Parekh 2001, 167). The curatorial decisions made by algorithmic recommendation systems significantly affect the diversity of musical content listeners are able to encounter and consume, and at present there is a serious deficit of diversity. Yet posing the question of diversity simply in terms of whether the content of a given streaming service is “diverse enough” is too limited. Instead, we suggest, it is more productive to think of diversity in AI and algorithmic music technologies along four interwoven lines: 1) the social dimension of AI production and the need to diversify the demographic makeup of those responsible for the design of recommendation systems; 2) the question of whose music and which music is made visible and audible on digital music platforms; 3) the nature of the interaction between global services and local musical cultures, traditions and practices; and 4) the challenge of incorporating into the design and affordances of recommendation systems diverse conceptions of both listening subjects and communities of use, so as to combat the extreme uniformities currently built into system design (see section 1, above) and better nurture the diverse musical subjectivities and musical cultures engaging with the technologies.

With respect to the social dimensions of diversity, *design* is clearly a key stage in which engineers “script” envisaged uses into their technologies (Akrich 1992), configuring expectations about both users and patterns of use. As we argued in section 1, these decisions are likely to be influenced by the social identities of the designers—along lines of gender, race, class and ability—since designers tend to conflate their identities and experiences with those of their imagined users. In short, who gets to engineer matters. If there is a lack of diversity among the people designing AI-based recommendation systems, this shapes not only which musics are attended to, but the variety

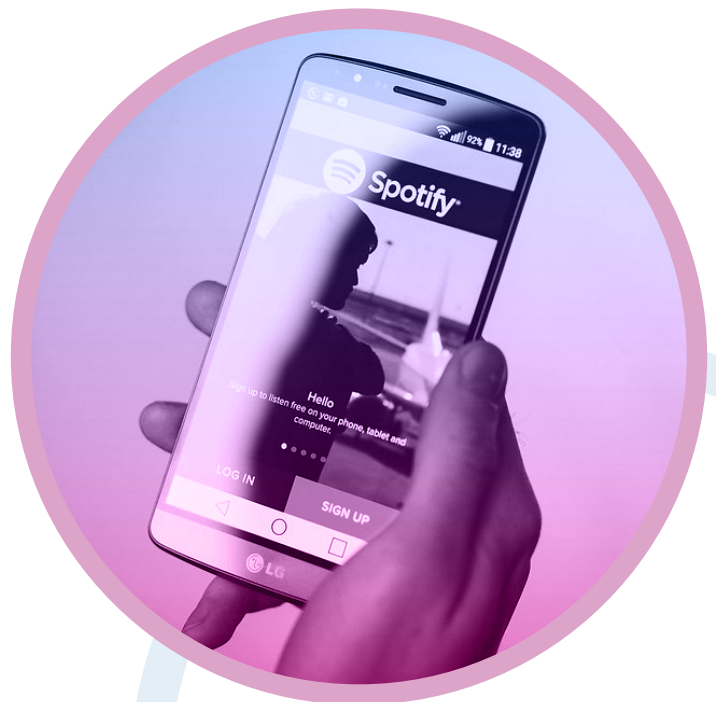
of users’ tastes, needs, and preferences addressed in design practices. Such decisions are further amplified by the motives of the particular platform or company involved (e.g. to gain advertising revenue, increase subscriptions, etc.). Along with advocating greater social diversity among the community of designers in the leading commercial streaming services, then, we call for *greater diversity in the institutional forms of the services themselves*—specifically, the growth of non-profit and publicly-oriented recommendation systems aimed at enhancing human musical discovery and flourishing in ways that go far beyond the profit motives of existing services.

This speaks to the second issue, the question of whose music and which music is made visible and audible on digital music platforms. There are many dimensions at play here, including at the level of content (e.g. genre, region of origin, release date, or popularity), at the level of producers (e.g. gender, age, race, class, and other identity markers of the creators themselves), and at the level of the listeners favoured by a particular service’s music collection (e.g. prioritized demographics, usage habits, identity markers, etc.). In their goal to reproduce what is already popular and to cater to individualized musical preferences, AI-based recommendation systems offer narrow visions of the variety of musical cultures that exist globally and regionally. Given that most services not only elevate Western pop music as the most visible and discoverable content but also as the universal model upon which curation and recommendation systems are built, the profound challenges posed by “other” musics have not yet been registered or fully incorporated into existing services and technologies. Take, for example, the many traditional, non-Western, electronic and computer musics in which melody, harmony, tempo, and rhythm are not primary aesthetic features and instead timbre, gesture, microtonality, melisma, spatialization, or other rhythmic subtleties prevail. Given that these subtleties are much more difficult to extract, analyze, and quantify, AI-based recommendation technologies cannot account for these musics as easily as they can musics that share aesthetic characteristics with mainstream Western pop. How can recommendation services, instead of ignoring the aesthetic features listed above, *take their cues* from such aesthetic markers of musical difference? How can a diversity of musics, especially those traditional and non-Western musics that have resisted incorporation into the global archives of

recorded music, offer models on which to build more diverse curatorial processes as well as different kinds of listening and engagement with musical cultures? This issue is partly about whether or not listeners are being recommended diverse content through interaction with a given system; it is equally about whether or not a service prioritizes different aesthetic features and models of musical production and consumption in its core models and processes.

The third aspect of diversity concerns the need to scrutinize the relationship between global services aiming to serve a universalized global audience and local musical cultures, traditions and practices. From one perspective, since online streaming services boast catalogues composed of many millions of songs, it appears that users can browse and listen to music from all around the world. In this light, it would seem that such globally present services enable artists from every market to have their music exposed to people far beyond their local areas. Yet, without doubt, many local musics remain “undiscovered” and many markets are unaccounted for, and this correlates with both structural and more mundane barriers for artists wanting to be represented on these crowded platforms (e.g. the prevalence of English-language musics, royalty structures, etc.). Certainly, online streaming platforms appear to have the potential to act as a powerful force boosting the distribution and availability of local and regional musics. But in fact, artists from larger, established markets with greater resources and support from the traditional structures of the music industries find themselves more advantaged than ever before. The “winner-takes-all” economics of many streaming services mean that established artists receive far more plays than independent and local musicians, with the effect that only a fraction of regional or local musics will ever travel beyond a country’s borders. Moreover, the fine and subtle gradations of musical difference characteristic of diverse musics may be eroded by the powerful but blunt search features built into algorithmic recommendation systems which, as noted above, have very limited aesthetic diversity built into their design. Indeed, in stark contrast to the global reach achieved by certain musicians through these services, the effect of the extreme centralization of the global platforms is that it may become harder for local musicians to have their music heard *even in their own communities*. Recommendation systems therefore have the potential to act as a neocolonialist force in music, trained on data in which dominant user demographics are over-represented, and using the tastes and preferences embedded in this data to guide the music consumption of other musical cultures. Even when streaming platforms begin to collect data from new markets that they enter, their understanding of it is shaped by the early adopters in those markets, who are more likely to be privileged and elite, their tastes shaped by exposure to Western popular music. Exacerbating these problems, streaming platforms have as yet not been directly incentivized to take action on the distribution of local music. By pursuing goals of, for example, maximizing engagement and growing user-bases, they are therefore likely to affect the consumption of local music in seriously adverse ways.

