Introducing the TISMIR Education Track: What, Why, How?

MEINARD MÜLLER (iD)

SIMON DIXON (iD)

ANJA VOLK (iD)

BOB L. T. STURM (iD)

PREETI RAO (iD)

MARK GOTHAM (iD)

*Author affiliations can be found in the back matter of this article

ABSTRACT

This editorial introduces the new education track for the Transactions of the International Society for Music Information Retrieval (TISMIR) and aims to provide guidance to both prospective authors and users of this track's material regarding its context, goals, and scope. To begin, we offer TISMIR-specific context, including the journal's history, its unchanged scope and remit, and the motivations behind introducing the new track. This context is supplemented by broader insights into developments in the field of Music Information Retrieval (MIR), the personal pedagogical experiences of the authors, and the rapid, extensive development of Open Educational Resources across various domains. We highlight the key characteristics of educational articles in general and explore why the music domain may provide an intuitive and motivating setting for education across various levels and disciplines. The education track aligns with existing tracks in terms of TISMIR's dedication to scientific research in MIR, broadly defined as the processing, analyzing, organizing, and creating of music and music-related information using computational methods. Educational articles within this track maintain the high standards expected in terms of scientific rigor, clarity of language, and compelling presentation. However, they differ in their focus on a tutorial-style delivery and their emphasis on existing MIR research methods, techniques, principles, and practical matters relevant to the diverse interests of the MIR community. Through this editorial, our objective is to offer guidance, clarify review criteria, and stimulate discussion on crafting effective educational articles, thereby laying the foundation for a broader discourse on education within MIR and beyond.

CORRESPONDING AUTHOR:

Meinard Müller

International Audio Laboratories Erlangen, Germany

meinard.mueller@audiolabs-erlangen.de

1 INTRODUCTION

Over the past 25 years, the field of music information retrieval (MIR) has developed and advanced in terms of problem scope, methodologies, and applications. Initially focused on simple tasks and methods with small datasets, MIR has now expanded to encompass a wide range of concepts, models, and algorithms, greatly enhancing our ability to access, analyze, understand, and create music. Due to the intricate and diverse nature of music, MIR plays a crucial role as a discipline that connects technical fields like signal processing, machine learning, and information retrieval with a multitude of other fields such as musicology, library sciences, cognitive sciences, psychology, philosophy, ethics, and law. Furthermore, MIR provides an excellent training ground for engineers in a variety of fields not directly related to music, as the skills learned are highly transferrable to other multimedia domains.

The International Society for Music Information Retrieval (ISMIR) is a non-profit organisation seeking to advance research in the MIR field. Since its inauguration in 2000, the ISMIR annual conference has become the flagship event of the society, bringing together academic, industrial and artistic researchers and practitioners from around the world to present their latest research findings, exchange ideas, and foster collaborations. The open-access ISMIR conference proceedings reflect the state of the art in the field and have garnered recognition, being well-cited in the research community.[1]

In 2018, the Transactions of the International Society for Music Information Retrieval (TISMIR) was established to provide a platform for the dissemination of the highest quality and most substantial scientific research in MIR (Dixon et al., 2018). Complementing the ISMIR conference proceedings, TISMIR allows for the publication of more comprehensive and in-depth articles. Through a multi-stage peer review and revision process, authors can obtain high-quality feedback from MIR experts, ensuring the rigor and quality of the published research. Besides research articles that present unpublished original research outcomes and applications, TISMIR also supports two other tracks: overview articles that provide comprehensive reviews of broad MIR research problems, and dataset articles that showcase novel data gathering and annotation efforts (an essential, and often undervalued part of the field).

Computational research does not operate only in isolation but within a context of market relevance, not solely in the proverbial 'ivory tower.' The significance of music-related research has become increasingly evident with the widespread adoption of digital music platforms and streaming services. Two commonly acknowledged truths underscore the ubiquity of music: firstly, its prominent role in every society, across all times and places; and secondly, the progressive increase in access

and availability of music, particularly in the modern era. Against this backdrop, it is crucial to highlight the unprecedented level of ubiquitous access facilitated by recent technological advancements, notably in audio data compression during the 1990s, and the subsequent emergence of streaming services and platforms. These services empower us to engage with music seamlessly, whether as passive listeners or active participants, regardless of time or location.

1.1 EDUCATION

Partly related to this ubiquity of music computing and the attendant access to music that it brings, MIR has not only emerged as a distinct *research* field but has also become increasingly attractive in *education* worldwide. Prominent universities now offer specialized courses in MIR and music processing across various disciplines such as computer science, engineering, and digital humanities.[2]

Music, with its ability to engage and fascinate people, serves as a powerful tool for interactive learning, providing tangible examples, hands-on exploration, and opportunities for experimentation (Müller et al., 2021). The interdisciplinary nature of MIR exposes students to a wide range of subjects, fostering a holistic understanding of both technical and cultural aspects. Furthermore, the practical applications of MIR, including music recommendation systems and audio transcription, resonate with learners, enhancing their comprehension and skill development through real-world applications.

As MIR gains prominence in research, educational literature on the subject is gradually becoming more available, although the pace of this growth remains relatively slow. Examples include textbooks on MIR topics in recent years such as those by Müller (2007); Müller (2015, 2021), Lerch (2012, 2022), Weihs et al. (2016), and Knees and Schedl (2016). Moreover, many articles published in the ISMIR proceedings and other venues often meet a high educational standard, offering material suitable for teaching and learning specific MIR topics. Insightful educational discussions on specific MIR topics are frequently present in introductory chapters of PhD theses published in the MIR field. Furthermore, ISMIR regularly hosts several tutorials at its annual conferences, which are well-attended and warmly received by the ISMIR community.[3] Educational materials are also available in informal blogs or interactive programming notebooks provided on platforms like GitHub.

However, most of these educational resources on MIR research are scattered across platforms and venues, making them difficult to locate. Additionally, some sources do not guarantee a high standard of quality, as they do not employ any review process and so lack feedback mechanisms by MIR experts for quality control. As such, each prospective user must vet them from scratch. Finally, even well-intentioned users face

challenges in properly referencing these materials due to factors such as the lack of academic publication, ambiguous licensing, and unclear authorship and origins.

Addressing many of these issues, a peer-reviewed and academically published educational article holds a distinct place in the educational ecosystem. It is particularly suitable for providing supportive information (e.g., 'what' and 'how to' for non-routine tasks) and procedural information (e.g., 'how to' for routine tasks), while accompanying notebooks or websites can complement these materials for practicing tasks (van Merriënboer, 2019).

1.2 TISMIR AS A VENUE FOR EDUCATION

To meet the growing need for high-quality educational resources in MIR, we have established an education track within the TISMIR journal. This track provides a dedicated platform for disseminating educational content tailored to the needs of students, educators, and researchers in MIR. Our aim is to enhance the accessibility, credibility, and impact of educational materials in the field by offering structured, peer-reviewed articles. Educational articles in TISMIR should systematically introduce fundamental theories, common practices, and/or applications in well-defined areas of interest to readers across various MIR-related fields, whether those fields are mature or emerging.

Examples may include specific signal processing or machine learning techniques, discussions on musicological or cognitive theories, introductions and reviews of recent MIR tasks, conceptual insights, or discussions on specific applications, reflecting the diverse and interdisciplinary interests of the MIR community. Through collaborative efforts among researchers, educators, and practitioners, our education track seeks to advance the pedagogical landscape of MIR and cultivate a community of lifelong learners dedicated to excellence in computational music research.

