

HarmonAI: A Visual Interface for AI-assisted Music Harmonization

Simeon Rau
University of Stuttgart
Stuttgart, Germany

Benedikt Brachtel
Folkwang University of the Arts
Folkwang, Germany

Michael Sedlmair
University of Stuttgart
Stuttgart, Germany




Figure 1: HarmonAI’s main visualizations to explore AI-generated harmonizations: 1) Tabular overview visualizes metric values and time-based progressions, e.g., many harmonizations feature typical chord transitions, often including D, G, and A chords. 2) Piano rolls for melody and chords, enriched with a lead sheet covering the chord name and functionality, and consonance and typicality (background colors). 3) We use coloring based on the circle of fifths, so similar colors sound harmonic.

Abstract

We present HarmonAI, a visual and analytical approach for AI-assisted harmonization in music composition. HarmonAI combines manual music composition with generative AI while keeping control in the user’s hands, who curates and selects from AI-generated suggestions. With a two-step approach, we first generate multiple harmonization candidates for a given melody using a generative model, producing a set of increasingly complex music samples. In the second step, we provide a structured, ranked tabular overview of the set and visualize musical characteristics to explore and investigate samples and identify interesting harmonizations. We assume that using rankings and visualizations based on carefully selected metrics, rather than sequentially listening to all suggestions, will help to identify interesting samples from a large set quickly. We discuss our metric selection and visual design, and illustrate HarmonAI’s utility in a usage scenario.

CCS Concepts

• **Human-centered computing** → **Information visualization; Scenario-based design.**

Keywords

Music Visualization, Visual Analytics, Generative-AI Usage

ACM Reference Format:

Simeon Rau, Benedikt Brachtel, and Michael Sedlmair. 2026. HarmonAI: A Visual Interface for AI-assisted Music Harmonization. In *Proceedings of The 19th International Symposium on Visual Information Communication and Interaction (VINCI '26)*. ACM, New York, NY, USA, 5 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Composing music is a complex task in which composers connect melody, harmony, and instrumentation. Writing harmonies requires expertise in complex music theory, which sets a high bar for innovation and exploration. In contrast, common chord progressions set an easy start, but one might get stuck in the same loop.

As a bridge, artificial intelligence (AI) can offer alternative harmonizations to broaden one’s horizons. But often, controlling generative AI can be hard, and results might miss a composer’s needs. Instead of repeatedly generating for a “good” result [18], existing research [33] proposed a broad sampling approach. Building on that concept, HarmonAI offers extensive material through harmonization samples for exploration and for refining a composer’s choices.

Curating a large set of AI-generated harmonizations by listening to all samples or manually analyzing them is tedious and infeasible. Instead, we provide a visual overview to analyze the set of harmonizations concurrently and avoid trial-and-error selection by visualizing harmonization information, thereby focusing only on more interesting samples. Although visualizations cannot entirely

replace listening, they can guide to interesting samples [34] or even let composers imagine sound (notational audiation [6]).

Therefore, HarmonAI leverages a user-centered, visual analytics approach to explore AI-generated harmonizations of varying complexity, thereby facilitating serendipitous discoveries and fostering creativity [38]. To enable analysis of harmonic intention and features, we carefully select, adjust, and calculate metrics from the literature and our own expertise, and display them visually in a tabular overview [15] with color schemes that resemble harmonic consonance. Our goal is for users to imagine the sample's characteristics/sound using our visualizations, similar to auditory imagery [20] and notational audiation [6]. We illustrate HarmonAI's functionality through a usage scenario, in which we explore and analyze AI-generated harmonizations and report on how users can draw conclusions from our metrics.

In summary, we contribute 1) the design of a visual analytics approach to explore AI-generated harmonizations, 2) a usage scenario illustrating its functionality, and 3) the source code of our web-based prototype (HarmonAI: github.com/snrau/HarmonAI).

2 Related Work

Harmonic assessment and analysis are common in musicology and are often applied to existing compositions [35]. In modern times, research proposes metric calculations [41, 45, 48] to capture the essence of standard harmonic rules and connections [32], such as consonance between notes [3, 9, 16]. In our work, we use these metrics, along with others, to visualize the harmonic capabilities of AI-generated outputs.