Finally, we want to raise the issue of how to incorporate diverse conceptions of both listening subjects and communities of use into system design so as to combat the extreme uniformities and the de-contextualized and reductive models of listeners built into music streaming platforms. As noted in section 1, a theory of the human subject is built into the very design of recommendation systems: a human subject who is existentially overwhelmed by the abundance of the global digital music archive, whose evolving taste is structured by “similarity” to what they already know and like, who is utterly individualized, and who seeks to maximize her/his listening events. If this is the driving model of the listening subject, it is one that reduces the cultural diversity of musical practices around listening. While listeners surely have a wide variety of propensities to seek out diverse content—some may be satisfied by a narrow window of content, some may actively seek out new repertoire—these differences will vary across cultures, regions, and demographics and such variety could be better integrated into the design of AI-based recommendation. Furthermore, the design of online streaming services often obfuscates or eliminates important cultural contexts for making sense of a diversity of musics. “Playlists,” for example, can divorce songs from the contexts in which they were meant to appear or position them relative to other pieces in ways that were unintended or that may even be harmful. The architecture of the machine learning algorithms that underlie recommendation systems is focused on the individual—yet understanding and modelling musical cultures and histories is a highly social matter. The cultural and social contexts and communities in which both listeners and musics are embedded are largely absent from current services and technologies. Our question for the service providers and designers is: how could this be otherwise?



## 5. AI'S LONG-TERM IMPACT ON CULTURAL CONSUMPTION

---

The focus in existing debates around AI is often on short-term effects. Our conviction is that it is now urgent and overdue to raise the unpredictable, yet very real, medium- and long-term consequences of algorithmic recommendation and curation as they shape individual discovery, exposure, and consumption, as well as how these processes unfold at a large scale across populations, regions and cultures. Recommendation algorithms automate the cultural and musical space within which individuals navigate music and culture through personalized tracking, customized recommendations, and targeted social identity profiling. What are the effects of these processes on the unpredictable and nonlinear exchanges—with friends, family, and peer groups, with local and translocal, online and offline subcultures—that were previously central to the formation and evolution of individual and collective musical and cultural tastes? Musical identity was formerly intensely entangled in other kinds of evolving relationship with the world, other cultural and social experiences. These are processes that cannot (and arguably should not) be captured by the rule-governed, autarchic relationship to music proffered by AI-driven recommendation. Do recommendation systems, in their drive to forge a path between an individual's musical taste “now” and “what's next” foreclose on or displace these other kinds of contingent, world-embedded, and world-associated musical experiences? We insist that the “reading” of the individual that comes from algorithms and data, stripped of those wider cultural and social ecologies, fuels an individuated, de-culturalized curation that fails to respond to, and ignores, the socially located nature of musical experience.

Cultural homogeneity has long been a concern for media scholars focused on the industrialized production of culture (Adorno 1978 [1938]), and algorithmic curation represents a contemporary intensification of cultural standardization and atomization. AI-driven recommendation encourages every user to have their own individual customized soundtrack, and the mechanics for creating these individuated commodities involves reducing the individual to broad, blunt socio-demographic categories (e.g. women, aged 20-24, who listen to hip-hop). Further, current models of AI and algorithmic curation focus on a client-service type model: how can the algorithm best match what I, as an individual consumer, want? AI-driven services are more than happy to answer these kinds of questions while ignoring their social equivalents: how should we, as various cultures and

communities, find connections through music? Instead, Spotify and like platforms encourage social features only when they can be tracked and linked back to individuals and data-generating events. The result is that musical growth and development as they are cumulatively governed by the algorithm are radically individualized and disembedded in order to be made computable by a series of reductions, formalisms, and rule-governed uniformities.

The likely long-term effects of these mechanisms at the heart of recommendation are powerful and highly disruptive: again, we urgently need research on these questions to feed into new approaches to system design. How will detaching people's taste-forming musical habits from their wider ecologies disrupt ordinary forms of social and cultural relationality, and the linking of music to other forms of experience? And what are the likely biographical and historical consequences of such changes? How will the links between our musical selves and the social and musical communities around us shift if AI-driven algorithmic recommendations are not altered to take into account the hyperlocal needs of the countless musical scenes and subcultures that are responsible for building collective musical cultures?



## 6. WHO AND WHAT ARE WE OPTIMIZING FOR?

---



Working through the themes and challenges highlighted in the previous sections raises key questions around how AI-driven recommendation systems should be optimized, and what notions of “quality” they might seek to embody. The answer to these questions will invariably look different depending on the position one occupies (e.g. musician, music label, data scientist, recommendation system designer, platform provider, etc.). We nonetheless offer the following initial thoughts about rethinking optimization and quality.

The “quality” of an algorithm refers in computer science to how accurately and efficiently it transforms any given input into the desired output. When using the search terms “Kendrick” and “Lamar,” for example, a high-quality algorithm should return Kendrick Lamar’s latest hit songs or albums, rather than, say, Anna Kendrick’s song for the *Trolls* soundtrack. Yet in fact, the term “quality” encompasses much more. For a start, there are platform-centric models of quality. In this case, quality might refer to algorithms that produce results that engender low skip rates, or that weed out spam and false positive results. Platform-centric notions of quality might also be based on opportunities for optimal data extraction, or for sponsored content and targeted advertising. In contrast, user-centric models of quality might emphasize results that deliver fast and enjoyable results that fit with a user’s listening profile. Quality here might mean the right balance of exposure to new and familiar artists, a fitting flow of content that reflects the user’s desired mood, emotions, or other contextual factors. The quality of an algorithm might be conceived differently still from an artist-centric perspective. Artists might describe a quality recommendation algorithm as one that is transparent in terms of the results it provides, or that is as likely to lead to independent as mainstream artists, or that engages with wider cultural factors when providing its results. The existence of various models of quality points, then, to issues of “fairness”: who determines which model of quality prevails? How transparent is knowledge about the “qualities” sought? And how fairly do recommendation algorithms work in terms of not discriminating against certain groups, genres, individuals, or modes of user engagement?