1.3 ORGANIZATION OF EDITORIAL

In the following sections we reflect on what may constitute a well-crafted educational article. To achieve this we highlight key characteristics of educational articles in general and clarify review criteria (Section 2). Additionally, we explore why the music domain may provide an intuitive and motivating setting for education across various levels and disciplines (Section 3). Lastly, we provide examples of topics, concepts, and applications that could serve as an orientation aid for potential educational TISMIR articles (Section 4). In the concluding Section 5, we present the benefits of dedicating time and effort to a TISMIR education track submission. We wish to emphasize that this editorial is not intended to be comprehensive and should not constrain creativity or restrict the perspectives of potential authors. Instead, our objective is to offer guidance and foster discussions

on crafting impactful educational articles for MIR and beyond.

2 KEY CHARACTERISTICS OF EDUCATIONAL ARTICLES

High-quality educational articles in the sciences typically exhibit several essential components, effectively conveying complex scientific information and promoting understanding and learning among readers. In this section, we review some of these characteristics, which are applicable to educational articles across various domains. Section 4 will further explore concrete examples within the context of MIR.

2.1 SCOPE

The objective of a TISMIR educational article is to introduce an MIR-related subject in a clear and accessible way, accommodating a broad and interdisciplinary readership. An educational paper should systematically introduce fundamental theories, common practices, and applications in a well-defined, mature or emerging area, and preferably be of interest to readers across various MIR-related fields. A practical approach to defining the scope of an educational paper is to establish learning objectives, which are declarative statements that specify what learners are expected to know and do (Orr et al. 2022).

It is crucial to distinguish educational papers from conventional research articles, which typically present innovative techniques and novel findings. Educational articles should refrain from introducing new research or techniques, instead offering fresh perspectives on existing concepts and methodologies. In other words, educational articles should prioritize enhancing understanding rather than highlighting research innovations. Articles will succeed in making a novel contribution, not in the introduction of new techniques (for example), but in presenting a previously unavailable *educational perspective* on the subject matter. As a result, such an article may serve as a basis for a lecture in, e.g., information retrieval, machine learning, or signal processing. This perspective may also involve presenting a more STEM-oriented introduction to a subject usually taught in the humanities or social sciences. Conversely, it could involve examining the applicability of a specific engineering technique in addressing particular problems or needs within the humanities.

2.2 COMPREHENSIBILITY/CLARITY

Across various domains, educators would likely agree that effective educational articles should be accurate, clear, engaging, and relevant, regardless of the subject matter (Bloom and Engelhart, 1956; Krathwohl, 2002; Starr et al., 2008). A wishlist of key characteristics may include the following aspects:

- *Accuracy.* An educational article should contain *accurate* and *current* scientific information, ensuring that readers can trust the information provided and rely on it for learning purposes.
- *Clarity.* An educational article should provide *clear* and *comprehensible* explanations of scientific concepts, theories, and methodologies in a manner that is easily understandable.
- *Jargon-Free.* An educational article should be free of unnecessary jargon, and explain technical terms for readers with limited background knowledge.
- *Links to Further Reading.* While it may not always be possible to explain every concept from scratch, an educational article should provide clear links to external resources explaining any assumed knowledge. These links should be carefully vetted, open access if possible, and adopt the same educational standards outlined here.
- *Engaging Presentation.* An educational article should engage readers through an interesting presentation, including examples, visuals like diagrams or graphs, and interactive elements where possible. Real-world examples and analogies should be used to simplify complex concepts and improve understanding. If suitable, incorporating music examples in digital symbolic representations or as audio recordings can also be beneficial.
- *Practical and Relevant.* An educational article should demonstrate practical relevance by discussing how scientific concepts apply to concrete and real-world MIR applications, ensuring readers perceive the value of the information presented. Furthermore, it may illustrate how techniques and abstract concepts can be applied and scaled to solve problems in other domains beyond music.

Effective pedagogical strategies in educational articles, such as presenting the material in small steps, providing worked examples, summarizing key points, providing exercises for reflection, and suggesting further reading, are vital for facilitating learning (Rosenshine, 2012). Moreover, proper references and citations are essential for enhancing the credibility of content and facilitating deeper exploration of the topic for readers. They support claims made within the article and offer additional resources for further understanding and research. These are only a few key characteristics and suggestions. Additional insights are offered by Heard (2022) regarding writing high-quality scientific educational articles.

2.3 INTENDED AUDIENCE AND USABILITY
An educational article should clearly state its purpose, audience, and how it can be used. This helps tailor the content and language accordingly. Clarifying its potential application in lectures, seminars, or practical settings also aids readers in understanding how to engage with the information. This ensures the article effectively serves its purpose and maximizes its impact on the intended audience.

The audience of an educational article can vary based on the subject matter and its intention. Generally, the audience comprises individuals interested in learning about a specific topic, including students, researchers, educators, and professionals. Educational articles are typically written to be accessible to readers with diverse levels of expertise, making them suitable for a broad audience.

When preparing an educational article for TISMIR, it may be helpful to envision it as providing a robust foundation for lectures, seminars, or lab sessions, targeting researchers new to the field, students, and teachers. For example, it could serve as comprehensive notes for a 60–90 minute lecture. Additionally, it should delve into specific techniques to assist researchers or practitioners in applying them to their respective research problems. Furthermore, when exploring an MIR task, it should encourage readers from diverse disciplines to explore potential applications within the music domain.

2.4 RESOURCE-RELATED PRINCIPLES
The TISMIR education track appreciates the FAIR principles of *Findability, Accessibility, Interoperability*, and *Reusability* (Wilkinson et al., 2016) and similar resource-related principles such as *Availability* and *Reproducibility*, and expects adoption of the same in submissions to the track.

Some of these principles are discussed by McFee et al. (2019) in the context of music signal processing research, and here we explicitly acknowledge the importance of these issues not only for research but also for educational articles. *Availability* ensures that the article is accessible to its intended audience, whether through open access publishing or through widely available platforms. *Reproducibility* refers to the ability to replicate the experiments or methods described in the article, which is crucial for validating the findings and ensuring that others can build upon the research. These aspects contribute to the credibility and usefulness of the educational article, allowing readers to access and verify the information presented.

In view of applicability and reproducibility, an educational article might offer links to test datasets, present a concise reference implementation of fundamental techniques, or incorporate supplementary multimedia materials like audio, images, sheet music, and symbolic data, if relevant. In this context, it is crucial that both the article and any supplementary material are *open source*, ensuring transparency, reproducibility, and trust. This fosters innovation and advancement in education, promoting equity and collaboration among researchers. Overall, open source and open access significantly enhance the dissemination, accessibility,

and impact of educational materials, benefiting academia and society alike.

Connecting many of the above considerations is that of *license*. It is essential (and required of a TISMIR education track submission) that the associated code is not only open in a general sense but specifies the exact nature of that openness with a clear and explicit license. Anyone seeking to make secondary use of the materials needs to have the confidence that it is permitted and to know if they have to acknowledge the source (and how to do so). The same goes for any data involved in those materials. For example, the opacity of license has held symbolic MIR back (Gotham, 2021).

