

ARTIFICIAL INTELLIGENCE AND MUSIC DISCOVERY

By Stephanie Bonjack and Nicole Trujillo

INTRODUCTION

We live in surprising times. Who would have guessed at the number of social, political, biological, environmental, and technological watershed moments that have happened since 2020? It is mind-blowing. And yet, at every significant moment, we find music and musicians. Music as a healing balm, music as a rallying cry, music as a creative output for difficult times. Now, we turn our attention to something that seems new, to have emerged out of these turbulent waters, and that is artificial intelligence. Once again, music and musicians are there. In this article, we seek to explore the relationship between musicians, libraries, and artificial intelligence. We examine the tools already in use for music discovery and how librarians have been employing artificial intelligence to enhance these tools. We look at ways librarians can be proactive in the development and integration of artificial intelligence in library products. Finally, we ask how artificial intelligence might improve the discovery and access of music and music research, both for scholarly inquiry and for performance.

MUSIC AND THE MACHINE

The launch of ChatGPT and similar software has brought renewed interest in the field of artificial intelligence (AI), especially in the field of machine learning. In machine learning, large datasets are fed to specific programs that will then learn to do higher function outputs, such as analysis and chat, by using learning algorithms and probabilistic outcomes. While much of the media coverage has centered on ChatGPT's human-like ability to summarize and converse, machine learning and other forms of AI have been used in software since the 1950s.¹

AI and Computer Music

At first glance, AI may seem very far removed from the field of music, especially classical music, which often relies on historical precedent. And yet, the use of artificial intelligence has intrigued music theorists and composers for over half a century. Schüler traces the history of the use of computers for generative music purposes to 1957, in which Frederick P. Brooks ingested hymn tunes into an early computer at Harvard University and synthesized new melodies from the data set.² According to Schüler, the 1960s saw the expanded use of computers

Stephanie Bonjack is the Head of the Howard B. Waltz Music Library at the University of Colorado Boulder, where she is an Assistant Professor in the University Libraries.

Nicole Trujillo is the Access & Discovery Librarian at the University of Colorado Boulder, where she is an Assistant Professor in the University Libraries.

¹ Michael J. Wooldridge, *A Brief History of Artificial Intelligence: What It Is, Where We Are, and Where We Are Going* /, First U.S. edition. (Flatiron Books, 2021), 15.

² Nico Schüler, "From Music Grammar to Cognition of Music in the 1980s and 1990s: The Surplus History of Computer-Based Music Analysis," *Muzikoloski Zbornik* 43, no. 2 (2007): 371–96. xx

for music analysis, often with folk songs, which were simpler and shorter than multipart art music.

In the 1970s, the use of computers for music analysis and composition (or Computer Music) expanded dramatically, as did the scholarly record for this new field. The International Computer Music Association (ICMA) hosted its first conference in 1974. The organization continues today and its conferences are a mix of “professional paper presentations and concerts of new computer music compositions.”³ Shortly after the founding of ICMA, MIT Press began publishing *Computer Music Journal* in 1977. In its first issue, Dexter Morrill published “Trumpet Algorithms for Computer Composition,” a primer on designing and using a computer program to produce natural and unnatural sounding brass tones for use by composers.⁴ From its inception, this journal has served as a publication output on AI in music theory and composition. More recently, the International Conference on Computational Intelligence in Music, Sound, Art and Design, which is part of the larger EvoStar conference group, provides an international stage to those “working on the application of Artificial Intelligence techniques in creative and artist fields.”⁵ At the 2023 conference, papers were presented on topics such as: using GuitarPro tablature for live music generation⁶ and the application of Neo-Riemannian Theory in generating film soundtracks⁷. It is clear that artificial intelligence has been a continuous thread in the interdisciplinary field that is Computer Music. It has engaged music theorists, composers, and scholars in wide-ranging areas of scholarship and music creation. The advent of new AI tools, like ChatGPT, will surely bring a new expansion to this field.

Music Information Retrieval

Related to this very large and established field is the newer field of music information retrieval (MIR). MIR grew as a field in the late 1990s as part of a growing interest in online music recognition and searching in digital libraries. This interest led to the establishment of the first International Symposium on Music Information Retrieval in 2000, later The International Society of Music Information Retrieval Conference.⁸ Conference proceedings examine specific AI technologies and how they play a role in music generation, music description, and music algorithms. For example, a 2021 paper by Parada-Cabaleiro, et al, proposed “a score-based

³ “International Computer Music Conference Proceedings - About,” accessed June 16, 2023, <https://quod.lib.umich.edu/i/icmc/about>.

