ResonaVis: Visualizing Interactive Music Data to Support Reflective Music Composition for Therapeutic Contexts

Abhishek Karwankar, Elise Ruggiero, Daniel Stevens, Matthew Louis Mauriello

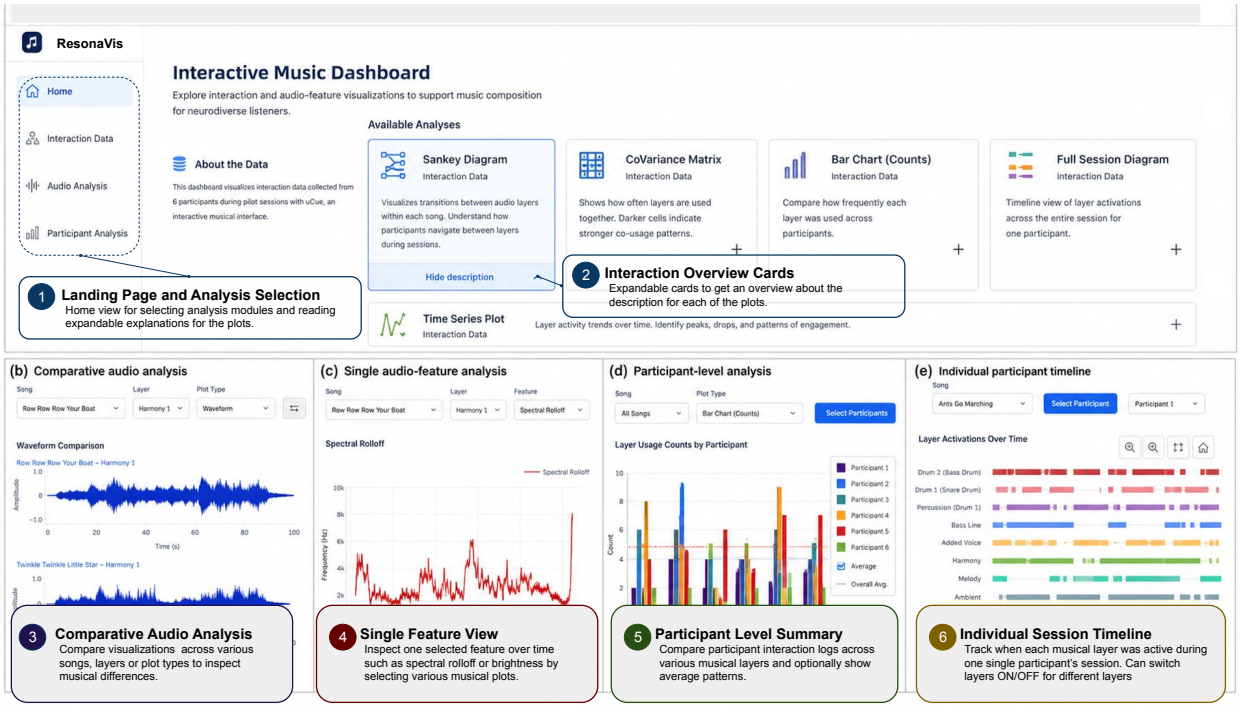

Figure 1: ResonaVis combines interaction-data and audio-feature visualizations to help composers explore participant engagement with musical layers, including analysis selection, comparative audio inspection, participant-level summaries, and session timelines.

Abstract—Designing music for therapeutic contexts requires navigating complex relationships between musical structure and listeners' sensory responses, yet composers often lack structured representations of these interactions to inform compositional decisions, relying instead on intuition. We present ResonaVis, an interactive visualization system to help composers analyze interaction and audio data from prior sessions with children with Autism Spectrum Disorder (ASD), informing future compositions. ResonaVis integrates audio features and interaction logs to capture how children engage with layered musical compositions, representing this engagement through coordinated visualizations of temporal transitions, layer co-occurrence, rhythmic activity, and spectral characteristics. Rather than prescribing strategies or directly supporting therapy sessions, the system surfaces patterns in past session data, supporting data-informed reflection during iterative composition. We evaluated ResonaVis through a mixed-methods study with eight music students and a follow-up case study with two experienced composers. Results demonstrate good usability (SUS = 72.23), exceeding benchmarks for early prototypes, and indicate that participants were able to identify interaction patterns, reason about layer relationships, and make informed compositional decisions. Despite moderate cognitive demands, participants reported high perceived performance and low frustration, suggesting productive engagement. Participants' confidence in interpreting interaction and acoustic data when composing for individuals with ASD also increased significantly across multiple dimensions—rhythm and dynamics, pitch and timbre, and spectral characteristics (all $p < 0.05$). Finally, qualitative findings suggest the potential to shift composers from intuition-driven toward more adaptive, data-informed composition. This work contributes a visualization design space for therapeutic music interaction data, an integrated system for compositional reflection, and empirical evidence that visualization tools can support analytical reasoning and confidence in data-informed creative practices.

Index Terms—Music; Autism; Software; Prototyping; Interactive interfaces

1 Introduction

Music therapy is a widely recognized intervention for children with Autism Spectrum Disorder (ASD), supporting emotional regulation [25], communication [26], and social engagement [63]. However, some therapeutic approaches rely on passive listening or facilitator-guided participation, limiting opportunities to adapt music to individual sensory preferences and behavioral needs. Because children with ASD exhibit highly individualized responses to rhythm, timbre, and musical structure, therapeutic outcomes depend on how effectively music can

• *All authors were affiliated with the University of Delaware at the time of work. E.R. is now affiliated with the Florida State University.*
• *Emails: karwabhi@udel.edu, emrugg@udel.edu, stevens@udel.edu, mlm@udel.edu*

Manuscript received xx xxx. 202x; accepted xx xxx. 202x. Date of Publication xx xxx. 202x; date of current version xx xxx. 202x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxxx/TVCG.202x.xxxxxxx

be personalized [16, 39]. While prior work has focused on the benefits of music exposure, less attention has been given to how data generated during interactive sessions can inform compositional decisions.

Interactive music systems such as uCue [36] enable children to explore music by activating and modifying instrumental layers, producing rich temporal interaction logs that capture when musical elements are added, sustained, or removed. These logs encode behavioral responses to musical structure and offer a potential foundation for data-driven composition. However, they are high-dimensional, multi-layered, and temporally complex, making them difficult to interpret through raw data or listening alone. Composers must reason about patterns such as layer co-occurrence, temporal evolution, and relationships between interaction behavior and acoustic features—tasks that are challenging without appropriate visual representations. Prior work has shown that visualizing interaction data can reveal meaningful patterns and support analytical reasoning that would otherwise remain hidden [30].

Interactive visualization systems, particularly visual analytics dashboards, are well-suited to support such tasks by enabling users to explore heterogeneous data, identify relationships, and derive insights through coordinated views [38, 65, 68]. While dashboards have been widely adopted in domains such as healthcare, education, and business intelligence, their role in supporting creative practices—especially in music composition and therapeutic contexts—remains underexplored. This gap motivates an open question: *can interactive visualization support composers in interpreting complex interaction logs to inform therapeutic music composition?* To address this, we design and evaluate *ResonaVis*, an exploratory visual analytics system that investigates how visual representations can support compositional reasoning.

ResonaVis integrates multiple coordinated visualizations—including temporal session timelines, layer transition diagrams, co-occurrence matrices, and acoustic feature plots—to enable exploration of interaction patterns across musical layers and listening sessions. The system was developed through co-design sessions and iterations with composers experienced in therapeutic music composition, ensuring alignment with real-world analytical and creative tasks. We investigate the following research questions: *(RQ1) How can interactive visualization techniques support the process of composing music for children with ASD? (RQ2) In what ways can a visual dashboard help composers identify and interpret relevant musical patterns from session data?* And, *(RQ3) what design implications emerge from music composers engaging with visualization dashboards for therapeutic music composition?*

We evaluate ResonaVis through a mixed-methods exploratory study with eight music students trained in music composition/education and familiar with music therapy, alongside an observational case study with two experienced composers. Participants used ResonaVis to analyze uCue session data and reported significant increase in confidence in composing therapeutic music, particularly in musical aspects such as rhythm and dynamics ($p = 0.0004$), pitch and timbre ($p = 0.038$), spectral energy and brightness ($p = 0.004$), as well as interpretation of visualizations to support music compositional decisions ($p = 0.032$). The system supported the identification of interaction patterns not easily perceived through listening alone and achieved a System Usability Score (SUS) of 72.23 over the two studies, indicating good usability.