2.5 REVIEW PROCESS

Every educational article submitted to TISMIR will be vetted by the standard peer-review process, similar to any other TISMIR submission. However, different review criteria are employed compared to other article types. These criteria consider the key characteristics discussed before and include:

- Quality of language and presentation style.
- Emphasis on deepening understanding rather than research novelty.
- Introduction of a novel educational viewpoint that was previously unavailable.
- Appeal to a wide audience or to a well-defined target audience.
- Ease of reproducibility.
- Inclusion of instructive graphics, figures, and other multimedia materials.
- Consistency in mathematical and conceptual formalism where these are necessary.

We primarily support submissions that provide resources and educational materials on MIR. While TISMIR is not an educational journal, we remain open to other types of pedagogical papers addressing topics such as curriculum design, modes of assessment, diversity/inclusion, learning analytics, and plagiarism within the context of MIR.

All educational articles undergo a pre-screening by the TISMIR editorial board before being sent for peer review. Submissions may be rejected if they are out-of-scope or poorly written. However, authors of rejected submissions may be encouraged to consider submitting to the regular track if their submission contains novel research results. Authors may also be encouraged to resubmit after a thorough rewrite and restructuring, initiating a new review process.

Finally, it is important to note that the length of an educational submission should adhere to the same rules as regular TISMIR submissions, with a maximum of 8000 words. In exceptional cases where a paper is likely to exceed this limit, authors are encouraged to contact the journal to discuss available options.

3 MUSIC DOMAIN: OPPORTUNITIES FOR EDUCATION

Music is present in all facets of our lives, playing a vital role in fostering creativity. Given its widespread influence, the realm of music offers an ideal and intuitive environment for teaching across a range of disciplines including computer science, engineering, and computational humanities. Leveraging music as a pedagogical tool enables educators to boost student engagement, motivation, and comprehension while also facilitating interdisciplinary connections. In the following sections, we explore these aspects in greater depth.

3.1 ENGAGEMENT AND MOTIVATION

Many people have a personal connection to music and are inherently motivated to learn about it. Using music as a context for teaching technical concepts can enhance engagement and motivation, leading to deeper learning and understanding. Furthermore, music encourages creativity and innovation, making it an exciting domain for exploring new ideas and approaches. For example, within the realm of generative music, interactive installations, or intelligent music composition systems, one can delve into the study of machine learning and signal processing techniques. Finally, music can be interactive, allowing users to actively engage with the content through activities such as playing instruments, composing, remixing tracks, or participating in live performances. Interactive music technologies enable users to manipulate and control musical elements in real-time, enhancing engagement and creativity.

3.2 INTERDISCIPLINARY NATURE

The interdisciplinary nature of music and music processing spans a broad spectrum of fields and disciplines that intersect with the study and application of music. Below, we offer a selection of interdisciplinary connections to fields closely related to MIR.

- *Music Theory*. This field involves the study of the principles and structures underlying musical composition, which can include harmony, melody, rhythm, and form (Eerola, 2024). Thus, music theory provides the theoretical framework and conceptual basis for many MIR aspects, including analysis, feature extraction, retrieval, transcription, and generation of music.
- *Ethnomusicology*. Examining its influence on identity, community, and society, ethnomusicology explores the cultural, social, and historical dimensions of music (see, e.g., Rao et al., 2023; Duan et al., 2023; Tzanetakis, 2014). Drawing insights from this field, MIR can aid in tasks such as collecting and annotating music collections, developing inclusive and culturally diverse music representations, and

facilitating cross-cultural comparisons of musical practices and traditions.

- *Psychology.* The psychology of music explores how music affects human cognition, emotion, perception, and behaviour (Deutschm, 2013). It encompasses areas such as music cognition (Honing, 2021), psychoacoustics (Fastl and Zwicker, 2007), user experience, and music therapy (Agres et al., 2021).
- *Acoustics and Audio Engineering.* Acoustics is the study of sound and its transmission, while audio engineering focuses on the recording, processing, and reproduction of sound. These fields are essential for understanding the physical properties of musical sound and for developing audio technologies and equipment (see Fletcher and Rossing, 1998, for an example reference in this field).
- *Computer Science.* Serving as the technical cornerstone and providing essential tools for the efficient and effective analysis, organization, and retrieval of music information, computer science stands as a core discipline within MIR. Key subfields include machine learning, pattern analysis, information retrieval, data representation, and algorithm development.
- *Engineering.* Techniques particularly in the realm of digital signal processing (DSP) play a crucial role in the analysis, synthesis, and manipulation of music recordings. Leveraging methods such as the Fourier transform, filter design, and time–frequency decompositions, DSP facilitates tasks like spectral analysis, feature extraction, filtering, noise reduction, and audio effects processing (see, e.g., McFee, 2023; Zölzer, 2002).
- *Music Technology.* This field involves the design and implementation of hardware and software tools for creating, performing, and recording music. It encompasses areas such as electronic music instruments, digital audio workstations (DAWs), music software applications, and interactive music systems.
- *Music Education.* Music education involves teaching the theory, performance, history, and technology of music across various cultural and academic contexts. MIR can learn from music education by incorporating pedagogical strategies and user experience insights to develop tools that are pedagogically relevant and easy to use. Meanwhile, music education can benefit from these advanced technologies to enhance virtual digital curricula and other educational materials.

Overall, the interdisciplinary nature of music reflects its multifaceted role in human culture, society, and technology, enabling us to draw on insights and methodologies from diverse fields to advance our understanding and appreciation of music.

3.3 DATA MULTIMODALITY AND RICHNESS

Music, being a highly multimodal and rich multimedia domain, engages various modalities including auditory, visual, tactile, emotional, and cultural aspects (see, e.g., Li et al., 2019; Essid and Richard, 2012; Müller et al., 2012). These modalities are closely related to MIR in the following ways:

- *Auditory Modality.* Music primarily engages the auditory sense, conveying information through sound waves, including pitch, rhythm, melody, harmony, and timbre.
- *Visual Modality.* Music data often includes visual elements such as musical scores, sheet music, graphical representations of sound (e.g., spectrograms), and music videos. Visual elements provide additional information about the structure, dynamics, and performance of music and vice versa.
- *Tactile Modality.* Musicians and listeners often experience tactile sensations while playing instruments, feeling vibrations, or tapping along with the rhythm. Additionally, tactile feedback can be incorporated into music interfaces or haptic devices for immersive musical experiences.[4]
- *Emotional Modality.* Music evokes emotions and expresses feelings through its tonal and rhythmic elements. Although emotional responses to music can vary widely among individuals and cultures, music remains a powerful medium for conveying and eliciting emotions (Cañón et al., 2021; Yang and Chen, 2011).
- *Cultural Modality.* Music reflects cultural aspects of society, including language, traditions, rituals, and social norms. Different musical styles, genres, and practices emerge from diverse cultural backgrounds, contributing to the richness and diversity of the musical experience.

3.4 REAL-WORLD APPLICATIONS

In recent years, there has been a significant increase in commercial breakthroughs through MIR technology. These range from early examples like audio fingerprinting (Wang, 2003), as used by *Shazam*, to a wide range of MIR technologies employed by streaming services and other digital service providers. Real-world applications are instrumental in advancing the field of MIR by offering opportunities for validation, feedback, innovation, and collaboration, thus driving research and education by addressing practical challenges. In the following, we outline some key application areas without attempting to be exhaustive.