Beyond these perspectives, we might think more broadly about what a culture-centric model of quality would look like. Machine learning-based recommendation systems that create person-

alized bubbles (especially those where “on-device personalization” learns individual preferences and finesses them over time) could—as argued in the previous section—pose serious risks in the medium- and long-term to the existence of shared cultures and even to public cultural knowledge. Personalization is seductive: it seems effortlessly to provide more of “what I most want.” But if we believe in notions of shared and public culture, a narrowing space of “what I want” may need to be punctuated by, or extended to include, other cultural products and experiences, other modes of shared engagement. What would be entailed in designing AI algorithms that define “quality” in terms of supporting and fostering local music scenes, that aim to strengthen the bonds within and between various musical communities, or that favour social connection around music and the arts? What would it mean to take as a central optimization goal the question of how representative of musical and cultural diversity are the dominant platforms? How can we incentivize mechanisms and algorithms that will broadly promote and represent cultural diversity? The logic of these questions points not only to “who” but to “what” we are optimizing for in service design.

The various models of quality mentioned, and the visions of who and what we should be optimizing for, might contain tensions; they may even be incompatible. But if we adopt a culture-centric view of optimization, this is likely to promote a greater diversity of algorithms and AI-based technologies and related innovations in terms of ideas of quality and optimization. Ultimately, the CIFAR workshop envisioned a world in which algorithms, creators, and coders seek to represent and empower both enterprise and cultural diversity and autonomy. We need more reflection on what data means and what can be done with it in these regards; for example, tags, labels, and inventory data could be democratized so as to make them more culturally informed and contextually relevant. Culture must inform curation just as curation inventories and shapes culture; machine learning is a powerful tool that could be employed to optimize an array of social, cultural, and economic values and goals. We urge recommendation platforms, then, to find innovative ways to enable communities to benefit from digitization, and to implement new models of quality and optimization in platform design fuelled by cultural and social scientific insights.

# 7. RECOMMENDATIONS FOR REGULATORY AND POLICY INTERVENTIONS

---



Given the competing interests behind the development and implementation of AI technologies for cultural curation, there are clear tensions between corporate private interests and public interests. Recognition of this fact is signalled by an emerging transnational political consensus that self-regulation by the platforms has reached its limits: hence the need for new thinking about regulation and co-regulation. Recent events in the EU and UK (via Ofcom public consultations), the US (Congressional hearings on platforms and political speech), and Australia (by the Competition and Consumer Commission around social media platforms' publisher responsibilities) signal a growing political and public will to make the major platforms (Netflix, Spotify, social media, etc.) and their AI-related technologies more accountable in terms of transparency and their role in the curation of culture. In this light, we outline suggestions for both general and specific regulatory interventions governing AI-informed cultural curation—among them matters of transparency and legibility in relation to AI technologies and algorithmic processes, and diversity in relation to both content and creative sources. We argue that such interventions, and the principles underlying them, are likely to lead to more equitable, empowering, and creative outcomes for audiences and creators. In the spirit of a white paper, we offer “blue skies” thinking relatively unencumbered by political realism.

## **Acknowledging and regulating curation and publishing**

Thus far, with regard to platforms involved in domains traditionally covered by media, cultural, and communication policies—for example, social media and video streaming services—the main pressures towards regulation have come in relation to issues like “fake news,” misinformation, hate speech, and so on. This has been even more true as social media and other content platforms have made recent decisions to remove controversial posts or even shut down accounts of major public figures and politicians. Debates about these issues reveal how existing internet law and regulation tend to define tech companies as neutral intermediaries (equivalent to telecommunications companies), thereby sheltering them from regulation, rather than as publishers, editors, or curators (Gillespie 2018). The difficulties surrounding regulation in these areas make for an even more unpropitious environment for regulation of cultural (including music) search and recommendation—which

are generally considered relatively unproblematic compared with the obvious “public sphere” functions attributed to news and free speech. Nonetheless, given the needs for regulatory intervention set out in this white paper, we contend that this is the first basic issue that must be faced by policy-makers and regulators: the platforms' publisher and curator functions must be acknowledged along with due public obligations stemming from them, and appropriate regulatory principles and frameworks advanced.

## **Concentration and vertical integration**

Increasingly vast quantities of “digital content”—a euphemism for digitized forms of culture, arts, and information—are being curated by a small number of monopolistic platforms. As well as concentration, the vertical integration of curation platforms is also an obvious area for potential regulatory attention. In contrast to the belief that digital distribution would bring about an industry structure characterized by a “long tail” (Anderson 2006), content aggregators like Amazon, Hulu, and Netflix have entered into digital content creation, vertically integrating their operations. Napoli (2016, 343) notes how recent research finds that “the extent to which recommendation systems use inputs such as demonstrated popularity in making their recommendations can actually lead to greater concentration in consumption.” Effectively, “users are directed to popular content, which in turn enhances the demonstrated popularity of that content, which in turn further increases the likelihood of that content being recommended” (ibid., 344). At the same time, content aggregators unable to obtain licenses to quality content on terms they find favourable tend to get into producing content as an alternative; hence, “licensing challenges beget vertical integration,” while “vertical integration creates disincentives for meeting the licensing fees sought by content creators” (ibid., 347). Moreover, successful content creators have also gotten into distribution (e.g. HBO, CBS). This vortex of forces favouring vertical integration suggests that the internet has not “dramatically reconfigured the incentives for serving niche audience interests” (ibid., 351). Indeed, vertical integration has played a significant role in undermining diversity of content as well as audiences' access to it.