Generative artificial intelligence is increasingly used in a composer's workflow by generating melodies, harmonizations, instrumentation, and lyrics in symbolic form, audio, or text [21, 23]. In our work, we focus on the harmonization with symbolic representation. For this, models generate editable accompaniment in note form [5, 17, 19, 42] or chord progressions [7, 44, 47]. In our prototype, we use a pretrained model [47] to generate chord progressions, although our visual approach is model-agnostic.

Visual interfaces improve user interaction and provide visual information in the composition workflow [23, 26]. In particular, visualizations are often used to reveal harmonic structure [4, 30, 31] using color to display metrics [25, 28] and reveal harmonic differences [12, 29]. Visualization is used to make AI models more usable for beginners [1] and experts [27]. Rau et al. [33] proposed an approach to visually analyze large sets of melodies using a similarity-based map commonly used in visual analytics [22]. We take inspiration from this approach and adopt a tabular structure [10, 15] with colors to simultaneously analyze many samples and foster serendipitous, unconventional discovery.

3 Design

As an author team of visualization researchers, hobby musicians, music teachers, and a professional composer, we leveraged our expertise in participatory design [39] and closely followed Design Study Methodology [36]. We designed HarmonAI mainly for experts who want to go beyond standard chord progressions and their comfort zone, but exploring generated standard harmonizations could also benefit beginners.

Workflow and Tasks: We intend composers to use HarmonAI when a melodic theme is already present, but the accompaniment is still developing. A potential workflow using HarmonAI could then look like this (Figure 2): At first, the composer loads the melody and samples numerous chord progressions using generative AI. Using the tabular overview and visualized metrics, the composer selects interesting samples and inspects them in the colored lead sheet. After finding usable chord progressions, the composer exports them for use in other programs, such as digital audio workstations (DAW). With this workflow at hand, we extracted the following user tasks: Sample a large number of different harmonizations at once using generative AI (T1), systematically explore results to gain inspiration and novel ideas (T2), identify typical or uncommon harmonizations (T3), and analyze characteristics and harmonic intention (T4).

To support these tasks, we selected metrics (Section 3.1) and chose visualization layouts and designs (Section 3.2) to facilitate systematic exploration, identify harmonic characteristics, and support detailed analysis. Further, we allow playback at any time to maintain user control and support export to common formats.

3.1 Data & Metrics

Our primary data are generated harmonizations as chord progressions for a given melody. Both are stored in a symbolic representation, such as MusicXML or MIDI. In our prototype, we use a pre-trained generative AI model, Autoharmonizer [47], to generate multiple harmonizations (T1). To enable serendipitous discoveries of out-of-the-ordinary harmonizations, we use different model parameter values to influence result complexity during sampling.

For analysis and comparison of the generated harmonizations (T2, T3, T4), composers might want to gather information about harmonic characteristics, such as the interplay between chords and melody. Therefore, we use the samples' symbolic data to calculate various metrics for the user to investigate. We focus on metrics commonly used in musicology [28], such as Roman numeral chord designations for a key, and metrics found in literature [4, 48], such as tonal distance between chords. To distinguish standard harmonization from atypical harmonization, we introduce our own metric of a chord progression's typicality. After evaluating the meaningfulness, we implemented the metrics shown in Figure 3. We automatically estimate the key of samples, but also allow manual changes, so the composer can check different settings. To analyze the harmonic connection between notes, we classify notes as: within the chord, in the key's scale, a proposition, or outside the key. Propositions are short notes resolving into scale notes to build tension.

3.2 Visual Design

A composer's workflow using HarmonAI starts with loading a melody. We display that melody in a piano roll, commonly used in DAWs, where the notes are individually colored blocks. Next, the composer generates a large number of AI-suggested harmonizations (T1), using all AI parameter configurations and filtering duplicates.

Getting an Overview: After generating harmonizations, the user can explore the generated samples to find interesting ones (T2). To start this process in a structured way, visual analytics research [22] often uses a similarity map [22, 37] or a tabular-based layout [15]. Similarity maps have been used for music data [33],




Figure 2: HarmonAI’s approach and workflow. Based on a melody, we use AI to sample harmonizations, calculate metrics, and visualize them in a tabular overview and a lead sheet. Users explore and investigate samples and export their favorites.