⁴ Dexter Morrill, “Trumpet Algorithms For Computer Composition,” *Computer Music Journal* 1, no. 1 (1977): 46–52.

⁵ “EvoStar 2022 – The Leading European Event on Bio-Inspired Computation,” accessed June 16, 2023, <https://www.evostar.org/2022/>.

⁶ Sara Adkins, Pedro Sarmiento, and Mathieu Barthet, “LooperGP: A Loopable Sequence Model for Live Coding Performance Using GuitarPro Tablature,” in *Artificial Intelligence in Music, Sound, Art and Design*, ed. Colin Johnson, Nereida Rodríguez-Fernández, and Sérgio M. Rebelo, Lecture Notes in Computer Science (Cham: Springer Nature Switzerland, 2023), 3–19, https://doi.org/10.1007/978-3-031-29956-8_1.

⁷ Simon Colton and Sara Cardinale, “Extending Generative Neo-Riemannian Theory for Event-Based Soundtrack Production,” in *Artificial Intelligence in Music, Sound, Art and Design*, ed. Colin Johnson, Nereida Rodríguez-Fernández, and Sérgio M. Rebelo, Lecture Notes in Computer Science (Cham: Springer Nature Switzerland, 2023), 67–83, https://doi.org/10.1007/978-3-031-29956-8_5.

⁸ “ISMIR Conferences,” ISMIR, accessed June 23, 2023, <http://localhost:4000/conferences/>.

machine learning approach for the classification of texture in Italian madrigals of the 16th century.”⁹ A 2020 conference proceeding examines the conference history and concludes that we are currently in a third phase of music discovery, the Phase of Recommender Interfaces and Continuous Streaming, where large music and personal datasets combined with AI methods allows music to be serendipitously found based on previous selections and local context.¹⁰

User-Based Music Research

Adjacent to Music Information Retrieval is the user-focused literature around patron music-seeking behavior. Clark and Yeager’s user study “Seek and You Shall Find? An Observational Study of Music Students’ Library Catalog Search Behavior” exemplifies this body of literature. Clark and Yeager demonstrate the difficulties students experience in finding specific pieces of classical and jazz music in the library online catalog, and the relative ease with which they find media in the same system.¹¹ Dougan’s earlier study along these lines finds similar results. She concludes that music researchers “need the best example that will meet their needs right then, and/or allow them to follow along to other sources from it.”¹² This cascade of sources and information-seeking behavior is explored by Chandler in his user study on how music directors and conductors find repertoire. He demonstrates the use of multiple search channels employed by conductors to find music, but libraries were not robustly utilized. Instead, personal networks comprised of other conductors as well as established repertoire lists were more commonly used.¹³ This study does not point to the library catalog as a barrier to information seeking, but the fact that some music populations conduct their music discovery almost entirely away from the music library is compelling. In Dougan’s study “Information Seeking Behaviors of Music Students,” she observes undergraduate frustration with finding scores in the library catalog as well as the use of multiple non-library sources for finding sound recordings and music scores.¹⁴

The moment is clearly ripe for a change in how music is discovered in library catalogs and in library systems. The literature demonstrates that these systems are not serving patrons

⁹ Emilia Parada-Cabaleiro et al., “Automatic Recognition of Texture in Renaissance Music,” *Proceedings of the 22nd International Society for Music Information Retrieval Conference* (Online, ISMIR, November 7, 2021), <https://doi.org/10.5281/zenodo.5624443>. 509.

¹⁰ Peter Knees, Markus Schedl, and Masataka Goto, “Intelligent User Interfaces for Music Discovery” 3, no. 1 (October 16, 2020): 165–79, <https://doi.org/10.5334/tismir.60>. 170.

¹¹ Joe C. Clark and Kristin Yeager, “Seek and You Shall Find? An Observational Study of Music Students’ Library Catalog Search Behavior,” *The Journal of Academic Librarianship* 44, no. 1 (January 1, 2018): 105–12, <https://doi.org/10.1016/j.acalib.2017.10.001>.

¹² Kirstin Dougan, “Finding the Right Notes: An Observational Study of Score and Recording Seeking Behaviors of Music Students,” *The Journal of Academic Librarianship* 41, no. 1 (January 1, 2015): 61–67, <https://doi.org/10.1016/j.acalib.2014.09.013>. 67.