This work contributes: *i) ResonaVis*, an interactive visual analytics system that integrates temporal interaction logs and acoustic features to support exploratory analysis of therapeutic music sessions. *ii)* Visualization design studies demonstrating how composers interpret complex interaction data to inform creative decision-making. *iii)* Empirical evidence showing how visualization supports insight generation and confidence in therapeutic music composition. *iv)* Design implications for visual analytics systems supporting creative practices in music and other domains involving complex temporal interaction data.

2 Related Works

This work builds on research in music therapy for ASD, interactive and technology-supported music composition, and visual analytics systems for creative sensemaking. We situate our contribution at the intersection of data visualization and music composition to empower composers to interpret listener interaction data and inform therapeutic music design.

2.1 Music Therapy and Interactive Music Systems for ASD

Music therapy is a well-established intervention for children with ASD, supporting communication, social interaction, and emotional regulation [6, 11, 42, 73, 75]. Prior work demonstrates that music provides a non-verbal medium for expression, enabling meaningful behavioral and emotional engagement [25, 37, 41]. However, a growing body of research highlights that therapeutic effectiveness depends on specific musical elements, including rhythm, pitch, timbre, and frequency, which shape sensory and emotional responses. Studies show that individuals with ASD often exhibit distinct sensitivities to these attributes, such as preferences for rhythmic predictability, lower brightness, and stable pitch ranges [1,34,51,64]. Similarly, observational and interview-based studies indicate that lower-frequency sounds and gentle timbral qualities can promote engagement and reduce sensory overload [12].

While this literature provides important insights into ASD-specific musical preferences, most work focuses on therapeutic outcomes rather than supporting the compositional process itself. In practice, composers must translate these principles into design decisions, often without tools that connect behavioral interaction data with musical structure. Our work builds on these findings by enabling composers to interpret interaction patterns and acoustic features through visualization, supporting alignment between compositional and therapeutic considerations.

2.2 Music Technology and Creative Confidence

Music composition is inherently a creative and exploratory process [33]. Advances in digital tools and interactive systems have expanded the design space for composers, enabling new forms of experimentation, iteration, and personalization [8, 17, 70]. Technologies such as digital musical instruments and interactive interfaces support adaptive and participatory music-making, particularly in contexts involving children or users with diverse needs [7, 48]. These systems facilitate rapid exploration of musical ideas and enable compositions to evolve dynamically based on user interaction [14, 61]. Beyond technological affordances, confidence plays a critical role in shaping compositional practice. Prior work shows that higher confidence is associated with increased creativity, willingness to experiment, and the development of distinctive musical styles [31, 66]. In therapeutic contexts, confidence is particularly important, as composers must balance artistic intent with sensitivity to user needs and therapeutic goals [50]. Nonetheless, existing tools primarily support music creation rather than reflective reasoning about how compositions align with user behavior or therapeutic outcomes. This gap suggests an opportunity for tools that not only enable composition but also support interpretation and decision-making. We explore how visual analytics can enhance composers' confidence by making complex interaction and acoustic data more interpretable.

2.3 Visual Analytics for Creative Sensemaking

Interactive visualization systems are widely used to support sensemaking in complex data environments by enabling users to explore patterns, relationships, and trends through coordinated views [38, 65, 68]. These systems support iterative exploration and abstraction across data and task levels, aligning with established models of visualization design and analysis [55]. Visual analytics research emphasizes the role of interactive systems in supporting analytical reasoning, insight generation, and decision-making [10, 60, 62]. Dashboards, in particular, have been adopted across domains such as healthcare and education to support monitoring, reflection, and intervention [5, 47, 53, 56]. In creative and therapeutic contexts, visualization has been used to represent emotional states, engagement patterns, and creative processes over time [15, 24]. However, most systems are designed for retrospective analysis rather than supporting ongoing creative decision-making. Within music-related applications, visualization has been used to analyze audio features and listener responses, such as mapping spectral, rhythmic, or perceptual attributes to emotional interpretations [49]. Prior work also explores visualizing performance and interaction traces, including gestural input, timing, and improvisational structure [2, 21, 23, 77]. While these approaches reveal meaningful patterns, they rarely connect interaction data back to compositional workflows. In therapeutic settings, this limitation is particularly significant, as interaction patterns may

serve as proxies for user preferences and sensitivities [40, 74]. Our work addresses this gap by framing visualization as a creativity support tool that enables composers to interpret interaction data and translate insights into compositional decisions.

2.4 Music Visualization and Multimodal Representations

Visualization techniques for music have traditionally focused on representing individual musical attributes such as waveform, pitch, and spectral structure [22, 35, 45]. These representations are widely used in analysis and education but often operate on isolated dimensions of musical data. Similarly, compositional support tools have explored visual representations of rhythm and harmony but typically within constrained or pedagogical settings [4, 9, 27]. Recent surveys highlight that most music visualization systems do not integrate multiple data modalities, such as combining acoustic features with interaction or behavioral data [45]. Emerging work has begun to explore visualization as a component of creative workflows, but these systems remain limited in scope and rarely address real-world therapeutic contexts [18, 52]. In contrast to traditional digital audio workstations (DAWs), which focus on direct music creation, our approach emphasizes reflective analysis through multimodal visualization. ResonaVis integrates temporal interaction logs with acoustic features such as spectral energy, pitch, and dynamics, enabling composers to explore how user behavior and musical structure interact over time. This multimodal perspective supports a deeper understanding of engagement patterns and provides a foundation for data-informed compositional reasoning. ResonaVis operationalizes this perspective by enabling exploratory analysis of complex session data to inform creative decision-making.

3 Iterative Design Process

We describe the dataset, design process, composition challenges, dashboard design, and how the visualizations address these challenges.

3.1 Data and Design Context

The dataset used in ResonaVis consists of approximately four hours of music interaction logs collected from six children with ASD during guided play sessions using the uCue interface [36, 46]. While derived from sessions with children with ASD, they were not involved in the system's design or evaluation; instead, the data serve as a proxy for understanding interaction patterns. During these sessions, children interacted with modular musical compositions by activating and deactivating instrumental layers in real time through a tangible button-based interface. Each interaction event corresponding to a button press or release was automatically recorded by the system with timestamps and layer identifiers, producing a temporally ordered log of user actions. Logs were exported as CSVs and are referred to as *Interaction Data*.

The interaction logs are meaningful because they provide a fine-grained, behavioral trace of how children engage with different musical elements over time. Unlike post-hoc observations or subjective reports, these logs directly capture moment-to-moment interaction patterns, including which layers are preferred based on how frequently they are activated and how long they are sustained. As such, they offer an objective basis for understanding how musical structure influences engagement in therapeutic contexts.

Importantly, these logs also serve as a proxy for engagement behavior. Sustained activation of layers can indicate one's continued interest or comfort with specific sounds, while frequent toggling or rapid transitions may reflect exploration, curiosity, or sensitivity to certain auditory stimuli. Patterns such as repeated co-activation of layers, temporal sequencing of musical elements, and duration of interaction provide insights into how children respond to rhythm, timbre, and musical complexity. However, because these behaviors are encoded as high-dimensional temporal data across multiple layers, they are difficult to interpret without appropriate visual representations.

In addition to interaction data, we incorporate the full collection of musical tracks used in uCue to extract audio features, referred to as *Musical Data*. Together, these two data sources inform the selection of ten visualizations: five derived from interaction behavior and five from acoustic characteristics. Figure 2 illustrates the ResonaVis workflow, showing how interaction and musical data are transformed into coordinated visual representations to support compositional analysis.

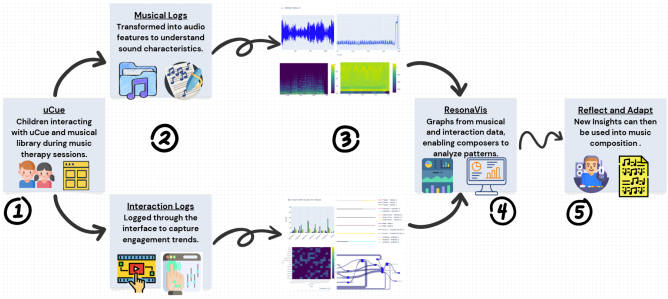

Figure 2: **ResonaVis workflow for data-driven music composition.** (1) Children interact with the uCue system during therapy sessions, generating interaction logs. (2) These logs are processed into interaction-based visualizations. (3) In parallel, musical data is transformed into acoustic feature visualizations. (4) ResonaVis integrates both interaction and musical representation. (5) Insights from this analysis inform the composition of new therapeutic music.