Figure 3: An overview of our metrics and nine generated harmonizations from the usage scenario (Section 4). Given the example premise to search for non-typical samples, the bottom three are the most interesting to investigate. We analyze the visualized metrics in increasing order to explore the potential of these three chord progressions.

but struggle to show many metrics simultaneously. As such, we decided to follow a tabular-based approach in LineUp [15], which allows users to interactively create, compare, and analyze multi-attribute data in rankings, enabling examination of how different attribute combinations vary. We adopt LineUp’s structure for AI-assisted music composition, where rows resemble harmonizations and columns show musical metrics and characteristics (Figure 1-1).

During exploration of generated samples, the composer might have ideas, such as atypical chord progressions (T3). To efficiently identify relevant samples among numerous generated results, we enable filtering and sorting based on selected metrics, thereby improving scalability. Furthermore, composers can personalize the visualizations, change the key, and listen to the most interesting harmonizations without relying solely on visual cues.

Harmonic Coloring: Because a composer might be interested in the harmonic consonance of notes and chords (T4), we use color based on chroma (C, C#, E,...). Research has proposed several approaches for this, using circular color palettes mapped to the circle of fifths [8, 28, 29] using rainbow color scales with perceptual problems for colorblind people [14]. To tackle this issue, we use a cyclic color palette [43] for the circle of fifths (Figure 1-3) throughout our visualizations. Through that, users can quickly identify consonant notes through color similarity. For example, a C major melody with only greenish/reddish notes is perceived as harmonically consonant, as reflected in the same color scheme of the tonic chord C-E-G.

Identify and Inspect Characteristics: While exploring harmonization, the composer may still need additional information to identify interesting samples (T2, T4). Therefore, we visualize all

our metrics based on their value type: single value, progression over time, or category. For example, for composers who are looking for single values, such as a sample’s typicalness, we display them as a horizontal bar or dot for quick comparison. For a composer, the key plays an important role because it determines harmonic connections, consonance, and a chord’s functionality. We apply the same color scheme as before and align major and minor keys accordingly, so C major and A minor seem similar. Time information is essential for composers and music. Therefore, we visualize most metrics at each time step with color to show a progression, allowing users to quickly spot, e.g., the Roman numerals over time. We use different color scales to emphasize a metric’s meaning, such as distinguishable categories or continuous values. For example, pitch consonance ranges from dissonant (red) to consonant (blue), allowing composers to find mainly consonant samples (T2).

Details and Music Representation: Once composers select an interesting sample to investigate further, they might want to examine the detailed relationship between chords and the melody (T4). For this, we use the piano roll with colored notes and a lead sheet, which compresses harmonization, both commonly used in musicology (Figure 1-2). We show the chords’ names and functionality (Roman numeral) within the key with the metric-based background colors. Through this coloring and time alignment, composers can connect the consonance between chords and specific notes.

Working With Samples: Within HarmonAI, visualizations are the primary medium, although listening remains a central element in music composition. Therefore, composers can also listen to selected samples, find interesting samples, and export their favorites.

4 Usage Scenario

During composing, finding a suitable harmonization for a melody, the main goal of HarmonAI, is dependent on the composer's intent. We therefore demonstrate HarmonAI's capabilities in a specific scenario in an easy setup with nine harmonizations for illustrative purposes: A composer imagines a consonant, good-sounding harmonization, but with twists. We start by loading a melody and generating harmonizations (T1) for 100 parameter configurations.

Firstly, we want to identify samples with anomalies (T3), so we rank by typicality score and can quickly spot the bottom three as atypical (Figure 3-1). Using the typicality progression (Figure 3-2), we can spot the atypical chord transitions (gray areas) appearing in different positions: the beginning, the end, and in the middle.

As this sounds interesting, we examine the consonance and harmonic relationship to the melody. Therefore, we take a look at the pitch consonant score and progression (Figure 3-3). The bottom sample shows many peach and white sections, indicating slight dissonance, and the Roman numerals (Figure 3-4) reveal alternating chords that we had not in mind, so we remove this candidate (T2).

We inspect the remaining two samples individually to analyze their characteristics and intention (T4). Firstly, the sample with the atypical beginning shows a normal, but dissonant ending by returning to the key's tonic in Roman numerals (Figure 3-4), typical for ending sequences. Looking at tone classes (Figure 3-5), many notes outside the key occur, and the tonal distance between chords and melody increases, indicating some tension (Figure 3-6). Thus, the sample is interesting when tension is reduced at the end.