¹³ Martin Chandler, “The Information Searching Behaviour of Music Directors,” *Evidence Based Library and Information Practice* 14, no. 2 (June 15, 2019): 85–99, <https://doi.org/10.18438/ebliip29515>.

¹⁴ Kirstin Dougan, “Information Seeking Behaviors of Music Students,” *Reference Services Review* 40, no. 4 (January 1, 2012): 558–73, <https://doi.org/10.1108/00907321211277369>.

well and many musicians form their own methods and information channels as workarounds to the existing tools. The inclusion of AI in library systems presents a possible solution to these issues or an enhancement of the current search process. Here are some areas in which we see that happening.

AI DIRECTIONS FOR MUSIC RESEARCH

Augment Library Interactions with Conversational Search

There's nothing more infuriating than trying to address a reference question you can't answer easily, especially if it's at a desk or over live chat. Pival suggests keeping a Bing Chat (or ChatGPT) open on your computer as a "guide on the side" to help you with reference questions.¹⁵ He advocates for Bing Chat, as it connects with the live internet and has a higher version of the GPT software running behind it than the free version of ChatGPT. Lund and Wang asked ChatGPT directly "How can ChatGPT be used to improve research and scholarship in academia?" and the AI suggested it could help with literature review by summarizing articles or providing lists of articles based on keywords.¹⁶ This seems like an especially helpful suggestion when one is trying to find content on a specific topic quickly. Keyword generation is a strength with this software. We recently asked ChatGPT to generate a list of double reed instruments like the Heckelphone (a niche instrument from the 19th century by a German maker). ChatGPT generated a list of standard double reed instruments in the modern orchestra, Western period instruments, as well as double reed folk instruments from Korea and China. The AI also gives you the option to regenerate the response, which in this case expanded the list.

One can see the application of conversational search in instruction settings as well. When working with students to search databases, using ChatGPT to identify keywords is an effective strategy. We asked the AI to suggest keywords for finding articles on trans voices, and it came back with subtle nuances we might not have brainstormed ourselves, which we were able to use to find targeted articles. Dolan and Freer have many things for instructors to try using ChatGPT, including: "creating a lesson plan, drafting learning objectives, brainstorming classroom activities, summarizing content into bullet points or slide titles, or explaining a topic in simple terms."¹⁷ They warn of "hallucinations," which are fake content generated by the AI that sounds real. Lists of articles are prone to this. Still, as a tool to enhance classroom teaching and learning, we agree with Dolan and Freer that ChatGPT is worth trying.

Create Tools to Assist in Metadata Clean up and Creation

¹⁵ Paul R. Pival, "How to Incorporate Artificial Intelligence (AI) into Your Library Workflow," *Library Hi Tech News* ahead-of-print, no. ahead-of-print (January 1, 2023), <https://doi.org/10.1108/LHTN-03-2023-0052>. 15.

¹⁶ Brady D. Lund and Ting Wang, "Chatting about ChatGPT: How May AI and GPT Impact Academia and Libraries?," *Library Hi Tech News* 40, no. 3 (January 1, 2023): 26–29, <https://doi.org/10.1108/LHTN-01-2023-0009>. 27.

¹⁷ Natalia Dolan and Jennifer Freer, "5 Things to Know about Generative Text AI Tools... That Might Be Outdated or Upgraded by the Time of Publication," *Academic BRASS* 18, no. 1 (Spring 2023). No page in document?

Music libraries generally rely on the MARC metadata standard to represent a large portion of library-owned physical scores and music recordings.¹⁸ Creating quality MARC records for these resources requires a deep knowledge of music, languages, and library cataloging rules. However current issues around staffing in libraries mean that many of these records are created with only the briefest of metadata, if at all.¹⁹ To overcome these challenges, music catalogers can use AI technology to complement existing workflows. For example, the MusicFormGenre tool was used by Brigham Young University Libraries to add Music Genres to thousands of catalog records, making those records browsable by instrument.²⁰ Other tools include Culturegraph, which was developed by the German National Library (DNB) and creates clusters of MARC records that can then be used for metadata analysis and authority work.²¹ Another is CCLitBox, which automatically assigns Colon Classification numbers to works on WikiData.²² Many of these tools are still in the experimental phase, with different machine learning techniques and data sets being analyzed for proof of concepts. Music librarians, with their specialist knowledge of specific MARC fields and desired music-specific search filters should be involved in these projects with the goal of enhancing and updating current MARC with useful search metadata.