- *Music Recommendation Systems*. MIR techniques are used to develop music recommendation systems that suggest personalized playlists, albums, or tracks to users based on their listening history, preferences,

and context (Knees and Schedl, 2016). Originally explored within the realm of digital music libraries, these systems have now gained broad usage across streaming platforms, online music stores, and personalized radio services, enhancing user satisfaction and engagement.

- *Music Classification and Tagging.* MIR algorithms are employed for automatically classifying and tagging music according to genre, mood, tempo, instrumentation, and other musical attributes (Nam et al., 2019). This facilitates music organization, search, and discovery, enabling users to find relevant content more efficiently.

- *Music Transcription and Score Following.* MIR research addresses the challenges of automatically transcribing audio recordings into symbolic representations, such as musical scores or MIDI files (Benetos et al., 2019). Score following techniques enable computers to synchronize with live performances and generate real-time annotations, facilitating applications in music education, performance analysis, and accompaniment systems (Dannenberg and Raphael, 2006).

- *Music Analysis and Visualization.* MIR methods enable the analysis and visualization of musical content in various forms, including audio waveforms, spectrograms, chord sequences, and structural annotations. These tools aid music scholars, composers, and performers in studying music structure, style, and expression, as well as in creating visualizations for educational and artistic purposes (Goto and Dannenberg, 2019).

- *Music Generation and Composition.* MIR techniques contribute to the development of algorithms and systems for generating new musical compositions, improvisations, or arrangements. These systems can be used for automated music composition, collaborative music-making, and creative exploration, inspiring new artistic expressions and fostering musical innovation.

- *Music Copyright and Intellectual Property Management.* MIR methods play a role in copyright enforcement, content identification, and intellectual property management in the music industry (Yesiler et al., 2021). Techniques such as audio fingerprinting and watermarking are used to detect unauthorized use of copyrighted music and protect the rights of content creators and distributors.

- *Digital Humanities.* The applications of MIR in digital humanities are multifaceted and significant. For example, MIR technology is increasingly used for the preservation of music-related cultural heritage and facilitates the analysis and comprehension of musical compositions and performances (see, e.g., van Kranenburg et al., 2019; de Valk et al., 2017; Rosenzweig et al., 2022). Furthermore, MIR

techniques are utilized in corpus-driven research to analyze extensive music collections, unveiling patterns, genres, styles, and temporal trends (see, e.g., Volk and de Haas, 2013; Weiß et al., 2019; Arthur, 2021; Henry et al., 2024).

- *Music Therapy and Healthcare.* MIR research contributes to applications in music therapy and healthcare, where music is used as a therapeutic intervention for physical, emotional, and cognitive rehabilitation (Agres et al., 2021). Automated music analysis and generation tools support the development of personalized music therapy interventions tailored to individual needs and preferences.

4 EXAMPLE SCENARIOS

4.1 TEACHING SIGNAL PROCESSING THROUGH MUSIC

As outlined by Müller et al. (2021), music signals offer abundant opportunities to enrich signal processing education, as they evoke intuitive understanding of signal properties that beginners can readily grasp without needing extensive technical knowledge. For instance, on very short time scales, one may want to know the onset position and the fundamental frequency of a played note. On longer time scales, questions may arise about the instrumentation and the chord being played. On even longer time scales, it becomes possible to explore higher-level concepts such as melody and rhythm. Therefore, music can serve as a means to make signal processing education more accessible by employing easily understandable initial questions, which encourage hands-on exploration and experimentation right from the beginning. In particular, music can provide tangible examples to demystify mathematically intricate concepts essential for comprehending advanced signal processing concepts.

In a typical introductory course on digital signal processing, a broad range of topics is covered, including sampling theory, Fourier analysis, convolution and filtering, and time–frequency representations, just to name a few. In their educational article, Müller et al. (2021) illustrate how music enhances signal processing education by providing practical applications and interactive learning experiences, using Fourier analysis as an example. As shown in Figure 1, musical signals provide intuitive examples that demonstrate the effectiveness of short-time Fourier analysis, unveiling noticeable changes in pitch, volume, and other musically significant attributes in spectrogram images. Similarly, there is value in educational materials that aid in comprehending and teaching more advanced signal processing concepts, such as the widely adopted constant-Q transform (CQT) introduced by Brown (1991), its harmonic extension
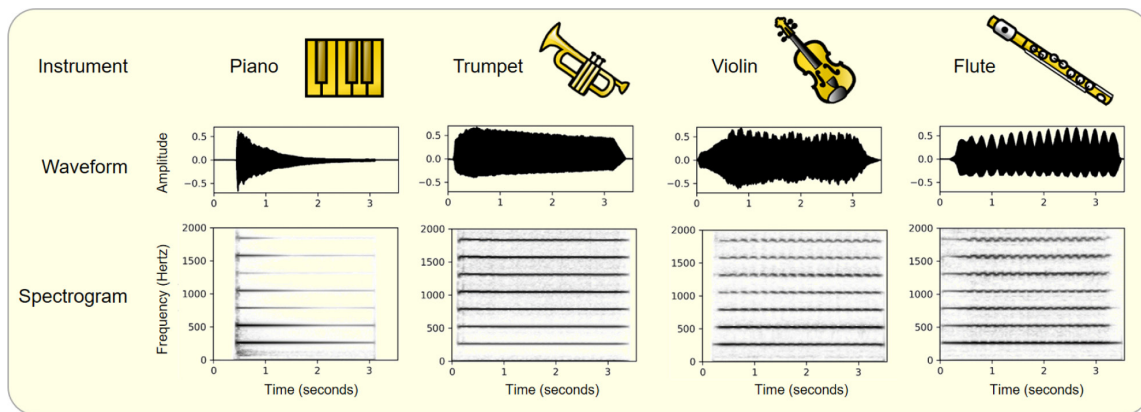
Figure 1 Example figure for an educational article, showing the waveform and spectrogram of various instruments playing the same note C4 with a fundamental frequency of 261.6 Hertz (figure adapted from Müller et al., 2021).

(HCQT) by Bittner et al. (2017), or feature representations derived from deep learning (Kim et al., 2020).

One of the most challenging and fascinating research problems in MIR is automatic music transcription, which seeks to convert an acoustic music signal into a form of musical notation using computational methods. As detailed in the educational article by Benetos et al. (2019), this task encompasses multiple subtasks, including multipitch estimation, onset and offset detection, instrument recognition, beat and rhythm tracking, interpretation of expressive timing and dynamics, and score typesetting. The diverse aspects of music recordings and their acoustic and musical properties provide abundant opportunities to enhance signal processing education. Integrating music into the curriculum makes it more accessible, engaging, and relevant for students with a musical background, thereby enriching their learning journey and deepening their comprehension of the subject.

In addition to providing motivation through tangible music-based scenarios, the availability of well-designed software packages that enhance accessibility to signal processing is indispensable for facilitating interactive learning (Guzdial, 2013). The MIR community has contributed several excellent toolboxes that provide modular source code for processing and analyzing music signals. Notable examples include *essentia* (Bogdanov et al., 2013), *madmom* (Böck et al., 2016), *Marsyas* (Tzanetakis, 2009), *Sonic Visualiser* (Cannam et al., 2006) and the *MIRtoolbox* (Lartillot and Toiviainen, 2007). While these toolboxes are mainly designed for research-oriented access to audio processing, the Python package *librosa* (McFee et al., 2015) and the FMP notebooks (Müller and Zalkow, 2019), alongside its Python package *libfmp* (Müller and Zalkow, 2021), adopt an explicit educational stance. These toolboxes provide interactive elements to a typical lecture-based signal processing classroom, allowing students to solidify their abilities to apply signal processing concepts in musical examples, thus bridging the gap between education and cutting-edge research.