Secondly, the sample with an atypical ending could be a promising candidate for continuing a section rather than ending it. Investigating the consonance (Figure 3-3) reveals that the chords sound mostly consonant, and the Roman numerals (Figure 3-4) reveal the most important chords of a key in the beginning (I, IV, and V). These do not conflict much with the melody, as the tonal distance between them seems low most of the time (Figure 3-6). Combined with tone classes, we can find a few propositions (peach notes) that build tension towards the next harmonic part. Finally, the chord tonal distance (Figure 3-7) reveals higher distances, which is interesting, due to the consonant sound. In the end, we prefer this sample, but would use both samples in our composition.

In summary, we used HarmonAI to generate and investigate harmonizations and to identify samples with different purposes and distinct harmonic relationships. With an initial idea, we illustrated how composers can come to our conclusions using our approach.

5 Discussion

Throughout our design process and usage scenario, we demonstrate that visual analytics support harmonization through generative AI. While visual analytics has been applied to music composition in prior work [30, 33], we utilize generative AI and prioritize metrics and analytical structure. Although music composition is a subjective, ill-defined, and creative task, harmonization often follows mathematical rules within music theory [32]. Consequently, the nature of AI-assisted harmonization aligns well with the strengths of a tabular and metric-based visual analytics approach [11]. Building on the rule-based nature of harmonization, during the careful

selection of metrics and computational methods, we learned that existing algorithmic representations do not fully capture the breadth of music theory, especially in uncommon cases. Such ambiguity arises, for example, when key and scale alter a chord's function and interpretation, and should be considered during design for user interaction. Consequently, our current metric set should be understood as a foundational starting point. A complete set for all compositional practices remains an open challenge.

Given our abstract metric-based representations, it is unclear how composers interpret and use them in a musical context. A useful reference concept is notational audiation [6], a skill that enables mental connections of visual representation and music sound, which we partially support through piano rolls. In the overview, visual representations primarily show abstract metrics and progressions. Despite this, we believe that over time, a skilled user may learn to map color-based progressions to their sound and harmonic intent. More broadly, such skills reflect the tendency to map visual patterns to underlying meaning in visual analytics [2, 24, 33]. In our case, we introduce a more specific and weaker form of auditory imagery [20], however, its effectiveness remains to be examined.

Limitations: A natural limitation of our work is that the selected metrics may not cover all individual needs and aspects of harmonization. Music is highly subjective, so an all-purpose solution is infeasible and requires customization. We argue that our metrics provide a useful foundation for many users. Related is the limited evaluation by other composers and users. While we followed Design Study Methodology [36] and leveraged our own expertise for participatory design [39], our decisions might introduce biases [13] and not generalize to every composer. We deliberately decided to demonstrate HarmonAI in a usage scenario, as a thorough evaluation of generalizability [40, 46] is out of scope.

6 Conclusion

In conclusion, we present HarmonAI, a visual analytical approach to augment composers' manual harmonization workflows with generative AI and a structured, tabular overview. As curators, composers use visualized metrics to find, investigate, analyze, and export interesting harmonization samples. We demonstrate HarmonAI in a usage scenario and identify interesting samples based on harmonic information. Throughout the process, we leverage our domain expertise to reflect on design considerations. In the future, we plan to conduct a thorough, longitudinal study with multiple composers to examine HarmonAI's functionality and flexibility within their adapted workflows and to assess the completeness of the metrics.

Ethics and Privacy Statement

We used AI during prototype development/coding (e.g., pretrained open-source AI [47]) and generative language models (ChatGPT, Grammarly) to improve grammar and readability across all sections. Common prompt: "Improve on writing and grammar: {part of text}." We reviewed and approved all revisions, and all scientific content, ideas, and conclusions are our own.

Acknowledgments

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 251654672 – TRR 161.