Integrate Music Holdings into Commercial Discovery Systems

Commercial music metadata also uses AI to create metadata for their records, and these additions often surpass what our current MARC records contain. For example, Spotify uses natural language processing to assign metadata terms to millions of records on the fly allowing it to match songs with emerging popular terminology. The company also uses AI technology to match audio files of brand-new songs with existing music audio to improve listener recommendations.²³ Google and Shazam have each amassed a large enough set of audio data to develop searches that use imaging of audio files or sound clips to create searches that use sound,

¹⁸ Chris Holden, Keith Knop, and Nara Newcomer, "Music Discovery: Past, Present, and Future," *Notes* 75, no. 4 (2019): 591–619. 608.

¹⁹ Meg Fisher and Pauline Rafferty, "Current Issues with Cataloging Printed Music: Challenges Facing Staff and Systems," *Cataloging & Classification Quarterly* 61, no. 1 (January 2, 2023): 91–117, <https://doi.org/10.1080/01639374.2023.2186999>. 106.

²⁰ Jeff Lyon and Greg Reeve, "Implementing the Genre and Medium of Performance Algorithm in a Local Catalog," (presentation at Music OCLC Users Group Annual Meeting, Norfolk, VA, February 26, 2020). accessed November 10, 2023, <https://musicoclcusers.org/meetings/moug-2020-norfolk-virginia/>

²¹ Angela Vorndran and Stefan Grund, "Metadata Sharing – How to Transfer Metadata Information among Work Cluster Members," *Cataloging & Classification Quarterly* 59, no. 8 (November 17, 2021): 757–74, <https://doi.org/10.1080/01639374.2021.1989101>.

²² Carlo Bianchini and Stefano Bargioni, "Automated Classification Using Linked Open Data. A Case Study on Faceted Classification and Wikidata," *Cataloging & Classification Quarterly* 59, no. 8 (November 17, 2021): 835–52, <https://doi.org/10.1080/01639374.2021.1977447>.

²³ Ipshita Sen, "How AI Helps Spotify Win in the Music Streaming World," *Outside Insight* (blog), May 22, 2018, <https://outsideinsight.com/insights/how-ai-helps-spotify-win-in-the-music-streaming-world/>.

not text, as the input.²⁴ One literally hums into the microphone to search! On the flip side, Apple's recent release of the Apple Music App relies on high quality hand inputted metadata by classically trained musicians to address many of the issues unique to Classical music.²⁵ This demonstrates that while AI can supplement some metadata creation roles, there is still a need even in the modern marketplace for a highly skilled workforce of catalogers. Both of these options are cost and time prohibitive for libraries.

Acknowledging the high bar of competing with large for-profit commercial projects, libraries have explored opportunities for integrating library-held resources into commercial search. In particular, library products have leveraged relationships with Google and integration into the Google Knowledge Panel to better connect patrons with local library resources. The Google Knowledge Panel is AI technology that pieces together a card of entity related information using a combination of user search term and location.²⁶ In 2017 OverDrive, a large ebook and audiobook provider, announced that OverDrive ebooks were now findable on the Google Knowledge Panel.²⁷ In 2020 EBSCO announced their purchase of the linked-data company Zepheira, whose expertise was essential in building a new product, EBSCO BiblioGraph, which integrates library MARC holdings into the Google Knowledge Panel. In 2022 OCLC followed suit, releasing news that they were working with Google to integrate library holdings into the Google Knowledge Panel and in Google Book Search.²⁸ So far these integrations have been book specific, but one could easily imagine a related integration for music streaming and music scores. Music librarians can advocate for large library lending platforms, such as Naxos or Alexander Street, to partner with Google for a similar integration, or work with OCLC or other cataloging partners to further integrate non-book entities into commercial search engines.

Evaluate and Suggest Improvements for Music Discovery Platforms

Since much of AI use in search is new, under development, and expensive, it is no surprise that many of our library-specific multimedia interfaces lag behind in incorporating AI into search. This is unfortunate, as users expect these interfaces to provide them with experiences

²⁴ Christian Frank, "The Machine Learning Behind Hum to Search," November 12, 2020, <https://ai.googleblog.com/2020/11/the-machine-learning-behind-hum-to.html>.