4.2 TEACHING MUSIC THEORY USING COMPUTATION

The *Open Music Theory* textbook (OMT)[5] is a natively-online open educational resource intended to serve as the primary text and workbook for undergraduate music theory curricula. As such, it explicitly targets individuals learning music theory, primarily undergraduates. That said, certain sections of the book utilize computational methods to enrich the study experience, potentially inspiring prospective contributors to the TISMIR education track.

For example, the textbook includes a *Harmony Anthology*[6] chapter. Traditional anthologies may be most familiar in the context of literature as a collection of works that are 'important' or 'typical' of a time and/or place (e.g., *American poetry after the Civil War*). Similar resources exist in music theory teaching, bringing together examples from a particular historical period and/or for a pedagogical purpose. For further discussion and many references, see Gotham (2019).

OMT's Harmony Anthology aims to offer an updated version of this resource type for the modern age, with two main differences. First, it operates on a much larger scale. Traditional anthologies typically provide two or three examples due to practical limitations in printing, binding, and selling. In contrast, the OMT Harmony Anthology gathers hundreds of examples, organized around specific chords such as the *Neapolitan Sixth*. Second, and relatedly, this larger scale shifts the focus away from a few 'exemplars' toward a broader collection, encouraging free exploration.

Exploring this collection underscores the subjective nature of analysis. Therefore, individual readers and analysts (including students) should anticipate disagreements and benefit from open debate. Readers can form their own opinions about whether a specific instance 'counts' as an example of the category. This effective method of building meta-cognition is not available in resources that only use clear-cut, unambiguous examples (see also Gotham, 2021).

Under the hood, the Harmony Anthology leverages a great deal of MIR work, both in the preparation of the data and in the extent of links serving that goal of free exploration. First, all the scores in the anthology (and many throughout the textbook) result specifically from open source dataset creation efforts that explicitly aim to service both MIR research and public-facing resources (as reported, for example, by Gotham and Jonas, 2022)). The fact that these scores are released under a CC0 (Creative Commons Zero) license, allowing anyone to use, adapt, and distribute the work without restrictions, is highly significant in terms of the FAIR principles discussed above. Second, the analyses are likewise the result of MIR dataset curation, this time from the meta-corpus 'When in Rome' (Gotham et al., 2023). Third, the anthology hosts numerous external links, to the score online (for immediate playback online), score on the Github repository (for download without any need for login), and to the extracted example as a small image, which can be browsed with a pop-up window (no re-direct or download required).

Finally, MIR scholars know very well that, even putting aside the subjectivity involved in an analyst making a commitment to a specific chord, the definition and categorization of these chords for anthology purposes bring a whole additional layer of ambiguity. For example, the term 'modal mixture' has been core to American music theory teaching for some time now but never received a robust definition until Gotham (2023). This anthology explicitly implements that algorithm and sheds light on the criteria for selection— again, this is an effective way to build students' meta-cognition.

At this stage, the connection to current research risks becoming too strong and distant from the educational criteria set out above. On the other hand, the absence of a robust definition in existing educational literature demonstrates the need for *something*, and the openness of the approach (complete with open-source code for the implementation available back at 'When in Rome') provides the kind of FAIR principles on which we place such a high premium.

4.3 TEACHING MATHEMATICS AND COMPUTER SCIENCE THROUGH MUSIC

'Music and mathematics' is perhaps the oldest form of an explicit connection between music and other disciplines. This connection traces its origins back through the medieval *quadrivium*, a four-part curriculum comprising arithmetic, geometry, music, and astronomy, all the way to antiquity. Moreover, it continues to be revived and modernized. Noteworthy for our present purposes is the special issue on 'Pedagogies of Mathematical Music Theory' in the *Journal of Mathematics and Music*, where the concept of teaching mathematics through music is explicitly discussed (Yust and Fiore, 2014).

Computer science, despite its shorter history, has also found value in leveraging the enduring connection between mathematics and music in its own distinct ways. Notable examples of using music to instruct basic coding include *SonicPi*[7]: a code library and GUI for live-coding music. SonicPi explicitly prioritizes the teaching aspect, influencing every decision, down to the naming of attributes and methods. For an example of how music may help teach basic coding, consider the 'for-loop'. While text and visual representations can be effective, arguably nothing captivates as much as musical data. This is particularly true considering our inherent affinity for repetition in music (Margulis, 2014).

A strong connection between MIR and computer science is evident with the availability of large datasets and computational resources. In this context, machine learning with deep architectures stands out as an advanced approach within MIR (see, e.g. Humphrey et al., 2013; Pons et al., 2018). In particular, the 'end-to-end' pipeline represents a significant departure from the established approach of engineering features and then training models. Instead, it integrates these two steps, optimizing features and models together, resulting in MIR systems that seem to perform very well. However, a considerable challenge emerges when it comes to evaluating the resulting 'black box' in meaningful ways. This task involves dissecting the system and striving to interpret the functions of its components. An educational article in TISMIR could offer valuable insights by presenting and discussing methods to uncover the workings of the black box.

Expanding on this, while much research focuses on explainable AI, particularly in fields involving text and image data, the music domain offers a unique perspective due to its complexity and multimodality, along with potential confounding factors arising from various versions of the same piece and other inherent intricacies in musical data. For instance, Weiß et al. (2020) explore local key estimation in classical music, using a cross-version dataset of Schubert's Winterreise performances. Their study reveals insights into model efficacy influenced by different training–test splits and uncovers correlations between errors and annotator disagreements, as well as musically explainable relationships among different keys. As another example, Cosme-Clifford et al. (2023) explore the behavior and constitution of a transformer model, drawing on significant expertise in music theory. This study highlights how music can enhance comprehension in computational domains, providing insights that traditional methods may overlook.

4.4 TEACHING CROSS-CULTURAL COMPARISON THROUGH COMPUTATION

Cross-cultural music studies play a significant role in broader cross-cultural education, facilitating the appreciation of cultures beyond one's own. Given the

variations in language, behavior, societal norms, and especially teaching methods, delving into the music of a new culture can pose significant challenges. This is where computational methods prove invaluable, acting as a bridge between different cultural artistic forms. By identifying the commonalities that underlie all music traditions, likely stemming from universal human cognitive processes, it becomes possible to establish connections that greatly enhance understanding of culturally distinctive genres. Examples of this approach are readily found in computational musicology, which not only makes genre-specific music concepts more universally accessible through objective descriptions of audio features and visualization but also benefits from human annotations to provide additional context and insight.

For instance, Clayton et al. (2023) provide examples of vocal performance in the north Indian *Dhrupad* style, highlighting its meticulously arranged patterns of interaction between syllabic singing and accompanying drumming. They illustrate how the lyrics of the composition are sung at different speeds (basic, double, triple, and quadruple) across various metric cycles with visual depictions of this phenomenon via the derived pitch contour that exhibits shape compression with precisely the same ratios. Tempograms computed on source-separated audio signals allow for the identification of distinct rhythmic episodes throughout the concert, where the singer and the drummer each subdivide the beat differently but in recognizable ratios associated with the intrinsic musical concept of *laykari* (rhythmic play). This analysis offers insight into the various ways in which a listener familiar with the cultural context may parse and appreciate the music.