References

- [1] Kazuhisa Akimoto and Junichi Hoshino. 2020. SketTune: Real-time input assistance for novices to compose music for self-expression. *Entertainment Computing* (2020), 100311. doi:10.1016/j.entcom.2019.100311
- [2] Basak Alper, Nathalie Henry Riche, Fanny Chevalier, Jeremy Boy, and Metin Sezgin. 2017. Visualization Literacy at Elementary School. In *Conf. Human Factors in Computing Systems (CHI)*. 5485–5497. doi:10.1145/3025453.3025877
- [3] Christopher Antila and Julie Cumming. 2014. The VIS Framework: Analyzing Counterpoint in Large Datasets. In *Conf. Int. Society for Music Information Retrieval (ISMIR)*. 71–76. doi:10.5281/zenodo.1417766
- [4] Tony Bergstrom, Karrie Karahalios, and John C. Hart. 2007. Isochords: visualizing structure in music. In *Graphics Interface Conf.* 297–304. doi:10.1145/1268517.1268565
- [5] Lancelot Blanchard, Cameron Holt, and Joseph Paradiso. 2025. AI Harmonizer: Expanding Vocal Expression with a Generative Neurosymbolic Music AI System. In *Int. Conf. New Interfaces for Musical Expression (NIME)*. 578–581. doi:10.5281/zenodo.15698966
- [6] Warren Brodsky, Yoav Kessler, Bat-Sheva Rubinstein, Jane Ginsborg, et al. 2008. The mental representation of music notation: notational audiation. *Journal of Experimental Psychology: Human Perception and Performance* (2008), 427–445. doi:10.1037/0096-1523.34.2.427
- [7] Yi-Wei Chen, Hung-Shin Lee, Yen-Hsing Chen, and Hsin-Min Wang. 2021. SurpriseNet: Melody harmonization conditioning on user-controlled surprise contours. In *Conf. Int. Society for Music Information Retrieval (ISMIR)*. 1–8. doi:10.5281/zenodo.5624422
- [8] Peter Ciuha, Bojan Klemenc, and Franc Solina. 2010. Visualization of concurrent tones in music with colours. In *ACM Int. Conf. Multimedia*. 1677–1680. doi:10.1145/1873951.1874320
- [9] W. Bas de Haas, Frans Wiering, and Remco C. Veltkamp. 2013. A geometrical distance measure for determining the similarity of musical harmony. *Int. Journal of Multimedia Information Retrieval* (2013), 189–202. doi:10.1007/s13735-013-0036-6
- [10] Chris Donahue, Shih-Lun Wu, Yewon Kim, Dave Carlton, Ryan Miyakawa, and John Thickstun. 2024. Hookpad Aria: A Copilot for Songwriters. In *Late-Breaking Demo Conf. Int. Society for Music Information Retrieval (ISMIR-LBD)*. <https://ismir2024program.ismir.net/lbd%5F489.html#lbd>
- [11] Alex Endert, M Shahriar Hossain, Naren Ramakrishnan, Chris North, et al. 2014. The human is the loop: New directions for visual analytics. *Journal of Intelligent Information Systems* (2014), 411–435. doi:10.1007/s10844-014-0304-9
- [12] Morwaread M Farbood, Egon Pasztor, and Kevin Jennings. 2004. Hyperscore: a graphical sketchpad for novice composers. *IEEE Computer Graphics and Applications* (2004), 50–54. doi:10.1109/MCG.2004.1255809
- [13] Yu Fu and John Stasko. 2023. HoopInSight: Analyzing and Comparing Basketball Shooting Performance Through Visualization. *IEEE Trans. Visualization and Computer Graphics (TVCG)* (2023), 858–868. doi:10.1109/TVCG.2023.3326910
- [14] Matthias Geissbuehler and Theo Lasser. 2013. How to display data by color schemes compatible with red-green color perception deficiencies. *Opt. Express* (2013), 9862–9874. doi:10.1364/OE.21.009862
- [15] Samuel Gratzl, Alexander Lex, Nils Gehlenborg, Hanspeter Pfister, and Marc Streit. 2013. LineUp: Visual Analysis of Multi-Attribute Rankings. *IEEE Trans. Visualization and Computer Graphics (TVCG)* (2013). doi:10.1109/tvcg.2013.173