3.2 Participatory Design

Participatory Design emphasizes collaboration between stakeholders and designers to create user-centered solutions [28, 54]. We adopted this approach to develop ResonaVis with music composers, ensuring alignment with real-world compositional practices. Beyond interface design, we focused on identifying *visual representations* and *analytical workflows* to support interpretation of interaction data from therapeutic sessions. Following a visualization design study methodology [78], we iteratively refined data abstractions, tasks, and visual encodings with domain experts. Prior work in music and affective technology highlights the importance of participatory approaches for aligning systems with users' perceptual and creative needs [29, 76].

3.2.1 Stakeholder Collaboration

Our team collaborated with two additional undergraduate researchers and music composers as a part of the University of Delaware's (UD) Undergraduate Research Program, distinct from the Study 1 evaluation participants, over a seven-week period through weekly design sessions. The collaboration began by eliciting composers' existing workflows, challenges, and strategies for reasoning about musical structure. We then introduced sample interaction logs from uCue and explored how such data might inform composition. Through iterative discussions, sketching, lo-fi and mid-fi prototyping, we co-developed candidate visual representations, refined their encodings, and evaluated their usefulness for compositional reasoning tasks. This process was informed by prior work in music visualization [45], and supported through feedback via meetings and asynchronous communication.

3.2.2 Design Challenges in Music Composition

The collaboration revealed recurring challenges in composing for diverse audiences, including understanding sound compatibility, balancing complexity, and interpreting listener preferences. Importantly, composers emphasized the difficulty of reasoning about *temporal interaction patterns* and *layer relationships* from raw data, highlighting a gap between available data and actionable insight.

Across discussions, three themes consistently emerged. First, composers struggled with *relational reasoning*, particularly identifying which layers or sounds work well together and how combinations evolve over time. Second, they noted challenges in *managing complexity*, where increasing the number of layers often led to diminished clarity or overwhelming auditory experiences. Third, composers expressed difficulty in *understanding and validating listener preferences*, especially given the variability across individuals and contexts.

These challenges collectively point to the need for visual representations that make interaction behavior interpretable, support comparison across compositions, and provide insight into both structural and perceptual aspects of music. A comprehensive list of identified challenges is provided in the appendix of the full paper at https://osf.io/xg6su.

3.3 Design Goals

Based on these challenges, we derived a set of design goals to guide the development of ResonaVis: *(DG1)* the system should support reasoning about relationships between musical layers, including co-occurrence and transitions, *(DG2)* it should enable analysis of temporal interaction patterns, capturing how engagement evolves over time, *(DG3)* it should integrate interaction data with acoustic features to support cross-modal reasoning, *(DG4)* it should facilitate comparison across sessions, songs, and participants, and *(DG5)* the system should support the generation of actionable insights, enabling composers to translate observed patterns into compositional decisions.

3.4 ResonaVis Design

Guided by these goals, ResonaVis integrates multiple visual encodings and interaction techniques to support exploratory analysis of musical and behavioral data (Figure 1).

3.4.1 Visual Encoding of Interaction Patterns

A central design goal was to transform raw interaction logs into visual encodings whose structure matches each analytical task, rather than encodings chosen for visual appeal alone. Guided by established task-encoding principles and refined through informal composer feedback, we selected each representation after weighing alternatives. To support reasoning about layer relationships (DG1), Sankey diagrams were chosen over chord diagrams or adjacency matrices because directed flow preserves transition order and remains legible across many co-occurring layers. Covariance matrices were favored over scatterplot matrices for representing pairwise layer co-usage, since composers prioritized rapid identification of strongly correlated pairs over precise distributional detail. To capture how engagement evolves over time (DG2), temporal session diagrams and time-series plots record layer activations over time, enabling detection of patterns such as sustained usage or rapid transitions; we aligned these plots on a shared, continuous time axis specifically to support the cross-layer and cross-song comparison required by (DG4), rather than disjoint or circular timelines that would obscure such comparisons.

3.4.2 Cross-Modal Visualization

To integrate interaction data with acoustic features (DG3), ResonaVis combines interaction data (e.g., layer activations, sequences, and durations), musical data (e.g., acoustic features such as loudness, spectral centroid, and frequency distribution), and composition structure (e.g., modular layers and their variations). By coordinating these modalities within a unified visual interface, composers can relate behavioral patterns—such as engagement spikes—to underlying musical properties, including changes in timbre or spectral energy. This cross-modal integration enables analyses that would not be possible when considering a single data source in isolation.

3.4.3 Visual Analytics Workflow

The workflow emerged from participatory design sessions and is structured to support comparison across sessions, songs, and participants (DG4). Composers begin by inspecting interaction data through timelines and summary visualizations to build an understanding of overall session behavior. They then detect patterns such as repeated layer usage, co-occurrence, and temporal trends, which provide insight into engagement and musical structure. This is followed by comparing behaviors across songs or participants to identify consistent preferences or variations. To deepen their analysis, composers relate interaction patterns to underlying musical features using acoustic visualizations (DG3), enabling connections between behavioral responses and properties such as timbre, pitch, and spectral characteristics. These ultimately inform compositional decisions (DG5), where composers modify or select layers based on the patterns identified which reflects a shift from passive log observation to active, visualization-supported exploration.

3.4.4 Supporting Insight Generation

The combination of visual encodings and analytical workflow enables composers to generate insights that are difficult to obtain through listening or raw data alone (DG5). These insights span both interaction behavior and musical characteristics, bridging user engagement with compositional structure. For example, composers identified *rhythmic stability* through layers that remained active over extended durations, as well as variations in activation frequency that reflected changes in musical dynamics. By integrating interaction data with acoustic features, composers further reasoned about *pitch and timbre relationships*, using representations such as spectrograms to understand how different timbral qualities influenced engagement. These insights - reported in detail in Section 4.5.2, directly informed compositional decisions across relational, acoustic, and temporal dimensions.

3.5 Implementation and Interface Design

The dashboard incorporates ten coordinated visualizations grouped into two categories: *Interaction Data* and *Musical Data*. All views were implemented using Plotly[1] and deployed via Dash[2], enabling interactive exploration and linked analysis across representations. Design decisions were guided by established usability principles [69], ensuring consistency, responsiveness, and interpretability.

Rather than introducing novel visualization techniques, our contribution lies in the *systematic integration and contextualization* of existing visual representations to support therapeutic music composition. The visualizations are designed to support three complementary analytical challenges. First, *relational analysis* is supported through views such as the Sankey diagram and covariance matrix, enabling composers to reason about layer transitions and compatibility. Second, *temporal analysis* is facilitated by timeline-based representations (e.g., full-session diagrams and time-series plots), which reveal interaction patterns and engagement over time. Third, *acoustic analysis* is enabled through signal-based visualizations (e.g., waveplots, MFCC (Mel-Frequency Cepstral Coefficients), and spectrograms), supporting reasoning about pitch, timbre, and spectral characteristics. Together, these coordinated views enable composers to move between interaction behavior and acoustic properties, supporting both exploratory analysis and composition-oriented decision-making. A detailed description of each visualization, is provided in the appendix of the full paper at https://osf.io/xg6su. Color schemes (Viridis) and layout choices were refined with stakeholders to ensure accessibility and perceptual clarity, resulting in a cohesive visual analytics system.

4 Study 1: Evaluation with Music Students

For this preliminary study, we selected music students (undergraduate and graduate) trained in composition and music education as our target population [44, 57]. This group was accessible for iterative testing, possessed sufficient knowledge of music structure to interpret the visualizations, and was pedagogically oriented toward developmental contexts relevant to therapy. Their novice-to-intermediate status also made them more open to tool support, allowing us to surface usability and conceptual challenges before involving professional composers.

4.1 Recruitment and Participants

We recruited 8 student composers (5 females, 3 males) via mailing lists and word of mouth. They were aged between 19 and 46 years (mean = 24.75, SD = 8.83) and had formal training in music composition or music education. Six of the eight participants reported familiarity with music therapy or similar practices. The study welcomed participants of all genders, ethnicities, and backgrounds. Importantly, there were no specific medical conditions or exclusion criteria beyond the stated requirements. The study design was approved by the Institutional Review Board (IRB) of UD (Protocol #2097325), ensuring that all

[1] https://plotly.com/
[2] https://dash.plotly.com/

Table 1: **Participant demographics and background.**

ID	G	Age	Deg	Major	Music	Vis
P1	F	46	BSc	English (Music minor)	Both	N
P2	F	22	MSc	Linguistics	Edu	Y
P3	F	19	BSc	Music Perf./Comp.	Both	Y
P4	F	20	BSc	Music + Psychology	Comp	Y
P5	F	20	BSc	Cognitive Sci., Viola Perf.	Both	N
P6	M	22	MSc	Music Composition	Both	Y
P7	M	24	MSc	Music Performance	Both	N
P8	M	25	MSc	Music Performance	Both	N

research activities were conducted in accordance with ethical standards. Informed consent was obtained from all participants, and individual identities were not linked to the data used in the analysis. All participant data was anonymized. Table 1 gives an overview of the participants. The session with P2 was a virtual Zoom session, conducted based on their availability. We also removed data for one participant (P5) from our analysis, as the session was canceled midway due to technical difficulties. Thus, our results focus on the seven participants who successfully completed the study session.