²⁵ Ben Cohen, "Apple Wants to Solve One of Music's Biggest Problems," *Wall Street Journal*, March 30, 2023, sec. Business, <https://www.wsj.com/articles/apple-music-classical-app-3325c6ba>.

²⁶ "About Knowledge Panels - Knowledge Panel Help," accessed June 23, 2023, <https://support.google.com/knowledgepanel/answer/9163198?hl=en>.

²⁷ Matt Enis, "Google Directs Users to Library Ebooks from OverDrive," *Library Journal*, accessed June 23, 2023, <https://www.libraryjournal.com/story/google-directs-users-library-ebooks-overdrive>.

²⁸ "OCLC and Google Now Connect Web Searchers Directly to Library Collections," OCLC, June 7, 2022, <https://www.oclc.org/en/news/releases/2022/20220413-google-search-links-directly-to-library-books.html>.

similar to those like Spotify and Netflix..²⁹ One way to begin critically examining these interfaces beyond usability studies is to ask vendors about the AI technologies they currently employ in their platforms and which ones are on the roadmap for future releases. If vendors are not forthcoming with this information, music librarians can use a checklist of technologies to measure a maturity scale of adoption. Research that identifies AI technology integrations valuable for discovery are invaluable resources for helping music librarians identify and prioritize checklist items. For example, a recent article by the developers of the PubMed medical search lists query manipulation, author name disambiguation, automatic indexing, and Learning to Rank Algorithms as important AI technologies used to improve their search.³⁰ These evaluations could then record the gap between commercial and library-centric products and help define the most productive pathways forward for improving user experiences.

CONCLUSION

Any mention of AI technologies needs to address the fact that these techniques have several criticisms, specifically about their ethicality. For example, if you use generative AI to compose a piece of music is it under your copyright, or that of the technology? Should Chat GPT be used in public communications, or will these communications come across as inauthentic? Will our library administration begin using machine learning technology to replace skilled cataloging labor with lesser (but cheaper) work? Do we feel sufficiently supported and knowledgeable to point out some of the inequalities programmed into the computer code that might lead to a racist “Are you looking for...” or “Other music you might like” result? We feel that our involvement in this area of research is necessary for guiding the future of AI tools used in libraries. Our datasets, which are used in many AI research projects, can be larger and more representative of diverse groups. Our library searches can be examined for potential biases, critiqued, and improved. Our professional participation in this field is essential for fully exploring the possibilities, good and bad, of new AI tools and futures.

Now is the time to identify underdeveloped areas of accessibility that might be addressed by AI. It is clear that discovering music, especially classical music scores, continues to be difficult for patrons. Is there a role in which AI might allow patrons to use more natural language to find what they are looking for? In the concert hall, supertitles are used for opera translations and are projected above the stage. But for every performance with a pre-concert lecture, why not use the auto-generation of captions to make it easier for everyone in attendance to understand what is being spoken on stage? Captioning assists not just hearing-impaired individuals, but

²⁹ Amalia Beisler, Rosalind Bucy, and Ann Medaille, “Streaming Video Database Features: What Do Faculty and Students Really Want?,” *Journal of Electronic Resources Librarianship* 31, no. 1 (January 2, 2019): 14–30, <https://doi.org/10.1080/1941126X.2018.1562602>. 25.

³⁰ Lucy Kiester and Clara Turp, “Artificial Intelligence behind the Scenes: PubMed’s Best Match Algorithm,” *Journal of the Medical Library Association : JMLA* 110, no. 1 (n.d.): 15–22, <https://doi.org/10.5195/jmla.2022.1236>. 110.

those for whom English is a second language.³¹ Also, captions generally reinforce learning, a desired outcome for all of us.

AI and tools like ChatGPT can be viewed in many different lights. They can be seen as threatening the existing order, or as tools to enhance and further our work. We are buoyed by the approach composers and music theorists have taken over the last half century, in which AI was seen as a tool to explore, help develop, and play with. The existence of AI didn't stop composers from composing, but it gave those who chose to use it additional materials and compositional fodder. It is possible AI might do the same for librarianship, if we choose to pursue it.

³¹ Mark Feng Teng, "Vocabulary Learning through Videos: Captions, Advance-Organizer Strategy, and their Combination," *Computer Assisted Language Learning* 35, no. 3 (2022): 518-550, doi: <https://doi.org/10.1080/09588221.2020.1720253>