Culture-specific attributes of music, such as the precise intonation of scale intervals, can be effectively demonstrated through computational analyses like continuous pitch contours, with platforms like *Music in Motion*[8] offering valuable educational resources. Future efforts should focus on simplifying and adapting visualizations generated by MIR to better meet the educational needs of the non-MIR music community. Another important objective is to create educational materials and interfaces to assist musicologists in leveraging MIR to complement traditional approaches to music study. Lastly, as mentioned in Section 4.2, MIR methods have the potential to enhance music pedagogy across cultures by contributing to the development of robust definitions of music concepts. This can play a significant role in cross-cultural educational contexts and foster further research in music across different cultures.

5 CONCLUSIONS

In this editorial, we introduced the new education track for TISMIR, aiming to provide guidance for prospective authors and readers regarding its context, goals, and scope. Rather than prescribing conclusive, one-size-fits-all solutions, this editorial serves as an invitation for initial reflections on what may constitute a well-crafted educational article. Our objective is not to impose limitations but to offer guidance and stimulate discussion on crafting impactful educational articles for MIR and beyond.

We strongly believe that the education track presents numerous benefits and opportunities for both authors and readers. Here are compelling reasons why one should consider writing an educational article and specifically why exploring education in the field of MIR is valuable:

- Firstly and most clearly, educational articles serve as invaluable resources for students, educators, and researchers, supporting both formal education and self-directed learning initiatives. By incorporating or linking to visual and acoustic materials, music examples, code samples, clear explanations, and exercises, these articles promote active learning and engagement among readers.

- Secondly, contributing to the education track allows authors to actively advance the MIR field by sharing their expertise, methodologies, and insights with the broader community. By disseminating knowledge and best practices, authors enrich the collective knowledge base and foster innovation and collaboration within the various disciplines encompassed by the MIR field. The interdisciplinary nature of MIR provides authors with a diverse range of topics and methodologies to explore, making it an exciting and dynamic domain for educational endeavors.

- Thirdly, writing an educational article fulfills a scholarly duty to disseminate knowledge, bridging research and teaching. It provides a platform to share findings and methodologies with a wider audience, ensuring that the research has a meaningful impact beyond academic circles.

- Last but not least, writing an educational article can be a valuable experience for researchers and PhD students. For instance, starting with a tutorial presented at an ISMIR conference, transforming tutorial materials into an educational article offers an opportunity to refine these resources into a sustainable format. Additionally, writing an educational article allows PhD students to reflect on general principles and fundamental techniques pertinent to their work. Such an article can also serve as a solid foundation for the introductory part of a PhD thesis. In this context, the review process provides crucial feedback from MIR experts, helping researchers and students enhance their work for quality and relevance in the field.

In conclusion, contributing to the TISMIR education track not only allows authors to explore the interdisciplinary nature of the field but also provides unique opportunities for leveraging music as a pedagogical tool. By creating engaging, multidisciplinary educational materials, authors can inspire learners to explore new ideas, develop new skills, and make meaningful contributions to the advancement of knowledge in MIR and beyond.

NOTES

1 https://scholar.google.com/citations?hl=en&view_op=search_venues&vq=music&btnG=.
2 For a list of these, see https://www.ismir.net/resources/research-centers.
3 For a list of these, see https://www.ismir.net/resources/tutorials.
4 For new musical interfaces see also the proceedings of the NIME conference, https://nime.org/.
5 https://viva.pressbooks.pub/openmusictheory/.
6 https://viva.pressbooks.pub/openmusictheory/chapter/anthology-harmony/.
7 https://sonic-pi.net/tutorial.html.
8 https://autrimncpa.wordpress.com/about/.

ACKNOWLEDGEMENTS

COMPETING INTERESTS

The authors have no competing interests to declare.

AUTHOR AFFILIATIONS

Meinard Müller orcid.org/0000-0001-6062-7524
International Audio Laboratories Erlangen, Germany

Simon Dixon orcid.org/0000-0002-6098-481X
Queen Mary University of London, UK

Anja Volk orcid.org/0000-0002-6755-9592
Utrecht University, Netherlands

Bob L. T. Sturm orcid.org/0000-0003-2549-6367
Royal Institute of Technology KTH, Sweden

Preeti Rao orcid.org/0000-0002-2842-8450
Indian Institute of Technology Bombay, India

Mark Gotham orcid.org/0000-0003-0722-3074
Durham University, UK

REFERENCES

Agres, K. R., Schaefer, R. S., Volk, A., van Hooren, S., Holzapfel, A., Bella, S. D., Müller, M., de Witte, M., Herremans, D., Melendez, R. R., Neerincx, M., Ruiz, S., Meredith, D., Dimitriadis, T., and **Magee, W. L.** (2021). Music, computing, and health: A roadmap for the current and future roles of music technology for health care and well-being. *Music & Science*, 4:1–32. DOI: https://doi.org/10.1177/2059204321997709

Arthur, C. (2021). Vicentino versus Palestrina: A computational investigation of voice leading across changing vocal densities. *Journal of New Music Research*, 50(1):74–101. DOI: https://doi.org/10.1080/09298215.2021.1877729

Benetos, E., Dixon, S., Duan, Z., and **Ewert, S.** (2019). Automatic music transcription: An overview. *IEEE Signal Processing Magazine*, 36(1):20–30. DOI: https://doi.org/10.1109/MSP.2018.2869928

Bittner, R. M., McFee, B., Salamon, J., Li, P., and **Bello, J. P.** (2017). Deep salience representations for F0 tracking in polyphonic music. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 63–70, Suzhou, China.

Bloom, B. S. and **Engelhart, M. D.** (1956). *Taxonomy of Educational Objectives: The Classification of Educational Goals. Handbook I: Cognitive Domain*. David McKay Company.

Böck, S., Korzeniowski, F., Schlüter, J., Krebs, F., and **Widmer, G.** (2016). madmom: A new Python audio and music signal processing library. In *Proceedings of the ACM International Conference on Multimedia (ACM-MM)*, pages 1174–1178, Amsterdam, The Netherlands. DOI: https://doi.org/10.1145/2964284.2973795

Bogdanov, D., Wack, N., Gómez, E., Gulati, S., Herrera, P., Mayor, O., Roma, G., Salamon, J., Zapata, J. R., and **Serra, X.** (2013). Essentia: An audio analysis library for music information retrieval. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 493–498, Curitiba, Brazil. DOI: https://doi.org/10.1145/2502081.2502229

Brown, J. (1991). Calculation of a constant q spectral transform. *Journal of the Acoustical Society of America*, 89(1):425–434. DOI: https://doi.org/10.1121/1.400476

Cannam, C., Landone, C., Sandler, M. B., and **Bello, J. P.** (2006). The Sonic Visualiser: A visualisation platform for semantic descriptors from musical signals. In *Proceedings of the International Conference on Music Information Retrieval (ISMIR)*, pages 324–327, Victoria, Canada.