4.2 Study Design

We follow a design study approach, grounding system development in domain-specific needs and iterative collaboration with composers, consistent with established visualization design study methodologies [67]. We combine qualitative and quantitative metrics, focusing on how well composers could interpret the visualizations and their overall satisfaction with the interface. To keep sessions under one hour, we selected a subset of visualizations for the structured evaluation tasks, focusing on relational and acoustic analysis (Section 3.5)—covering rhythmic attack, pitch, frequency, amplitude, overtone, brightness, darkness, timbre, and instrument type—based on discussions with our research team and stakeholders. Temporal analysis visualizations (e.g., full-session diagrams, time-series plots) remained available but were not the focus of a dedicated evaluation task; their value was instead assessed qualitatively through usability feedback (Section 4.5.3). Participants received verbal definitions and examples for shared understanding (sheet in Supplementary Materials). We did not test clinical outcomes, but explored whether the dashboard could scaffold composition decisions aligned with therapeutic goals. Task 1 focused on Interaction Data, and responses from this task served as input for Task 2.

During the study, participants interacted with two sections of the interface: the Interaction and Audio section, which included the Participant Analysis (S1) and Comparative Analysis (S2) interfaces, as shown in Figure 3. In S1, participants were required to select a song of their choice, choose the number of participants to include (with a preference for selecting all to obtain a net average), pick a plot for the challenge, and select the appropriate options from a dropdown menu. The observations made in S1 were then used as input for S2. In S2, participants selected the same song chosen in S1, selected a plot for the challenge, and input their observations from S1 into two comparative plots provided in this section. Finally, they were asked to write down their findings in response to specific questions related to their observations (see Supplementary Materials for the entire study design). Table 2 describes the three design challenges which we formulated collaboratively with our stakeholders using language commonly employed in music therapy [71]. These design challenges were used to assess the participants' understanding and the dashboard's ability to facilitate analysis. We calculate the raw NASA-TLX workload scores [32] after each challenge which captured participants' perceived cognitive and emotional workload, and the System Usability Score (SUS) [72] at the end to determine the overall usability.

4.3 Procedure

Participants engaged with ResonaVis through multiple coordinated visualizations. Each session lasted approximately one hour, and participants took part only once. Sessions were conducted in person at

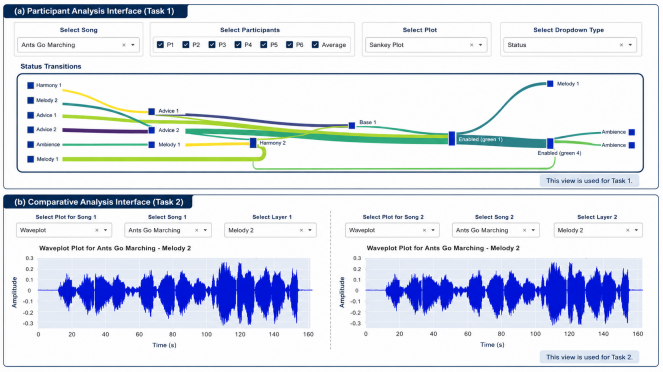

Figure 3: **Study interface used in evaluation tasks.** (a) Participant Analysis Interface (Task 1): users select a song, participants, plot type, and category to explore individual interaction patterns; shown here, a Sankey diagram of status transitions. (b) Comparative Analysis Interface (Task 2): users independently select a plot, song, and layer per panel for side-by-side comparison; shown here, waveplots for the same song and layer on a shared temporal axis (0–160s).

Table 2: **Design Challenges and associated visualizations.**

Q	Evaluation Question	Visualizations
1	How do rhythmic/dynamic characteristics of selected layers complement each other?	Covariance Matrix (S1), Waveplot (S2)
2	How does popularity of frequently used layers relate to spectral energy and perceived brightness?	Bar Graph (S1), Spectral Rolloff/Centroid (S2)
3	How do layer transitions correspond to frequency/pitch changes over time, and how might they affect listening experience?	Sankey Diagram (S1), Mel Spectrogram (S2)

the Sensify Lab in UD for all participants except P2, who participated remotely. In-person sessions were conducted on a 24-inch monitor (1920×1080 resolution), while the remote participant was instructed to use the largest available screen and share their display. All sessions were audio recorded with participants' consent.

The session began with a brief introduction to the dashboard, during which participants were guided through its features and interaction mechanisms. They were encouraged to freely explore the interface to become familiar with the visualizations before proceeding. This was followed by a short pre-session interview, where participants described their typical music composition process, including how they develop and refine musical ideas. Participants then completed a pre-session questionnaire assessing their confidence in composing music for children with ASD and their ability to interpret visual data. Following this, they engaged in the three design challenges designed to evaluate the dashboard. These design challenges required participants to analyze visualizations, extract insights, and relate these insights to potential compositional decisions. After each design challenge, participants documented their observations and provided feedback on the usefulness and clarity of the visualizations.To assess cognitive workload, participants completed the NASA-TLX after each challenge. At the conclusion of all challenges, they completed the SUS test to evaluate overall usability. Participants then repeated the confidence questionnaire to capture changes in perceived ability following interaction with the system. Finally, sessions concluded with a short debriefing, where participants reflected on their experience, discussed challenges encountered, and suggested potential improvements.

4.4 Data Collection and Analysis

In this study, data analysis was conducted using both quantitative and qualitative approaches. Quantitative analysis involved statistical examination of metrics such as SUS scores, and the Raw Task Load Index (scale of 10, following Hart [32], which validates its reliability) for each task. These metrics provided measurable insights into participant per-

formance and the overall usability of the dashboard. On the qualitative side, thematic analysis [13] was employed to analyze feedback from participants and their responses to study questions. This analysis helped identify common themes and suggestions, offering deeper insights into the participants' experiences and guiding potential improvements to the dashboard. The responses of the participants were double coded by the first author along with the music composers on the team.

4.5 Results

The results are organized into subsections covering pre-session interviews, task performance, insights from visualizations, workload and usability scores, confidence changes, and qualitative feedback. Together, these findings provide a holistic view of how participants engaged with the visualisations and ResonaVis. A post-session interview further examined whether their perspectives on composing for children with ASD had shifted.

4.5.1 Pre-Session Interview Results

We conducted a pre-session interview to understand participants' existing approaches to music composition and the decision-making processes underlying their work. Participants were asked to describe their typical workflow, from initial inspiration to refinement, including how they iteratively develop and evaluate their compositions. Analysis of the responses revealed three primary compositional approaches. First, some participants adopted an *inspiration-driven approach*, drawing from familiar musical pieces and using emulation or adaptation as a starting point (P1). Second, a majority of participants described a *segmented and iterative workflow*, in which individual musical elements (e.g., rhythm, melody, or harmony) were composed separately and refined through repeated listening, visual cues, and feedback (P2, P3, P6, P7, P8). Third, one participant emphasized a *therapeutic alignment approach*, revisiting and adjusting compositions to better match emotional goals and principles from music therapy (P4)—such as reducing sensory overload and promoting calmness through predictable rhythmic and tonal structures. Across these approaches, participants relied heavily on subjective judgment, prior experience, and iterative listening, with limited use of structured data or visual analysis to guide their decisions. This suggests that, prior to using ResonaVis, compositional strategies were largely experience-driven and heuristic-based rather than explicitly informed by multi-dimensional analysis of musical features. These baseline insights provide an important point of comparison for subsequent results, enabling us to examine how participants' workflows evolved from intuition-driven practices to more data-informed and adaptive compositional strategies after interacting with the system.