- [16] Christopher Harte, Mark Sandler, and Martin Gasser. 2006. Detecting harmonic change in musical audio. In *ACM Workshop on Audio and Music Computing Multimedia*. 21–26. doi:10.1145/1178723.1178727
- [17] Cheng-Zhi Anna Huang, Curtis Hawthorne, Adam Roberts, Monica Dinculescu, et al. 2019. The Bach Doodle: Approachable Music Composition with Machine Learning at Scale.. In *Conf. Int. Society for Music Information Retrieval (ISMIR)*. 793–800. doi:10.5281/zenodo.3527930
- [18] Cheng-Zhi Anna Huang, Hendrik Vincent Koops, Ed Newton-Rex, Monica Dinculescu, et al. 2020. AI Song Contest: Human-AI Co-Creation In Song-writing. In *Conf. Int. Society for Music Information Retrieval (ISMIR)*. 708–716. doi:10.5281/zenodo.4245530
- [19] Cheng-Zhi Anna Huang, Tim Cooijmans, Adam Roberts, Aaron Courville, and Douglas Eck. 2017. Counterpoint by Convolution. In *Conf. Int. Society for Music Information Retrieval (ISMIR)*. 211–218. doi:10.5281/zenodo.1416369
- [20] Timothy L Hubbard. 2010. Auditory imagery: empirical findings. *Psychological Bulletin* (2010), 302–329. doi:10.1037/a0018436
- [21] Shulei Ji, Xinyu Yang, and Jing Luo. 2023. A Survey on Deep Learning for Symbolic Music Generation: Representations, Algorithms, Evaluations, and Challenges. *ACM Comput. Surv.* (2023). doi:10.1145/3597493
- [22] Daniel Keim, Gennady Andrienko, Jean-Daniel Fekete, Carsten Görg, et al. 2008. Visual analytics: Definition, process, and challenges. In *Information visualization*. 154–175.
- [23] R. Khulusi, J. Kusnick, C. Meinecke, C. Gillmann, J. Focht, and S. Jänicke. 2020. A Survey on Visualizations for Musical Data. *Computer Graphics Forum (CGF)* (2020), 82–110. doi:10.1111/cgf.13905
- [24] Sukwon Lee, Sung-Hee Kim, Ya-Hsin Hung, Heidi Lam, et al. 2016. How do People Make Sense of Unfamiliar Visualizations?: A Grounded Model of Novice's Information Visualization Sensemaking. *IEEE Trans. Visualization and Computer Graphics (TVCG)* (2016), 499–508. doi:10.1109/TVCG.2015.2467195
- [25] Hugo Lima, Carlos Santos, and Bianchi Meiguins. 2019. Visualizing the Semantics of Music. In *Int. Conf. Information Visualisation (IV)*. 352–357. doi:10.1109/IV.2019.00066
- [26] Hugo Lima, Carlos Dos Santos, and Bianchi S. Meiguins. 2021. A Survey of Music Visualization Techniques. *Comput. Surveys* (2021), 1–29. doi:10.1145/3461835
- [27] Ryan Louie, Andy Coenen, Cheng Zhi Huang, Michael Terry, et al. 2020. Novice-AI Music Co-Creation via AI-Steering Tools for Deep Generative Models. In *Conf. Human Factors in Computing Systems (CHI)*. 1–13. doi:10.1145/3313831.3376739
- [28] Delfina Malandrino, Donato Pirozzi, Gianluca Zaccagnino, and Rocco Zaccagnino. 2015. A Color-Based Visualization Approach to Understand Harmonic Structures of Musical Compositions. In *Int. Conf. Information Visualisation (IV)*. doi:10.1109/iv.2015.21
- [29] Delfina Malandrino, Donato Pirozzi, and Rocco Zaccagnino. 2018. Visualization and Music Harmony: Design, Implementation, and Evaluation. In *Int. Conf. Information Visualisation (IV)*. 498–503. doi:10.1109/IV.2018.00092
- [30] Matthias Miller, Alexandra Bonnici, and Mennatallah El-Assady. 2019. Augmenting Music Sheets with Harmonic Fingerprints. In *ACM Symp. Document Engineering*. 1–10. doi:10.1145/3342558.334539
- [31] Matthias Miller, Daniel Fürst, Hanna Hauptmann, Daniel A. Keim, et al. 2022. Augmenting Digital Sheet Music through Visual Analytics. *Computer Graphics Forum (CGF)* (2022), 301–316. doi:10.1111/cgf.14436
- [32] W. Piston, M. DeVoto, and A. Jannery. 1987. *Harmony*. Norton. <https://books.google.de/books?id=MozDQgAACAAJ>