The degree of concentration of major platforms curating cultural material must therefore be a regulatory priority, and it appears fast to be becoming one. Regulation has been and can be designed to control and shape “network effects” for public benefit, not just corporate profit, as has been the case in a number of countries and regions—evidenced by recent anti-trust and competition investigations in Canada, the EU, and increasingly the US. Once again, the political will is intensifying to address such regulatory issues.

### Data collection, use, and trade

Another major area of increasing regulatory concern is the existence of massive structural asymmetries of information between platform operators and their users and suppliers. AI-driven curation depends on the collection and manipulation of, and trade in, vast amounts of personal data, both on consumers and on artists and creators. In this light, governments should foreground the rights, freedoms, and empowerment of the public and vulnerable communities over those of private corporations in regard to information about the gathering and processing of data taken both from the public and from creators.

The EU’s General Data Protection Regulation (GDPR) is currently the most developed regulatory intervention globally on the collection and use of personal data, introduced in 2018 to harmonize data privacy laws across Europe. How, and in what ways, has it been effective? This is still debated. Radical critics have emphasized the way the big data economy favours both corporate and state surveillance, seeing liberal concerns with privacy and data regulation as inadequate responses (Andrejevic 2020, Zuboff 2019). A worrying and under-discussed feature of GDPR implementation is the arbitrary imbalance in how GDPR has been applied across different sectors. Thus, academic research has been burdened with onerous GDPR compliance obligations when compared to both journalism and artistic practices; in marked contrast, platform corporations appear effectively to evade GDPR enforcement, collecting and transacting consumer data in de facto covert ways with impunity, and without effective transparency or accountability to consumers and citizens. This has been accompanied by a systematic attrition of former channels for consumer accountability—direct, interactive means to ask for additional information, register complaints, or request compensation. The power to erect such asymmetries can surely be attributed in part to the “first mover” advantages stemming from the unimpeded establishment of monopolies in the early decades of internet-based cultural commerce.

A specific issue regarding data protection and privacy concerns is the use by streaming platforms of large bodies of user data in combination with datasets from other sources derived from the commercial trade in data, a process that can lead to de-anonymization of putatively anonymous data. Moreover, when data is held and re-used by platform companies to profile particular groups of consumers, this increases the risk that this data becomes de-anonymized or that profiling will infringe

on individual and privacy rights. We suggest that the porosity of borders between databases and the trading and sharing of databases across multiple platforms—sometimes owned by the same entity—point to the need for regulatory attention to be triggered when such databases reach a certain size or aggregation, itself requiring transparency.

### Transparency and legibility—and their absence

Responses to the collection, use, and trade in big data by curation platforms has often taken the form of calls for greater “transparency.” But, ironically, this term is itself profoundly opaque, bundling together distinct if related issues that we intend to disentangle. The challenges covered by “transparency” are of at least three types: a) those to do with the provision of greater consumer information and empowerment; b) those to do with the nature and functioning of AI technologies, and the capacity of artists and consumers to exercise controls over them; and c) those associated with regulatory and/or governmental means of monitoring, auditing, and evaluating the platform companies and their services. All three challenges turn on institutionalizing means to empower consumers, artists, and regulators *vis à vis* curation platforms.

**a. Transparency concerning consumer information and empowerment:** An obvious feature of the use of personal data by music streaming services and their recommendation systems is that no transparency is currently provided at the interface level. Instead, a kind of ersatz “transparency” prevails: the uses to which customer data are put are covered in unreasonably lengthy and arcane end-user-license agreements that are virtually impossible for consumers to read and understand in their entirety. In this light, as means to redress the information asymmetry in data collection and use, regulation should take three complementary forms: it should create legal obligations on services to provide greater transparency over what personal data is being collected, why, and how it is being used; it should require platforms to present this information to users in easy-to-understand, brief, and straightforward ways—it should, in other words, make it available, readable, and comprehensible; and it should require platforms to restore effective older means whereby customers are empowered to ask for further information, register complaints, and/or request compensation through direct, human customer service interactions. Only by putting the onus on platforms to make genuinely transparent the algorithmic processes determining the uses made of personal data and their implications for privacy will users be empowered to be aware of such uses and implications, and to make informed decisions on that basis.

**b. Legibility regarding the nature and functioning of, and controls over, AI curation technologies:** It is often remarked that the functioning of machine learning technologies and recommendation algorithms are “black boxed.” Such opacity has been linked, again, to calls for transparency. However, calling AI curation technologies “black boxed” elides three distinct aspects of the phenomenon. It refers first to the commercial interest in maintaining proprietorial secrecy over algorithms

and how they function. Globally, the past decade has seen an accelerating growth of patenting of AI technologies, with an increase of over 400% worldwide in the number of published AI patent applications (IPO 2019, 1). In Europe, although algorithms are considered “computational” and “abstract” in nature and were previously deemed unpatentable, once applied to a “technical problem” or shown to have a real-world application, patents are now regularly permitted under the European Patent Convention. This aspect of opacity lies at the commercial core of curation platforms and, at present, it is likely to resist transparency imperatives.

A second meaning of “black boxing” is captured by Bruno Latour’s observation that “scientific and technical work is made invisible by its own success;” hence, “the more science and technology succeed, the more opaque and obscure they become” (Latour 1999, 304). If philosophically teasing and pleasing, the point has little obvious relevance for regulation; moreover, it is actually in tension with a little-discussed third and more compelling dimension of opacity, which follows.

A third aspect of the “black boxing” or opacity specifically of software technologies derives from the real sociotechnical difficulty, even for the most expert programmers, in deciphering on the basis of code itself how a given piece of programming works (Born 1996, 1997). This feature is a well-known problem within professional programming communities (e.g. Juhár 2019). It is this third meaning of opacity that is likely to be most amenable to regulation, and while it may seem to be in contradiction with the secrecy required by patenting, this is not so. For it is not about releasing source code; rather, it turns on recognition of the need to make the nature and functioning of AI-driven curation systems comprehensible and accessible to users. Achieving this depends, first, on the system’s functions being deciphered at a high level, and the resulting understanding then being translated and rendered *legible* for unskilled users in ways that invite and support non-trivial engagements with the system. This is an entirely new sociotechnical challenge for commercial platforms, and we propose the term *legibility*—as opposed to transparency—to capture this process of rendering algorithmic processes comprehensible and accessible to users and thereby open to users’ engagements and interventions. This is a design philosophy to which platforms, and the AI community at large, have not yet risen. If achieving this legibility might appear onerous, then it also potentially represents one of the most important regulatory interventions to correct existing asymmetries of information and power between platform operators and users in relation to recommendation systems, potentially empowering both consumers and artists. We would add here the generative potential for perspectives from science and technology studies (STS) to be brought directly into dialogue with AI curation engineers to achieve the kinds of legible, open systems we are advocating. Just this kind of interdisciplinary bridging between STS and engineering is envisaged by Peter-Paul Verbeek when writing on the politics and ethics of technological design, with the goal of fostering critical and reflexive paradigms in which “the ethics of engineering

design... take more seriously the moral charge of technological products and rethink the moral responsibility of designers accordingly” (Verbeek 2006, 379).