4.5.2 Insights Derived

We conducted a thematic analysis of participants' written responses, organized around the three analytical dimensions from Section 3.5: relational, acoustic, and temporal analysis. As noted in Section 4.2, structured evaluation tasks addressed relational and acoustic dimensions; temporal analysis is discussed qualitatively below. First, *relational reasoning* emerged through compatible layer combinations and transitions. Covariance-based visualizations enabled participants to interpret co-occurrence patterns as indicators of rhythmic alignment and dynamic complementarity, while Sankey diagrams supported interpretation of flow thickness and structure to understand how layer changes contributed to variation, contrast, and perceived smoothness in the auditory experience. Second, *acoustic reasoning* was observed as participants connected frequency-based representations (e.g., spectral centroid) to perceptual qualities such as calmness, intensity, and brightness, and related spectral energy distributions to sustained or reduced engagement. Third, while *temporal reasoning* was not the focus of a dedicated evaluation task, participants qualitatively reported that timeline-based views helped them identify engagement patterns over time that would be difficult to perceive through listening alone. Importantly, participants did not treat these dimensions in isolation. Instead, they integrated acoustic, relational, and dynamic cues to form holistic interpretations of musical structure and perception. This suggests that ResonaVis supports *multi-dimensional reasoning*, enabling users to connect low-level acoustic features with higher-level perceptual and therapeutic interpretations.

Role of ResonaVis: These insights were closely tied to the visual encodings provided by the system. Covariance and Sankey diagrams supported relational reasoning about co-occurrence and transitions, while spectral plots enabled acoustic interpretation of energy distribution and brightness. Together, these coordinated views supported users in constructing meaningful interpretations of complex musical relationships. Qualitatively, rhythmic alignment and spectral energy were among the most frequently observed themes, while transition-related insights captured more holistic aspects of the auditory experience. These findings align with observed confidence increases (Section 4.5.6) and NASA-TLX (Section 4.5.4) workload patterns, suggesting that the cognitive effort contributed to deeper understanding rather than overload. A detailed codebook is provided in the appendix of the full paper at `https://osf.io/xg6su`.

4.5.3 Usability Results

We evaluated the usability of ResonaVis using the System Usability Scale (SUS). The interface achieved a mean SUS score of 70.7 (SD = 17.8) in its first iteration, exceeding both the average score for initial systems (62) and the overall benchmark for web-based interfaces (68.05) reported by Bangor *et al.* [3], placing it in the *Good* range on associated adjective rating scale, above *OK* and approaching *Excellent*. This indicates that participants generally perceived the system as usable and approachable, even at an early stage of development. However, scores varied substantially across participants, ranging from 42.5 to 92.5. While several participants (e.g., P6: 92.5, P8: 85.0) reported high usability and found the system intuitive and well-integrated, others (e.g., P2: 42.5, P7: 52.5) indicated challenges in using the interface. Across analytical dimensions, usability varied by visualization type: relational views, particularly the Sankey diagram, required additional explanation before participants could confidently interpret layer transitions, whereas bar graphs and the covariance matrix were comparatively intuitive to read. Acoustic visualizations were generally accessible, supporting spectral reasoning with minimal guidance. Temporal views, while not formally evaluated in structured tasks, received qualitative endorsement: participants reported that full-session diagrams and time-series plots surfaced engagement patterns difficult to perceive through listening alone. This variability suggests that while the system is effective for some users, it may require additional support or refinement to ensure consistent usability, particularly for those less familiar with interpreting visualizations or multi-dimensional musical data. These observations align with the improvements in confidence (Section 4.5.6) and the moderate cognitive demands observed in the NASA-TLX results (Section 4.5.4), suggesting that the system supports meaningful analytical engagement even when usability is not uniformly optimal. Overall, the SUS results indicate that ResonaVis provides a strong foundation for usability in its current form, while also highlighting opportunities for improving learnability, consistency, and guidance to better support a wider range of users.

4.5.4 NASA-TLX for Each Challenge

We used NASA-TLX to assess perceived workload across tasks and identify opportunities to improve ResonaVis.

Challenge 1: Participants reported high mental demand (5.86) and effort (5.43), with low physical (0.43) and temporal demands (2.43). Despite this, performance was high (7.86) and frustration remained low (2.71), suggesting that while the task required sustained reasoning, participants could engage effectively without negative effect. This indicates that ResonaVis supports cognitively intensive compositional tasks while maintaining usability.

Challenge 2: Mental demand (5.43) remained similar to Challenge 1, but effort (5.00) and frustration (3.00) increased, alongside a drop in performance (6.43). Physical (0.43) and temporal demands (1.57) stayed low. This suggests greater cognitive complexity or ambiguity in interpreting the visualizations. The increased effort and frustration point to opportunities for improved guidance for interaction support.

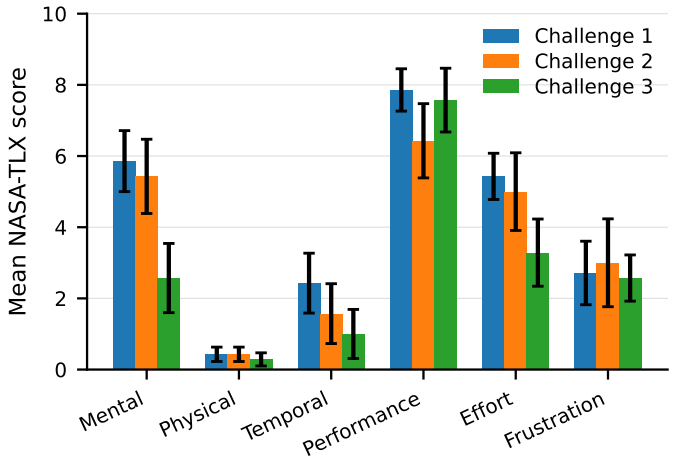

Figure 4: **NASA-TLX scores across the three challenges.** Error bars indicate standard error.

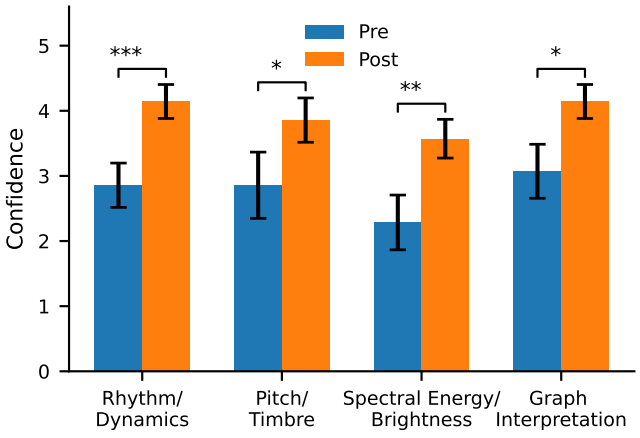

Figure 5: **Pre–post confidence gains:** Error bars show standard error; * indicate statistical significance (* p < 0.05, ** p < 0.01, *** p < 0.001).

Challenge 3: All workload dimensions were low, including mental (2.57), physical (0.29), and temporal demands (1.00), while performance remained high (7.57). Effort (3.29) and frustration (2.57) were also minimal, indicating that the task required little cognitive overhead. This reflects strong baseline usability for simpler tasks.

Cross-Challenge Insights: A clear gradient emerges: Challenges 1 and 2 required higher cognitive engagement, while Challenge 3 was lightweight. Notably, high mental demand in Challenge 1 did not reduce performance or increase frustration, indicating productive cognitive load. In contrast, Challenge 2 suggests that some forms of complexity are less well supported, highlighting the need for better interpretability and scaffolding. Figure 4 summarizes these results.

Relation to Confidence Gains: These workload patterns contextualize the significant confidence gains (Section 4.5.6). Although Challenges 1 and 2 required moderate to high cognitive effort, participants still showed increased confidence across all dimensions. This suggests that the effort was productive, supporting learning and deeper engagement. Rather than minimizing cognitive load, effective visual analytics systems should help users manage and leverage it.

4.5.5 Qualitative Feedback

We analyzed participants' reflections on their compositional strategies before and after using ResonaVis to understand how the system influenced their approach to designing music for individuals with ASD. The analysis reveals a clear shift from *rule-based* to *adaptive* composition strategies. Prior to using the system, participants largely relied on conservative heuristics grounded in assumptions about sensory sensitivities. Across themes, compositions emphasized simplicity, predictability, and reduced sensory complexity, including steady rhythms, limited dynamic variation, mid-to-low pitch ranges, and less bright timbres. These strategies reflected a precautionary approach aimed at avoiding overstimulation. After interacting with ResonaVis, participants demonstrated a transition toward *data-informed exploration*. Rather than adhering to static rules, they began incorporating controlled variation in rhythm, dynamics, timbre, and spectral properties based on observed interaction patterns and visual insights. Participants experimented with rhythmic transitions, and actively balanced spectral energy to achieve desired perceptual effects. Importantly, this shift did not reflect a move toward complexity alone, but toward *informed flexibility*. Participants continued to consider sensory sensitivity but adapted their decisions dynamically based on context. This transition from rigid heuristics to nuanced, context-aware reasoning suggests that ResonaVis supports more reflective and responsive compositional practices.