- [33] Simeon Rau, Frank Heyen, Benedikt Brachtel, and Michael Sedlmair. 2025. MAICO: A Visualization Design Study on AI-Assisted Music Composition. *IEEE Trans. Visualization and Computer Graphics (TVCG)* (2025), 4110–4125. doi:10.1109/TVCG.2025.3539779
- [34] Roger Redgate. 2019. Do you hear what I hear? Audiation and the Compositional Process. *Principles of Music Composing* (2019), 19–27. <https://research.gold.ac.uk/id/eprint/27208/>
- [35] Nicolas Ruwet and Mark Everist. 1987. Methods of Analysis in Musicology. *Music Analysis* (1987), 3–36. doi:doi.org/10.2307/854214
- [36] Michael Sedlmair, Miriah Meyer, and Tamara Munzner. 2012. Design Study Methodology: Reflections from the Trenches and the Stacks. *IEEE Trans. Visualization and Computer Graphics (TVCG)* (2012), 2431–2440. doi:10.1109/TVCG.2012.213
- [37] Michael Sedlmair, Tamara Munzner, and Melanie Tory. 2013. Empirical Guidance on Scatterplot and Dimension Reduction Technique Choices. *IEEE Trans. Visualization and Computer Graphics (TVCG)* (2013), 2634–2643. doi:10.1109/TVCG.2013.153
- [38] Ben Shneiderman. 2007. Creativity support tools: accelerating discovery and innovation. *Communications ACM* (2007), 20–32. doi:10.1145/1323688.1323689
- [39] Clay Spinuzzi. 2005. The Methodology of Participatory Design. *Technical Communication* (2005), 163–174. <https://api.semanticscholar.org/CorpusID:38600380>
- [40] Dan Stowell, Mark D Plumbley, Nick Bryan-Kinns, et al. 2008. Discourse analysis evaluation method for expressive musical interfaces. In *Int. Conf. New Interfaces for Musical Expression (NIME)*. Zenodo, 1–6. doi:10.5281/zenodo.1179631
- [41] David Temperley. 1999. What's the Key for Key? The Krumhansl-Schmuckler Key-Finding Algorithm Reconsidered. *Music Perception (MP)* (1999), 65–100. <http://www.jstor.org/stable/40285812>
- [42] John Thickstun, David Leo Wright Hall, Chris Donahue, and Percy Liang. 2024. Anticipatory Music Transformer. *Trans. Machine Learning Research (TMLR)* (2024), 1–20. doi:10.48550/arXiv.2306.08620
- [43] Kristen M. Thyng, Chad A. Greene, Robert D. Hetland, Heather M. Zimmerle, and Steven F. DiMarco. 2016. True Colors of Oceanography: Guidelines for Effective and Accurate Colormap Selection. *Oceanography* (2016), 9–13. doi:doi.org/10.5670/oceanog.2016.66
- [44] Hiroaki Tsushima, Eita Nakamura, Katsutoshi Itoyama, and Kazuyoshi Yoshii. 2018. Interactive arrangement of chords and melodies based on a tree-structured generative model. In *Conf. Int. Society for Music Information Retrieval (ISMIR)*. 145–151. doi:10.5281/zenodo.1492364
- [45] Anja Volk. 2008. Persistence and change: Local and global components of metre induction using Inner Metric Analysis. *Journal of Mathematics and Music* (2008), 99–115. doi:10.1080/17459730802312399
- [46] Marcelo Mortensen Wanderley and Nicola Orio. 2002. Evaluation of Input Devices for Musical Expression: Borrowing Tools from HCI. *Computer Music Journal* (2002), 62–76. doi:10.1162/014892602320582981
- [47] Shangda Wu, Yue Yang, Zhaowen Wang, Xiaobing Li, and Maosong Sun. 2024. Generating chord progression from melody with flexible harmonic rhythm and controllable harmonic density. *EURASIP Journal Audio, Speech, and Music Processing* (2024), 1–13. doi:10.1186/s13636-023-00314-6
- [48] Yin-Cheng Yeh, Wen-Yi Hsiao, Satoru Fukayama, Tetsuro Kitahara, Benjamin Genchel, et al. 2021. Automatic Melody Harmonization with Triad Chords: A Comparative Study. *Journal of New Music Research* (2021). doi:10.1080/09298215.2021.1873392