Cañón, J. S. G., Cano, E., Eerola, T., Herrera, P., Hu, X., Yang, Y., and **Gómez, E.** (2021). Music emotion recognition: Toward new, robust standards in personalized and

context-sensitive applications. *IEEE Signal Processing Magazine*, 38(6):106–114. DOI: https://doi.org/10.1109/MSP.2021.3106232

Clayton, M., Rao, P., and **Rohit, M. A.** (2023). Rhythm and structural segmentation in Dhrupad Bandish performance. In *Indian Art Music: A Computational Perspective*, pages 215–238. Sriranga Digital Software Technologies, India.

Cosme-Clifford, N., Symons, J., Kapoor, K., and **White, C.W.** (2023). Musicological interpretability in generative transformers. In *Proceedings of the International Symposium on the Internet of Sounds*, pages 1–9. DOI: https://doi.org/10.1109/IEEECONF59510.2023.10335202

Dannenberg, R. B. and **Raphael, C.** (2006). Music score alignment and computer accompaniment. *Communications of the ACM, Special Issue: Music Information Retrieval*, 49(8):38–43. DOI: https://doi.org/10.1145/1145287.1145311

de Valk, R., Volk, A., Holzapfel, A., Pikrakis, A., Kroher, N., and **Six, J.** (2017). MIRchiving: Challenges and opportunities of connecting MIR research and digital music archives. In *Proceedings of the International Workshop on Digital Libraries for Musicology (DLfM)*, pages 25–28. DOI: https://doi.org/10.1145/3144749.3144755

Deutsch, D. (2013). *The Psychology of Music*. Academic Press, 3rd edition.

Dixon, S., Gómez, E., and **Volk, A.** (2018). Editorial: Introducing the Transactions of the International Society for Music Information Retrieval. *Transactions of the International Society for Music Information Retrieval (TISMIR)*, 1(1):1–3. DOI: https://doi.org/10.5334/tismir.22

Duan, Z., van Kranenburg, P., Nam, J., and **Rao, P.** (2023). Editorial for TISMIR special collection: Cultural diversity in MIR research. *Transactions of the International Society for Music Information Retrieval (TISMIR)*, 6(1):203–205. DOI: https://doi.org/10.5334/tismir.179

Eerola, T. (2024). *Music and Science – Guide to Empirical Music Research*. SEMPRE Studies in the Psychology of Music. Routledge, London, UK.

Essid, S. and **Richard, G.** (2012). Fusion of multimodal information in music content analysis. *Multimodal Music Processing (Dagstuhl Seminar 11041), Dagstuhl Follow-Ups*, 3:37–52.

Fastl, H. and **Zwicker, E.** (2007). *Psychoacoustics: Facts and Models*. Springer, 3rd edition. DOI: https://doi.org/10.1007/978-3-540-68888-4

Fletcher, N. H. and **Rossing, T. D.** (1998). *The Physics of Musical Instruments*. Springer, Berlin, Germany, 2nd edition. DOI: https://doi.org/10.1007/978-0-387-21603-4

Gotham, M. R. H. (2019). Moments Musicaux. In *Proceedings of the International Workshop on Digital Libraries for Musicology (DLfM)*, pages 70–78, New York, NY, USA. DOI: https://doi.org/10.1145/3358664.3358676

Gotham, M. R. H. (2021). Connecting the dots: Recognizing and implementing more kinds of "Open Science" to connect musicians and musicologists. *Empirical Musicology Review*, 16:34–46. DOI: https://doi.org/10.18061/emr.v16i1.7644

Gotham, M. R. H. (2023). Chromatic chords in theory and practice. *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 272–278.

Gotham, M. R. H. and **Jonas, P.** (2022). The Open-Score Lieder Corpus. In Münnich, S. and Rizo, D., editors, *Proceedings of the Music Encoding Conference*, pages 131–136. Humanities Commons.

Gotham, M. R. H., Micchi, G., Nápoles-López, N., and **Sailor, M.** (2023). When in Rome: A meta-corpus of functional harmony. *Transactions of the International Society for Music Information Retrieval (TISMIR)*, 6(1):150–166. DOI: https://doi.org/10.5334/tismir.165

Goto, M. and **Dannenberg, R. B.** (2019). Music interfaces based on automatic music signal analysis: New ways to create and listen to music. *IEEE Signal Processing Magazine*, 36(1):74–81. DOI: https://doi.org/10.1109/MSP.2018.2874360

Guzdial, M. (2013). Exploring hypotheses about media computation. In *International Computing Education Research Conference (ICER)*, pages 19–26, La Jolla, CA, USA. DOI: https://doi.org/10.1145/2493394.2493397

Heard, S. B. (2022). *The Scientist's Guide to Writing: How toWrite More Easily and Effectively Throughout Your Scientific Career*. Princeton University Press, 2nd edition.

Henry, L., Frieler, K., Solis, G., Pfleiderer, M., Dixon, S., Höger, F., Weyde, T., and **Crayencour, H.-C.** (2024). Dig that lick: Exploring patterns in jazz with computational methods. *Jazzforschung / Jazz Research*, 50/51. In press.

Honing, H. (2021). *Music Cognition: The Basics*. Routledge. DOI: https://doi.org/10.4324/9781003158301

Humphrey, E., Bello, J. P., and **LeCun, Y.** (2013). Feature learning and deep architectures: New directions for music informatics. *Journal of Intelligent Information Systems*, 41(3):461–481. DOI: https://doi.org/10.1007/s10844-013-0248-5

Kim, J., Urbano, J., Liem, C. C. S., and **Hanjalic, A.** (2020). One deep music representation to rule them all? A comparative analysis of different representation learning strategies. *Neural Computing and Applications*, 32:1067–1093. DOI: https://doi.org/10.1007/s00521-019-04076-1

Knees, P. and **Schedl, M.** (2016). *Music Similarity and Retrieval*. Springer Verlag. DOI: https://doi.org/10.1007/978-3-662-49722-7

Krathwohl, D. R. (2002). A revision of Bloom's taxonomy: An overview. *Theory into Practice*, 41(4):212–218. DOI: https://doi.org/10.1207/s15430421tip4104_2

Lartillot, O. and **Toiviainen, P.** (2007). MIR in MATLAB (II): A toolbox for musical feature extraction from audio. In *Proceedings of the International Conference on Music Information Retrieval (ISMIR)*, pages 127–130, Vienna, Austria.

Lerch, A. (2012). *An Introduction to Audio Content Analysis: Applications in Signal Processing and Music Informatics*. John Wiley & Sons. DOI: https://doi.org/10.1002/9781118393550

Lerch, A. (2022). *An Introduction to Audio Content Analysis*. Wiley/IEEE Press, 2nd edition. DOI: https://doi.org/10.1002/9781119890980

Li, B., Liu, X., Dinesh, K., Duan, Z., and **Sharma, G.** (2019). Creating a multitrack classical music performance dataset for multimodal music analysis: Challenges, insights, and applications. *IEEE Transactions on Multimedia*, 21(2):522–535. DOI: https://doi.org/10.1109/TMM.2018.2856090

Margulis, E. H. (2014). *On Repeat: How Music Plays the Mind*. Oxford University Press, New York, NY. DOI: https://doi.org/10.1093/acprof:oso/9780199990825.001.0001

McFee, B. (2023). *Digital Signals Theory*. Chapman and Hall/CRC. DOI: https://doi.org/10.1201/9781003264859

McFee, B., Kim, J. W., Cartwright, M., Salamon, J., Bittner, R. M., and **Bello, J. P.** (2019). Open-source practices for music signal processing research: Recommendations for transparent, sustainable, and reproducible audio research. *IEEE Signal Processing Magazine*, 36(1):128–137. DOI: https://doi.org/10.1109/MSP.2018.2875349