Cross-Participant Trends: Across participants, a consistent pattern emerged: initial strategies were precautionary and generalized, whereas post-session approaches were more individualized and adaptive. Participants reported increased willingness to experiment, refine, and iterate on compositional elements, indicating a shift toward more context-aware music design. These qualitative changes align with the observed increases in confidence (Section 4.5.6) and NASA-TLX workload patterns, suggesting that the cognitive effort required by the system contributed to deeper understanding rather than overload. A detailed breakdown of these changes is provided in the appendix of the full paper at `https://osf.io/xg6su`.

4.5.6 Changes in Confidence Level

Participants' confidence in composing music for individuals with ASD (0–5 scale) increased significantly after using ResonaVis. Confidence was evaluated across four dimensions: (i) rhythmic and dynamic elements, (ii) pitch and timbre, (iii) spectral energy and brightness, and (iv) interpretation of graphs and data.

Mean confidence increased across all four dimensions. Confidence in working with rhythmic and dynamic elements increased from 2.86 (SE = 0.34) to 4.14 (SE = 0.26; $p = 0.0004$), representing the largest gain. Similarly, confidence in pitch and timbre improved from 2.86 (SE = 0.51) to 3.86 (SE = 0.34; $p = 0.038$), while confidence in reasoning about spectral energy and brightness increased from 2.29 (SE = 0.42) to 3.57 (SE = 0.38; $p = 0.004$). Participants also became more confident in interpreting visual representations of musical data, with scores increasing from 3.07 (SE = 0.22) to 4.14 (SE = 0.26; $p = 0.032$). This improvement suggests that ResonaVis not only supported compositional reasoning but also helped users develop the ability to read and make sense of complex visual encodings.

While the sample size is small and confidence was measured using a Likert scale, the consistent direction of the mean changes indicates a positive trend. We note that participants also received verbal terminology training prior to the tasks (Section 4.2), and we cannot fully disentangle gains attributable to this exposure from gains attributable to interacting with the visualizations themselves; we view confidence as reflecting the combined effect of terminology familiarity and visual analytical practice. These results should therefore be interpreted as indicative rather than conclusive, motivating further validation with a larger and more diverse participant pool. Overall, the gains across all four dimensions indicate that ResonaVis facilitated a broader understanding of acoustic features and their role in therapeutic music composition. Notably, improvements in both compositional aspects (e.g., rhythm, timbre, spectral features) and analytical skills (e.g., graph interpretation) suggest that the system supports both creative and analytical engagement. Figure 5 illustrates these changes in confidence before and after using ResonaVis.

4.5.7 Post-Session Interview Results

Following the study, participants reflected on how the visualizations aligned with or differed from their existing compositional practices,

and whether these representations influenced their approaches to developing and refining music. Overall, participants reported that the visualizations provided new and actionable perspectives on composition. Several participants described how the system supported more informed decision-making prior to composing. For example, P3 noted that the visualizations enabled them to pre-select instruments and identify compatible layers for mixing, particularly when aiming to create soothing auditory experiences. Similarly, P6 and P8 highlighted the usefulness of comparative analysis features, which allowed them to quickly identify differences between layer combinations. P6 further indicated a strong intention to incorporate the tool into their regular compositional workflow, suggesting that ResonaVis can integrate effectively into existing creative practices for some users. At the same time, participants' responses revealed variability in how well the system aligned with individual compositional styles. While P1, P2, and P7 acknowledged the potential value of the tool, they indicated that they might not frequently use it due to differences between their established workflows and the visualization-driven approach. For instance, P7 emphasized a preference for composing through musical notation, suggesting that the absence of note-level representations limited the tool's relevance for their practice. However, participants also noted that increased familiarity with the visualizations, or alignment between the tool and their compositional approach, could improve its utility.

Cross-Participant Insights: Across responses, a key pattern emerges: ResonaVis is particularly effective as a *pre-compositional and analytical support tool*, helping users explore layer relationships, evaluate combinations, and reason about perceptual outcomes before or during composition. Rather than replacing existing workflows, the system appears to augment them by providing additional perspectives that are difficult to obtain through listening alone.

Implications for Adoption: These findings suggest that the usefulness of ResonaVis depends on the degree of alignment between its visualization-driven paradigm and the user's existing compositional style. For participants open to integrating visual and data-driven reasoning, the system supported deeper insight and workflow integration. For others, particularly those relying on notation-based or highly internalized processes, the system was perceived as complementary but not essential. This highlights the importance of designing flexible interfaces that can accommodate diverse compositional practices.

5 Study 2: Observational Case Study

This case study explored ResonaVis within the workflows of two experienced composers (P9 and P10) who have created music for children with ASD and were recruited via word of mouth and through the research team's existing contacts (IRB Protocol #2097325). By observing their interactions with the system and gathering reflections, we examined how engagement-driven visualizations could inform advanced compositional decisions and validate their design choices.

5.1 Method

We conducted two one-hour sessions with experienced composers. P9 (21, Female) was a composer who had previously contributed to the creation of a musical library for children with ASD. P10 (Age not reported, Female) was an experienced music composer and director who has composed for children with ASD. She was pursuing a Master's degree in piano performance and had no prior experience with data visualization tools. P9 participated in a virtual session, while P10 was an in-person session. Participants were introduced to ResonaVis's features and visualizations, followed by a task where they explored interaction logs and musical features (Study design in Supplementary Material). They were encouraged to verbalize their thought processes, identify patterns, and propose compositional adjustments. Both sessions were audio recorded and transcribed for thematic analysis.

5.2 Findings

5.2.1 Baseline Confidence and Composition Approach.

Both participants reflected on the challenges of composing for children with ASD, particularly in the early stages of their practice. P9 described initially relying on intuition and focusing on simple, calming compositions without explicitly considering how different musical features might influence engagement over time. Similarly, P10 reported low initial confidence when beginning to compose for this audience, noting uncertainty about what musical elements were appropriate. To address this, she relied on prior research on sensory preferences, particularly differences in high, medium, and low frequency ranges. Her compositions emphasized simplicity and clarity of layers, with a primary goal of producing calming musical experiences. However, like P9, she noted that aspects such as transitions and temporal dynamics were not systematically considered in her early work.

5.2.2 Experience Using ResonaVis.

After being introduced to ResonaVis, both participants found the system useful in supporting compositional reasoning, though in slightly different ways based on their backgrounds. P9's responses aligned closely with findings from Study 1, particularly in using engagement patterns to reason about when to introduce or delay layers to shape musical structure. She highlighted how the system consolidated "a lot of information" in one place, enabling efficient review of interaction logs without replaying full recordings and supporting rapid "what-if" exploration. P10, despite having no prior experience with data visualization, responded positively to the system and found it valuable for practical composition. She particularly appreciated the bar charts and covariance matrix, noting that they clearly conveyed which layers were most used and how different musical elements co-occurred. She was also able to relate the frequency distributions shown in the mel spectrogram to her compositional decisions, using them to reason about pitch ranges and refine transitions between sections. Reflecting on this process, she noted, *"I can almost imitate the sounds while looking at the transitions in the Sankey Chart (While imitating some musical transitions after this quote)"* indicating how the visualizations supported a mapping between visual patterns and auditory expectations. Across both participants, integrating interaction logs and audio features enabled efficient identification of patterns such as frequently activated layers and transitions, supporting more deliberate decisions about when to introduce, delay, or combine layers to shape the composition's energy and flow.

5.2.3 Perceived Value, Usability, and Suggestions.

Both participants perceived ResonaVis as a useful tool for supporting therapeutic composition. P9 emphasized that such a system could have improved her confidence and independence when she first began composing for this audience, stating, *"When I first started composing for this audience, I often lacked confidence in knowing whether my choices were appropriate. A tool like this could have provided useful guidance and helped me make decisions more independently."*

In terms of usability, P9 rated the system at 77.5 (SUS), while P10 rated it at 70, resulting in an average usability score of 73.75, indicating good overall usability. P10's rating is notable given her lack of prior experience with visualization tools. She described the system as practical and applicable to real-world composition workflows, noting *"It feels like a practical tool I could actually use in my composition process, not something abstract."* At the same time, P10 reported some difficulty in interpreting certain visualizations (similar to section 4.5.3), highlighting a learning curve for users without technical backgrounds. This aligns with P9's suggestion to include clearer explanations or onboarding support. Both participants emphasized the importance of making visualizations more accessible to musicians who may not be familiar with data-driven tools.