By proposing the regulation of a mandate for legibility into the design of AI-based recommendation and curation systems, we highlight how the technical processes and sociotechnical associations underpinning the systems should be made comprehensible and accessible in non-trivial ways to non-technical users. At the simple end, this might take the form of flagging that a playlist is automatically rather than humanly generated, or results from a combined practice; at the more complex end, it might mean making legible the sonic features and algorithmic processes underlying the system’s operationalization of “similarity.”

Greater legibility over the internal processes and functioning of algorithmic curation systems might fuel an additional step: it could enable, and be accompanied by, the regulatory goal of requiring the delegation of greater and more diversified controls over recommendation systems to users—to consumers, but also to artists and creative communities. It might even mandate the design of systems open to diverse, entirely unforeseen user practices (Akrich 1992). Thus, in addition to the previous consequences—e.g. flagging whether a playlist is automatically or humanly generated, or making legible the sonic features and algorithms fuelling “similarity”—legibility might engender system designs that empower users to select their own criteria or principles on which basis to organize or order a given musical space curatorially.

- i. **Regulating to give consumers controls:** Currently, recommender systems are optimized for user retention and content hours. Instead, regulation could attempt to empower users through specifying industry-wide measures intended to encourage more diverse modes of cultural circulation by diversifying the nature of user interfaces and their modes of address and by delegating varying degrees and kinds of controls to users. What if music recommendation was built not on a logic of “similarity” but of burgeoning kinds of diversity, to be specified, selected, or developed by the user? This would be to design “co-production” into recommendation systems, while also making the intellectual and creative work of curation itself more legible. It would celebrate, valorize, and expand upon the creative role of curator, DJ, or playlist creator.
- ii. **Regulating to give artists and creative communities controls:** Automated recommendation and curation would also benefit from regulatory attention to the problematic ways in which decisions are currently made about cultural expression without direct and conscious (not data-generated) input from the originating communities. Hence, curation systems might also be regulated for increased legibility for artists, creators, and source communities as key stakeholders—giving them levers to control, for example, which elements of their creative outputs are distributed, the contexts in which recommendation



of these outputs is or is not appropriate, the kinds of metadata attached to the outputs, or where outputs are placed within the system. Artists and source communities might also select how they choose to be remunerated. Such approaches could carry wider entailments for representation and labour; for they suggest that curation platforms might be required to have a duty to creators and source communities (especially vulnerable ones) to diversify the social representation driving their practices, not just through ever-wider data gathering techniques but by hiring creators from marginalized communities into decision-making, curating, and coding positions. The premise here is that greater social diversity of designers and coders is likely to have effects on the nature of the algorithms and software developed. Of course, we would refute any mechanical essentialism here: such effects cannot be guaranteed.

**c. Transparency to enable regulatory means of monitoring, auditing, and evaluating platforms:**

This final sub-category addresses how transparency must be required of corporate platforms in order for services, targets, and performance to be monitored, audited, and evaluated by regulators. Transparency of this kind encompasses both the previous sub-categories. The challenge is that both kinds of information (consumer data and algorithmic processes) have been considered proprietary and form the basis of the platforms' business models. From the vantage point of this white paper, however, given the current imbalance of information and power, we assert a new principle: if existing business models entail the erection of what have now come to be seen as unacceptable proprietary walls around the aggregation, reuse, and sale of personal data, or unacceptably opaque algorithmic processes, then the logical implication is that these models will have to be systematically challenged by regulation, and changed.

A more radical approach to regulating the structural asymmetries of power built into the platforms' big data economy, implied by this principle, would be fundamentally to question the basic premise that personal data should routinely be commercially expropriated and traded at all—with or without the knowledge and legal consent of the individuals involved. Is it time to assert the contrary principle, both legal and ontological: that personal data are—or should be made—*inalienable*, a derivative of the person from whom they emanate and whom they reductively describe or perform (Prey 2018)?

Today, transparency and accountability are often invoked in bad faith as primary solutions to serious ethical concerns about the behaviour of corporate platforms and big tech companies. Yet, as Marilyn Strathern argues, mandates of transparency often fail to bring about insight into actual institutional or technological performance, since “everything that is brought to the surface is then deliberately hidden inside again” (Strathern 2000, 314-8). Mandates of transparency therefore often exacerbate the very problems they are intended to mitigate: demands for “more user information” can be satisfied by cynical “data dumps,” or through interfaces that provide only

superficial access or insight. Accountability, in turn, through its key instrument, auditing, often generates performances of accountability that are ruled, ironically, by unintelligible “KPIs” (key performance indicators); indeed, “auditing claims to deliver an ideal of transparency,” yet its criteria are vague and “the audit process itself is often opaque, closed to public scrutiny” (Born 2005, 238).

To achieve the equitable outcomes for artists and audiences sought by this report, regulation for legibility, transparency, and accountability must therefore anticipate the likelihood of ambiguities, evasions, and trickster moves in compliance. Policy must be unambiguous about its aims to make curation and recommender systems accessible, legible, responsive, and socially, culturally and ethically responsible through specific and actionable guidelines that go beyond the usual platitudes about making them “more transparent” or “more diverse.”

**Diversity of content, and of sources**

As a recent Canadian government report puts it, “quantity of content does not equal diversity of content.” Thus, “although more and more cultural content is being produced and experienced, it may be more challenging for citizens to find or be exposed to a diversity of content” (GCR 2019, 2). Relying on existing algorithmic curation and recommendation systems, users may be exposed only to content sourced from culturally dominant countries and regions, biasing exposure away from minority genres and cultures. Moreover, such algorithms lead to “filter bubble” or “echo chamber” effects whereby consumers are provided with content that stems from the system's profiling of their existing tastes and interests, narrowing exposure and reducing the diversity of cultural encounters—without the mechanisms for that narrowing or filtering being publicly available or understood. Artists and musicians lacking significant industry backing, moreover, face challenges of raising the visibility of their work and gaining fair or sustainable remuneration for it. The effect is that only a minority of artists and source communities gain the means and incentives to distribute and sell their cultural goods in the digital environment.