McFee, B., Raffel, C., Liang, D., Ellis, D. P., McVicar, M., Battenberg, E., and **Nieto, O.** (2015). Librosa: Audio and music signal analysis in Python. In *Proceedings of the Python Science Conference*, pages 18–25, Austin, Texas, USA. DOI: https://doi.org/10.25080/Majora-7b98e3ed-003

Müller, M. (2007). *Information Retrieval for Music and Motion*. Springer Verlag. DOI: https://doi.org/10.1007/978-3-540-74048-3

Müller, M. (2015). *Fundamentals of Music Processing – Audio, Analysis, Algorithms, Applications*. Springer Verlag. DOI: https://doi.org/10.1007/978-3-319-21945-5

Müller, M. (2021). *Fundamentals of Music Processing – Using Python and Jupyter Notebooks*. Springer Verlag, 2nd edition. DOI: https://doi.org/10.1007/978-3-030-69808-9

Müller, M., Goto, M., and **Schedl, M.,** editors (2012). *Multimodal Music Processing*, volume 3 of *Dagstuhl Follow-Ups*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany.

Müller, M., McFee, B., and **Kinnaird, K.** (2021). Interactive learning of signal processing through music: Making Fourier analysis concrete for students. *IEEE Signal Processing Magazine*, 38(3):73–84. DOI: https://doi.org/10.1109/MSP.2021.3052181

Müller, M. and **Zalkow, F.** (2019). FMP Notebooks: Educational material for teaching and learning fundamentals of music processing. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 573–580, Delft, The Netherlands.

Müller, M. and **Zalkow, F.** (2021). libfmp: A Python package for fundamentals of music processing. *Journal of Open Source Software (JOSS)*, 6(63):3326:1–5. DOI: https://doi.org/10.21105/joss.03326

Nam, J., Choi, K., Lee, J., Chou, S., and **Yang, Y.** (2019). Deep learning for audio-based music classification and tagging: Teaching computers to distinguish rock from Bach. *IEEE Signal Processing Magazine*, 36(1):41–51. DOI: https://doi.org/10.1109/MSP.2018.2874383

Orr, R. B., Csikari, M. M., Freeman, S., and **Rodriguez, M. C.** (2022). Writing and using learning objectives. *CBE—Life Sciences Education*, 21(3):fe3,1–6. DOI: https://doi.org/10.1187/cbe.22-04-0073

Pons, J., Nieto, O., Prockup, M., Schmidt, E. M., Ehmann, A. F., and **Serra, X.** (2018). End-to-end learning for music audio tagging at scale. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 637–644, Paris, France.

Rao, P., Murthy, H. A., and **Prasanna, S. R. M.,** editors (2023). *Indian Art Music: A Computational Perspective*. Sriranga Digital Software Technologies, India.

Rosenshine, B. (2012). Principles of instruction: Research-based strategies that all teachers should know. *American Educator*, 36(1):12–39.

Rosenzweig, S., Scherbaum, F., and **Müller, M.** (2022). Computer-assisted analysis of field recordings: A case study of Georgian funeral songs. *ACM Journal on Computing and Cultural Heritage (JOCCH)*. DOI: https://doi.org/10.1145/3551645

Starr, C., Manaris, B., and **Stalvey, R. H.** (2008). Bloom's taxonomy revisited: Specifying assessable learning objectives in computer science. *SIGCSE Bulletin*, 40(1):261–265. DOI: https://doi.org/10.1145/1352322.1352227

Tzanetakis, G. (2009). Music analysis, retrieval and synthesis of audio signals MARSYAS. In *Proceedings of the ACM International Conference on Multimedia (ACM-MM)*, pages 931–932, Vancouver, British Columbia, Canada. DOI: https://doi.org/10.1145/1631272.1631459

Tzanetakis, G. (2014). Computational ethnomusicology: A music information retrieval perspective. In *Proceedings of the International Computer Music Conference (ICMC)*. Michigan Publishing.

van Kranenburg, P., de Bruin, M., and **Volk, A.** (2019). Documenting a song culture: The Dutch Song Database as a resource for musicological research. *International Journal on Digital Libraries*, 20(1):13–23. DOI: https://doi.org/10.1007/s00799-017-0228-4

van Merriënboer, J. J. (2019). *The Four-Component Instructional Design Model*. Maastricht University, The Netherlands.

Volk, A. and **de Haas, W.** (2013). A corpus-based study on Ragtime syncopation. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 163–168, Curitiba, Brazil.

Wang, A. (2003). An industrial strength audio search algorithm. In *Proceedings of the International Conference on Music Information Retrieval (ISMIR)*, pages 7–13, Baltimore, Maryland, USA.

Weiß, C., Mauch, M., Dixon, S., and **Müller, M.** (2019). Investigating style evolution of Western classical music: A computational approach. *Musicae Scientiae*, 23(4):486–507. DOI: https://doi.org/10.1177/1029864918757595

Weiß, C., Schreiber, H., and **Muller, M.** (2020). Local key estimation in music recordings: A case study across songs, versions, and annotators. *IEEE/ACM Transactions on Audio,*

Speech, and Language Processing, 28:2919–2932. DOI: https://doi.org/10.1109/TASLP.2020.3030485

Weihs, C., Jannach, D., Vatolkin, I., and **Rudolph, G.** (2016). *Music Data Analysis: Foundations and Applications*. CRC Press. DOI: https://doi.org/10.1201/9781315370996

Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L. B., Bourne, P. E., Bouwman, J., Brookes, A. J., Clark, T., Crosas, M., Dillo, I., Dumon, O., Edmunds, S., Evelo, C. T., Finkers, R., Gonzalez-Beltran, A., Gray, A. J. G., Groth, P., Goble, C., Grethe, J. S., Heringa, J., 't Hoen, P. A. C., Hooft, R., Kuhn, T., Kok, R., Kok, J., Lusher, S. J., Martone, M. E., Mons, A., Packer, A. L., Persson, B., Rocca-Serra, P., Roos, M., van Schaik, R., Sansone, S.-A., Schultes, E., Sengstag, T., Slater, T., Strawn, G., Swertz, M. A., Thompson, M., van der Lei, J., van Mulligen, E., Velterop, J., Waagmeester, A., Wittenburg, P., **Wolstencroft, K., Zhao, J.,** and **Mons, B.** (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data*, 3(1). Article number: 160018.

Yang, Y.-H. and **Chen, H. H.** (2011). *Music Emotion Recognition*. CRC Press. DOI: https://doi.org/10.1201/b10731

Yesiler, F., Doras, G., Bittner, R. M., Tralie, C. J., and **Serrà, J.** (2021). Audio-based musical version identification: Elements and challenges. *IEEE Signal Processing Magazine*, 38(6):115–136. DOI: https://doi.org/10.1109/MSP.2021.3105941

Yust, J. and **Fiore, T. M.** (2014). Introduction to the special issue on pedagogies of mathematical music theory. *Journal of Mathematics and Music*, 8(2):113–116. DOI: https://doi.org/10.1080/17459737.2014.951188

Zölzer, U. (2002). *DAFX: Digital Audio Effects*. John Wiley & Sons, New York, NY, USA. DOI: https://doi.org/10.1002/0470846046