5.3 Summary

Overall, the case study demonstrates that ResonaVis can support both experienced and non-technical composers by bridging domain knowledge with data-driven insights. While P9 leveraged the system to refine existing strategies, P10 used it to ground her compositional decisions in visual evidence, particularly in relation to frequency and transitions. Together, their feedback highlights the system's practical value, while also underscoring the need for improved interpretability to support broader adoption among musicians.

6 Discussion

We explore how visualization supports data-driven music composition by enabling composers to interpret patterns and generate insights. Rather than evaluating clinical outcomes, we focus on how visual representations shape compositional reasoning and decisions.

6.1 RQ1: Visualization-Supported Compositional Insights

Our findings show that visualization enables composers to reason about musical structure through observable patterns in rhythm, spectral properties, and temporal transitions (Sections 4.5.2 and 4.5.5). Participants consistently identified relationships such as rhythmic alignment across layers, balance in spectral energy and brightness, and structured transitions between musical elements. These insights were not derived from listening alone, but from recognizing patterns in visual encodings such as timelines, covariance matrices, and flow-based representations.

These observations align with prior work highlighting the importance of rhythmic predictability, timbral balance, and low-brightness sound textures in therapeutic music contexts [1, 34, 51, 64]. However, our contribution lies in showing how such domain knowledge can be surfaced and operationalized through visualization. Rather than prescribing compositional rules, ResonaVis enables composers to discover and validate these relationships through data-driven exploration. From a VIS perspective, this demonstrates how visualization can bridge low-level acoustic features and high-level compositional intent, supporting insight generation in creative domains.

6.2 RQ2: Interpretation and the Role of Interaction Data

Our results further show that the value of ResonaVis emerges not only from individual visualizations, but from the integration of interaction data with audio features within a unified workflow (Section 5.2.3). Participants emphasized that interaction logs—capturing when and how layers were activated—provided essential context for interpreting musical properties. Without this temporal and behavioral grounding, acoustic features were perceived as static and less actionable.

This integration enables a shift from analyzing *what* the music contains to understanding *how* it unfolds through interaction. Visualizations such as Sankey diagrams and timelines supported tracing engagement patterns, identifying co-occurrence across layers, and reasoning about transitions over time. These findings align with prior work on visualization as a tool for exploratory analysis [45], but extend it to a setting where interpretation is inherently creative and context-dependent.

Importantly, usability metrics (SUS scores of 70.71 and 73.75) indicate that the system is learnable and falls within the range typically considered acceptable in HCI studies [3]. However, usability alone does not fully capture the system's impact. Participants also demonstrated consistent increases in confidence across multiple compositional and analytical dimensions (Section 4.5.6), suggesting that ResonaVis supports not only interaction with the interface but also a deeper understanding of musical structure and data representations.

While the sample size is small n=7 (Study 1) and n=2 (Study 2), such sizes are common in visualization and HCI research, where studies emphasize depth of insight and formative evaluation over statistical generalization [20, 43, 58]. Prior work has also shown that parametric tests remain robust for small samples and Likert-scale data when effects are consistent, supporting the interpretation of observed confidence gains as indicative of meaningful trends [19, 59]. Accordingly, we interpret these results in conjunction with consistent directional improvements across participants.

Taken together, these findings suggest that, in creative workflows, the primary value of visualization lies in supporting interpretability and contextual reasoning. By linking interaction behavior with acoustic features, ResonaVis enables users to connect abstract data patterns with concrete compositional decisions, highlighting the importance of integrated, interpretable representations over isolated feature presentation.

6.3 RQ3: Implications for Visualization in Creativity

We demonstrate how visualization can support creative, interpretive tasks by linking ResonaVis design goals with observed user behavior:

Support cross-modal reasoning. Aligned with our goal of integrating interaction and acoustic data (Section 3.4.2), participants used combined views to relate engagement patterns to musical structure (e.g., sustained layer use with stable spectral energy), enabling deeper insight beyond isolated features (Sections 4.5.2 and 4.5.5).

Design for interpretability over completeness. Although the system provides feature-rich views, participants consistently favored interpretable representations, while more abstract features required effort but led to measurable confidence gains (Section 4.5.6), reinforcing our emphasis on intuitive encodings.

Scaffold domain-specific understanding. Consistent with our goal of supporting non-expert interpretation, participants initially struggled with abstract features (e.g., spectral properties) (Section 4.5.6) but demonstrated significant post-use confidence improvements (Section 4.5.6), indicating that visualization can facilitate learning when grounded in domain context.

Enable iterative exploration. Reflecting our exploratory workflow design, participants moved iteratively between inspection, pattern detection, and refinement, highlighting the importance of supporting non-linear analysis in creative tasks. These implications suggest that visualization systems for creative domains should prioritize integrated, interpretable, and workflow-oriented designs over isolated analytics.

6.4 Limitations and Future Work

This study has several limitations. The uCue dataset used for visualization is modest, serving as formative design material rather than a representative sample. This evaluation focuses on early-stage ideation and exploratory analysis, not long-term compositional outcomes. The visualizations represented only a subset of musical parameters, primarily rhythm, dynamics, spectral energy, and brightness, and may not encompass the full range of factors relevant to composition. Finally, the participant pool is limited in size and diversity, which may restrict generalizability, and findings may reflect a narrow slice of compositional approaches. Relatedly, this study does not include clinical trials evaluating the therapeutic benefits, if any, of ResonaVis-composed music for children with ASD; our focus is on the visualization system's capacity to support composers' analytical and compositional reasoning, with clinical efficacy left to future work.

Future work should also explore both technical and methodological extensions. On the technical side, incorporating adaptive components—such as pattern recognition or recommendation modules—could support more proactive guidance for insight generation. At larger data scales, aggregation and filtering mechanisms will be needed to maintain legibility across encodings; child-level comparison across sessions is a natural extension of the current session-level design. Developing alternative visual encodings that better represent simultaneity and hierarchical musical structure remains an open direction. On the methodological side, longitudinal studies are needed to understand how visualization supports full composition workflows over time. Expanding participation to include professional composers and therapists will further ground the system in real-world contexts.

7 Conclusion

ResonaVis demonstrates the potential of visualization-driven approaches to support composers in reasoning about music for therapeutic contexts by making key elements—such as rhythm, dynamics, pitch, spectral energy, and brightness—more interpretable. Through interactive visualizations, the system surfaced patterns that enabled participants to reflect on compositional decisions, including aligning rhythmic and dynamic structures, balancing spectral properties, and designing smoother transitions to achieve cohesive musical outcomes. While these observations are consistent with prior work on therapeutic music, our contribution lies in showing how such relationships can be revealed and operationalized through visual analytics. More broadly, this work highlights how integrating interaction and audio data can support composer-centered workflows that move beyond passive analysis toward active exploration and decision-making. Future work will extend this approach by incorporating richer analytical methods, refining visualization designs, and exploring additional musical dimensions.

Acknowledgments

This work was supported by the National Science Foundation via the Accelerating Research Translation Program, Grant No. 2331440, and UD Institute for Engineering-Driven Health. We would like to thank Connor Fone and Justin Eichenberg for co-designing the dashboard with us, Liam Stapley for revamping the UI, Dr. Kathleen McCoy for her advice, and Sensify Lab members for their feedback on the manuscript. ChatGPT, Grammarly, and Overleaf's autocorrect were used to improve grammar, clarity, organization, and formatting. All suggestions were reviewed and revised by the authors, who take full responsibility for the final content.

Supplemental Materials

Our Supplementary Material contains study design for both the studies, the visualizations and their descriptions, and conversation prompts which were used to facilitate better conversation with the participants. They are available at https://osf.io/xg6su.