Regulation for diversity in AI-driven cultural curation platforms should take into account two related dimensions of the problem. The first concerns the diversity of cultural content made available and promoted via curation and recommendation. A key challenge here revolves around the scientific fields responsible for developing the techniques embodied in recommendation systems—in music, the field of music information retrieval (MIR). It is an accepted strand of criticism within the MIR community that the techniques and parameters it employs tend to derive from and reflect globally dominant areas of commercial popular music. Yet those techniques and parameters are applied in powerful curation algorithms as though they were applicable universally—with inevitable *de-pluralizing* effects. So, one imperative to increase the diversity represented in music recommendation is for MIR radically to expand its musical referents and parameters (which is already an incipient debate in the field itself) (Born 2020).

The second diversity vector potentially amenable to regulation concerns the *sources* of digital content and the need for a much wider range of source communities, regions, and creative labour to be represented and made available. However, this requires care and caution. Marginalized communities, Indigenous communities, hypervisible (Reddy 1998), and historically exploited communities also require protection from incursions on their autonomy. Curation—including both inclusion and classification without consent—can be a violent or disrespectful act. It may be that the UNCHR’s recognition of rights of cultural survival, and of the need for IP protections for traditional knowledge and traditional cultural expression, can provide leverage here through the principle that traditional knowledge and cultural survival involve the right to make decisions about which aspects of one’s culture are put into circulation or otherwise made available. This might well become possible through regulating for the legibility of AI-based curation and recommendation systems, in that, as detailed earlier, such systems might be required in future to give artists and artistic communities levers to exert controls over which elements of their creative outputs are distributed (if any), in what contexts, how they are placed, and how they are remunerated. Indeed, if we understand cultural survival and, even better, flourishing to imply not only the survival of works of culture but the existence of conditions to enhance the survival and flourishing of the communities from which such cultural works derive, it follows that members of originating cultural communities should be given due representation in the terms of participation and decision-making roles. It means, finally, that those artists and communities who deem the online curation of their cultural output to be deleterious to their ongoing flourishing should have the means to veto its curation.

At base, AI-driven curation poses the core challenge of updating traditional content regulation for the era of machine learning, so as to address automated cultural curation and correct existing biases in representation. But it also requires content regulation to be updated by reference to the present public embrace of principles of equality and equity for under-represented and marginalized artists and communities. In principle, then, as well as acting on the negative critique of biases and of the mis- and under-representation of some areas of content, regulatory intervention should espouse positive expansion of both kinds of diversity outlined above: diversity of content, and of sources. The two are not reducible to each other, and such interventions would increase exposure to, and the visibility and discoverability of, under-represented types of content and source communities by mandating editorial processes built into the functioning of algorithms themselves. This is the novel challenge posed by AI-based curation.

### **Artists’ and creators’ interests—remuneration beyond copyright**

Numerous problems exist presently with respect to how artists and creators are remunerated by curation platforms (Hesmondhalgh 2020). As an indicator, the distribution of royalties for music streaming sites favours superstar artists. Most streaming sites work on a percentage of revenue model (although Deezer

has recently moved to what they call a “user-centric payment system”) such that artists receive a percentage of revenues generated by the site that corresponds to the percentage of streams they generate. Given that users pay a fixed amount to stream an unlimited number of songs, the system can drown out smaller artists and leave them with a small portion of the revenues for any streams their music receives. Exacerbating this situation is the lack of transparency in how royalties are calculated and distributed. Recommendation systems are often designed in such a way that they perpetuate the rich-get-richer effect, compounded by the “cold start” problem, strategies such as “hit song science,” and the targeting of music to consumers based on existing music consumption—all of which narrows the opportunity for new music discovery and weakens the incentives for new and innovative artists to emerge on these platforms. As the Canadian government report cited before puts it, in the area of remuneration significant obstacles exist “in the form of a lack of data from platforms, a lack of standards for rights management, and the sheer scale and power of online platforms, which far outpaces that of individual creators or collectives” (GCR 2019, 6-7).

As a result, a number of regulatory options might be considered in order to achieve greater equality, fairness, and transparency with regard to the remuneration of artists, source communities, and creative industries, thus boosting their economic sustainability. Copyright and its limits should be one focus: presently, individual artists have little power in negotiations over copyright terms with the big tech companies and major curation platforms, as well as in their attempts to have copyright infringements enforced. In this light, policy might support the creation of copyright collectives to represent large groups of creators and increase their bargaining power in dealings with the platforms and tech firms. In addition, a burden should be placed on platforms to be more transparent in terms of remuneration as a prerequisite for greater equity. Thus, in the context of a chronic lack of understanding about how the platforms’ remuneration schemes function, regulation should oblige platforms to provide the necessary information in accessible and comprehensible form so that content creators are equipped with the knowledge and skills to derive due remuneration from their works. Copyright itself as a source of remuneration might even be reconsidered, given its limitations—for example, its individualistic framing of the economics of creativity, such that it cannot reflect the complex divisions of labour of many areas of creative practice, and its inability to protect the cultural and intellectual property of collective creative communities, such as those of Indigenous groups. Policy might propose alternative economic models to complement copyright, such as taxation of curation platforms to provide funds to support artists and creative communities, or remunerating the presence of content on platforms as opposed to the current focus on consumption or use. In sum, innovation to develop alternative economic models is urgently required from regulators, designed so as to offer incentives for adopting, and synergies with, the principles outlined in previous sections.

## 8. DIRECTIONS FOR FUTURE RESEARCH

---

We want finally to acknowledge the pressing need for empirical research to expand our understanding of the major trends in AI-influenced curation that we have touched on. To take four examples:

1. We need research that explores alternative economic models such as those outlined in the previous section, with an eye to their likely consequences in terms of supporting the sustainability and flourishing of a more diverse and extensive population of artists and creative communities than at present.
2. We require deep and longitudinal studies of the evolution of content diversity on platforms like Spotify, Apple Music, and Amazon, as well as the forces weighing against or in favour of greater content diversity.
3. The world and its cultural resources will benefit, as we have argued, from a more diverse institutional ecology supporting AI-based curation and recommendation, including alternative design philosophies and market strategies that might be pursued by non-profit and public service media organizations guided by ethical principles.
4. Above all, we require horizon-scanning research on the medium- and long-term consequences for both artists and users of the radical experiments in cultural curation, distribution, and experience unleashed by AI technologies—consequences that have largely been overlooked by research and policy to date.

# WORKS CITED

- Adorno, T. W. (1978 [1938]). "On the fetish character in music and the regression of listening." In *The Essential Frankfurt School Reader*. A. Arato and E. Gebhardt (eds). Oxford: Blackwell.
- Akrich, M. (1992). "The de-description of technical objects." In *Shaping Technology / Building Society: Studies in Sociotechnical Change*. W. Bijker and J. Law (eds). Cambridge, MA: MIT Press: 205-224.
- Anderson, C. (2006). *The Long Tail: Why the Future of Business is Selling Less of More*. New York: Hachette Books.
- Andrejevic, M. (2020). *Automated Media*. New York, NY: Routledge.
- Barry, A. (2001). *Political Machines: Governing a Technological Society*. London: Athlone Press.
- Benjamin, R. (2019). *Race After Technology: Abolitionist Tools for the New Jim Code*. Cambridge: Polity Press.
- Bijker, W., T. Hughes, and T. Pinch, eds. (1987). *The Social Construction of Technological Systems*. Cambridge, MA: MIT Press.
- Bijker, W. and J. Law, eds. (1992). *Shaping Technology / Building Society: Studies in Sociotechnical Change*. Cambridge, MA: MIT Press.
- Born, G. (1996). "(Im)materiality and sociality: The dynamics of intellectual property in an artificial intelligence research culture." *Social Anthropology* 4(2): 101-116.
- Born, G. (1997). "Computer software as a medium: Textuality, orality and sociality in an artificial intelligence research culture." In *Rethinking Visual Anthropology*. M. Banks and H. Morphy (eds). New Haven, CT: Yale University Press: 139-169.
- Born, G. (2005). *Uncertain Vision: Birt, Dyke and the Reinvention of the BBC*. London: Vintage.
- Born, G. (2020). "Diversifying MIR: Knowledge and real-world challenges, and new interdisciplinary futures", *Transactions of the International Society for Music Information Retrieval* 3(1): 193-204. DOI: <http://doi.org/10.5334/tismir.58>
- Broussard, M. (2018). *Artificial Unintelligence: How Computers Misunderstand the World*. Cambridge, MA: MIT Press.
- Callon, M. (1998). "An essay on framing and overflowing: Economic externalities revisited by sociology." *The Sociological Review* 46(1): 244-269.
- Cheney-Lippold, J. (2011). "A new algorithmic identity: Soft biopolitics and the modulation of control." *Theory, Culture & Society* 28(6): 164-181.
- Cheney-Lippold, J. (2017). *We Are Data: Algorithms and the Making of Our Digital Selves*. New York, NY: New York University Press.
- Couldry, N. and U. Mejias. (2018). "Data colonialism: Rethinking big data's relation to the contemporary subject." *Television and New Media* 20(4): 336-349.
- Deo, A. (2022). "Oral traditions in the aural public sphere: Digital archiving of vernacular musics in North India." In *Music and Digital Media: A Planetary Anthropology*. G. Born (ed). Durham, NC: Duke University Press.
- Douglas, M. (1986). *How Institutions Think*. New York, NY: Syracuse University Press.
- Drott, E. (2020). "Copyright, compensation, and commons in the music AI industry." *Creative Industries Journal*. DOI: 10.1080/17510694.2020.1839702
- Durham, B. (2018). *Regulating Dissemination: A Comparative Digital Ethnography of Licensed and Unlicensed Spheres of Music Circulation*. PhD thesis, Faculty of Music, University of Oxford, UK.
- European Union (2018). Directive (EU) 2018/1808 of the European Parliament and of the Council (Audiovisual Media Services Directive).
- Gershon, I. (2011). "Neoliberal agency." *Current Anthropology* 52(4): 537-555.
- Giddens, A. (1984). *The Constitution of Society. Outline of the Theory of Structuration*. Berkeley, CA: University of California Press.



- Gillespie, T. (2018). *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*. New Haven, CT: Yale University Press
- Gitelman, L. (2006). *Always Already New: Media, History and the Data of Culture*. Cambridge, MA: MIT Press.
- Government of Canada Report (GCR) (2019). *Report on the International Meeting on Diversity of Content in the Digital Age*. Ottawa, Government of Canada Report for UNESCO.
- Hesmondhalgh, D. (2020). "Is music streaming bad for musicians? Problems of evidence and argument." *New Media & Society*. DOI: <https://doi.org/10.1177/1461444820953541>
- Intellectual Property Office (IPO) (2019). *Artificial Intelligence: A Worldwide Overview of AI Patents and Patenting by the UK AI Sector*. London: Intellectual Property Office.
- Juhár, J. (2019). *Supporting Source Code Annotations with Metadata-Aware Development Environment*. 2019 Federated Conference on Computer Science and Information Systems (FedCSIS), IEEE.
- Latour, B. (1988). "Mixing humans and nonhumans together: The sociology of a door-closer." *Social Problems* **35**(1): 298-310.
- Latour, B. (1999). *Pandora's Hope: Essays on the Reality of Science Studies*. Cambridge, MA: Harvard University Press.
- Marvin, C. (1988). *When Old Technologies Were New: Thinking about Electric Communication in the Late Nineteenth Century*. New York, NY: Oxford University Press.
- Napoli, P. M. (2016). "Requiem for the long tail: Towards a political economy of content aggregation and fragmentation." *International Journal of Media & Cultural Politics* **12**(3): 341-356.
- Nieborg, D. B., and T. Poell. (2018). "The platformization of cultural production: Theorizing the contingent cultural commodity." *New Media & Society*. DOI: <https://doi.org/10.1177/1461444818769694>
- Noble, S. (2018). *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York, NY: New York University Press.
- O'Riordan, K. (2017). *Unreal Objects: Digital Materialities, Techno-scientific Projects and Political Realities*. London: Pluto Press.
- Oudshoorn, N. (2003). "Clinical trials as a cultural niche in which to configure the gender identities of users: The case of male contraceptive development." In *How Users Matter: The Co-Construction of Users and Technology*. N. Oudshoorn and T. Pinch (eds). Cambridge, MA: MIT Press: 209–227.
- Oudshoorn, N., E. Rommes, and M. Stienstra. (2004). "Configuring the user as everybody: Gender and design cultures in information and communication technologies." *Science, Technology and Human Values* **29**(1): 30–63.
- Parekh, B. (2001). "Rethinking multiculturalism: Cultural diversity and political theory." *Ethnicities* **1**(1): 109-115.
- Prey, R. (2018). "Nothing personal: Algorithmic individuation on music streaming platforms." *Media, Culture & Society* **40**(7): 1086-1100.
- Reddy, M. T. (1998). "Invisibility/hypervisibility: The paradox of normative whiteness." *Transformations: The Journal of Inclusive Scholarship and Pedagogy* **9**(2): 55-64.
- Rose, N. (1992). "Governing the enterprising self." In *The Values of the Enterprise Culture: The Moral Debate*. P. Heelas and P. Morris (eds). London: Unwin Hyman.
- Schwartz, B. (2004). *The Paradox of Choice: Why More is Less*. New York, NY: Ecco.
- Seaver, N. (2015). "Listening." In *Computing Taste: The Making of Algorithmic Music Recommendation*. PhD thesis, Department of Anthropology, University of California, Irvine.
- Seaver, N. (2019). "Too much music: Overload and informatic cosmology." Unpublished manuscript, last modified August 2019.
- Srnicek, N. (2017). *Platform Capitalism*. Cambridge, UK: Polity Press.
- Strathern, M. (2000). "The tyranny of transparency." *British Educational Research Journal* **26**(3): 309-321.
- Sterne, J. (2012). *MP3: The Meaning of a Format*. Durham, NC: Duke University Press.
- Sturm, B. L. and G. Wiggins. (2019). "The mismeasure of music: On computerised music listening and analysis." In *The Oxford Handbook of Music and Corpus Studies*. D. Shanahan, A. Burgoyne, and I. Quinn (eds). Oxford: Oxford University Press.
- Sturm, B. L., M. Iglesias, O. Ben-Tal, M. Miron, and E. Gómez. (2019). "Artificial intelligence and music: Open questions of copyright law and engineering praxis." *Arts* **8**(3):1-15.
- Switched on Pop (2018). "AI Music: Can AI "Algorithms" Write Pop Songs? (with Taryn Southern)." Podcast. *Episode 84*. May 4, 2018.
- Verbeek, P.-P. (2006). "Materializing morality: Design ethics and technological mediation." *Science, Technology, and Human Values* **31**(3): 361-380.
- Zuboff, S. (2019). *The Age of Surveillance Capitalism*. London: Profile Books.

# ABOUT US



## About the Contributors

**Georgina Born** is Professor of Music and Anthropology and a Senior Research Fellow at Oxford University, and a Professorial Fellow of Mansfield College. Born's research combines ethnographic studies and theoretical writings on music, media, and cultural production. From 2010 to 2015 she directed the ERC-funded major research programme "Music, Digitisation, Mediation: Towards Interdisciplinary Music Studies." Born is a Fellow of the British Academy, and chairs the section on Culture, Media and Performance. She has recently been awarded a second ERC advanced grant for "Music and AI: Building Critical Interdisciplinary Studies."

**Jeremy Morris** is an associate professor in the Department of Communication Arts at the University of Wisconsin-Madison. Morris's research focuses on new media use in everyday life, specifically the digitization and commodification of cultural goods. He is the author of *Selling Digital Music, Formatting Culture* (2015) and co-editor of the collection *Appified: Culture in the Age of Apps* (2018).

**Fernando Diaz** is a research scientist at Google Research Montreal. His research focuses on the design of information access systems, including search engines, music recommendation services, and crisis response platforms. Previously, Fernando was the assistant managing director of Microsoft Research Montreal, and a director of research at Spotify, where he helped establish its research organization on recommendation, search, and personalization. Fernando's work has received awards at SIGIR, WSDM, ISCRAM, and ECIR. He is the recipient of the 2017 British Computer Society Karen Spärck Jones Award.


**Ashton Anderson** is an assistant professor in the Department of Computer Science, University of Toronto, and a faculty affiliate with the Vector Institute and the Schwartz Reisman Institute for Technology and Society. Anderson's research in computational social science encompasses a diverse range of questions at the intersection of AI, data, and society. His work has appeared in venues including the *Proceedings of the National Academy of Sciences*, *Management Science*, and the International Conference on Machine Learning.

## About the Schwartz Reisman Institute for Technology and Society

**The Schwartz Reisman Institute for Technology and Society (SRI)** was established through a generous gift from Canadian entrepreneurs Gerald Schwartz and Heather Reisman in 2019. SRI is a research and solutions hub dedicated to ensuring that powerful technologies like artificial intelligence are safe, fair, ethical, and make the world better—for everyone. SRI develops new modes of thinking in order to understand the social implications of technologies in the present age, and works to reinvent laws, institutions, and social values to ensure technology is designed, governed, and deployed to deliver a more just and inclusive world. SRI researchers range in fields from law to computer science, engineering, philosophy, political science, and beyond. SRI draws on world-class expertise across universities, government, industry, and community organizations to unite fundamental research on emerging technologies with actionable solutions for public policy, law, the private sector, and citizens alike.

## About CIFAR

**CIFAR** is a Canadian-based global research organization that convenes extraordinary minds to address the most important questions facing science and humanity. The AI & Society program, one of the objectives of the [CIFAR Pan-Canadian AI Strategy](#), develops global thought leadership on the economic, ethical, political, and legal implications of advances in AI. These dialogues deliver new ways of thinking about issues, and drive positive change in the development and deployment of responsible AI.



# **ARTIFICIAL INTELLIGENCE, MUSIC RECOMMENDATION, AND THE CURATION OF CULTURE**

---

A white paper by Georgina Born, Jeremy Morris,  
Fernando Diaz, and Ashton Anderson

June 1, 2021

[srinstitute.utoronto.ca](http://srinstitute.utoronto.ca)  
[hello@torontosri.ca](mailto:hello@torontosri.ca)

 @TorontoSRI