References

[1] R. Allen, E. Hill, and P. Heaton. Hath charms to soothe...' an exploratory study of how high-functioning adults with asd experience music. *Autism*, 13(1):21–41, 2009. doi: 10.1177/1362361307098511 2, 9

[2] S. Astrid Bin, F. Morreale, N. Bryan-Kinns, and A. P. McPherson. In-the-moment and beyond: Combining post-hoc and real-time data for the study of audience perception of electronic music performance. In *IFIP conference on human-computer interaction*, pp. 263–281. Springer, 2017. doi: 10.1007/978-3-319-67744-6_18 2

[3] A. Bangor, P. T. Kortum, and J. T. Miller. An empirical evaluation of the system usability scale. *Intl. Journal of Human–Computer Interaction*, 24(6):574–594, 2008. doi: 10.1080/10447310802205776 6, 9

[4] J. B. Begole, J. C. Tang, and R. Hill. Rhythm modeling, visualizations and applications. In *Proceedings of the 16th Annual ACM UIST '03*, pp. 11–20. ACM, New York, 2003. doi: 10.1145/964696.964698 3

[5] J. Bernard, D. Sessler, T. May, T. Schlomm, D. Pehrke, and J. Kohlhammer. A visual-interactive system for prostate cancer cohort analysis. *IEEE computer graphics and applications*, 35(3):44–55, 2015. doi: 10.1109/MCG.2015.49 2

[6] G. Bharathi, A. Venugopal, and B. Vellingiri. Music therapy as a therapeutic tool in improving the social skills of autistic children. *The Egyptian Journal of Neurology, Psychiatry and Neurosurgery*, 55:1–6, 2019. doi: 10.1186/s41983-019-0091-x 2

[7] S. Björk, S. Lundgren, and J. Holopainen. Game design patterns. In *Proceedings of DiGRA 2003 Conference: Level Up*, 2003. doi: 10.26503/dl.v2003i1.60 2

[8] G. Born. On musical mediation: Ontology, technology and creativity. *Twentieth-century music*, 2(1):7–36, 2005. doi: 10.1017/S147857220500023X 2

[9] W.-Y. Chan, H. Qu, and W.-H. Mak. Visualizing the semantic structure in classical music works. *IEEE Trans. Visual Comput. Graphics*, 16(1):161–173, 2009. doi: 10.1109/TVCG.2009.63 3

[10] R. Chang, C. Ziemkiewicz, T. M. Green, and W. Ribarsky. Defining insight for visual analytics. *IEEE Computer Graphics and Applications*, 29(2):14–17, 2009. doi: 10.1109/MCG.2009.22 2

[11] Q. Chen, C. Mao, L. Qi, Y. Luo, G. Yang, L. Wang et al. Music-based therapeutic interventions for medical school students with emotional regulation and mental health: a pre-post cohort study. *Frontiers in Psychology*, 15:1401129, 2024. doi: 10.3389/fpsyg.2024.1401129 2

[12] F. L. Cibrian, J. Mercado, L. Escobedo, and M. Tentori. A step towards identifying the sound preferences of children with autism. In *Proceedings of the 12th EAI International Conference on Pervasive Computing Technologies for Healthcare*, PervasiveHealth '18, pp. 158–167. ACM, New York, 2018. doi: 10.1145/3240925.3240958 2

[13] V. Clarke and V. Braun. Thematic analysis. *The journal of positive psychology*, 12(3):297–298, 2017. doi: 10.1080/17439760.2016.1262613 6

[14] N. Collins. *Introduction to computer music*. John Wiley & Sons, 2010. doi: 10.1017/S1355771810000221 2

[15] J. L. Cybulski, S. Keller, L. Nguyen, and D. Saundage. Creative problem solving in digital space using visual analytics. *Computers in Human Behavior*, 42:20–35, 2015. Digital Creativity: New Frontier for Research and Practice. doi: 10.1016/j.chb.2013.10.061 2

[16] S. Daniel, D. Wimpory, J. T. Delafield-Butt, S. Malloch, U. Holck, M. Geretsegger et al. Rhythmic relating: Bidirectional support for social timing in autism therapies. *Frontiers in Psychology*, 13:793258, 2022. doi: 10.3389/fpsyg.2022.793258 2

[17] M. H. De Lima, D. Keller, M. S. Pimenta, V. Lazzarini, and E. M. Miletto. Creativity-centred design for ubiquitous musical activities: Two case studies. *Journal of Music, Technology and Education*, 5(2):195–222, 2012. doi: 10.1386/jmte.5.2.195_1 2

[18] R. De Prisco, D. Malandrino, D. Pirozzi, G. Zaccagnino, and R. Zaccagnino. Understanding the structure of musical compositions: Is visualization an effective approach? *Information Visualization*, 16(2):139–152, 2017. doi: 10.1177/1473871616655468 3

[19] J. C. De Winter. Using the student's" t"-test with extremely small sample sizes. *Practical assessment, research & evaluation*, 18(10):n10, 2013. doi: 10.7275/e4r6-dj05 9

[20] L. Faulkner. Beyond the five-user assumption: Benefits of increased sample sizes in usability testing. *Behavior Research Methods, Instruments, & Computers*, 35(3):379–383, 2003. doi: 10.3758/BF03195514 9

[21] R. Fiebrink, P. R. Cook, and D. Trueman. Human model evaluation in interactive supervised learning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pp. 147–156. ACM, New York, 2011. doi: 10.1145/1978942.1978965 2

[22] J. H. Fonteles, M. A. F. Rodrigues, and V. E. D. Basso. Creating and evaluating a particle system for music visualization. *Journal of Visual Languages Computing*, 24(6):472–482, 2013. doi: 10.1016/j.jvlc.2013.10.002 3

[23] J. Foote. Visualizing music and audio using self-similarity. In *Proceedings of the Seventh ACM MULTIMEDIA '99*, pp. 77–80. ACM, New York, 1999. doi: 10.1145/319463.319472 2

[24] A. Franklin, S. Gantela, S. Shifarraw, T. R. Johnson, D. J. Robinson, B. R. King et al. Dashboard visualizations: Supporting real-time throughput decision-making. *Journal of Biomedical Informatics*, 71:211–221, 2017. doi: 10.1016/j.jbi.2017.05.024 2

[25] M. Geretsegger, C. Elefant, K. A. Mössler, and C. Gold. Music therapy for people with autism spectrum disorder. *Cochrane Database of Systematic Reviews*, (6), 2014. doi: 10.1002/14651858.CD004381.pub3 1, 2

[26] C. Gold, M. Voracek, and T. Wigram. Effects of music therapy for children and adolescents with psychopathology: a meta-analysis. *Journal of Child Psychology and Psychiatry*, 45(6):1054–1063, 2004. doi: 10.1111/j.1469-7610.2004.t01-1-00298.x 1

[27] M. Goto. Smartmusickiosk: music listening station with chorus-search function. In *Proceedings of the 16th Annual ACM UIST '03*, pp. 31–40. ACM, New York, 2003. doi: 10.1145/964696.964700 3

[28] J. C. Greene. Stakeholder participation in evaluation design: Is it worth the effort? *Evaluation and Program Planning*, 10(4):379–394, 1987. doi: 10.1016/0149-7189(87)90010-3 3

[29] F. Grond, R. Motta-Ochoa, N. Miyake, T. Tembeck, M. Park, and S. Blain-Moraes. Participatory design of affective technology: Interfacing biomusic and autism. *IEEE Transactions on Affective Computing*, 13(1):250–261, 2022. doi: 10.1109/TAFFC.2019.2922911 3

[30] Y. Han, G. D. Abowd, and J. Stasko. Intivisor: A visual analytics system for interaction log analysis. *IEEE Trans. Visual Comput. Graphics*, 31(3):1772–1784, 2025. doi: 10.1109/TVCG.2024.3370637 2

[31] D. J. Hargreaves, R. MacDonald, and D. Miell. Musical identities mediate musical development. *Music and music education in people's lives*, 1:124–142, 2018. doi: 10.1093/oxfordhb/9780199730810.013.0008_update_001 2

[32] S. G. Hart. Nasa-task load index (nasa-tlx); 20 years later. In *Proceedings of the human factors and ergonomics society annual meeting*, vol. 50, pp. 904–908. Sage publications Sage CA: Los Angeles, CA, 2006. doi: 10.1177/154193120605000909 5

[33] W. Henry. Creative processes in children's musical compositions: A review of the literature. *Update: Applications of Research in Music Education*, 15(1):10–15, 1996. doi: 10.1177/87551233960150010 2

[34] E. Hicks, J. Johnson, R. W. Black, and G. R. Hayes. Understanding noise sensitivity through interactions in two online autism forums. In *Proceedings of the ACM SIGACCESS ASSETS '23*, art. no. 19, 12 pp. ACM, New York, 2023. doi: 10.1145/3597638.3608413 2, 9

[35] R. Hiraga, R. Mizaki, and I. Fujishiro. Performance visualization: a new challenge to music through visualization. In *Proceedings of the Tenth ACM MULTIMEDIA '02*, pp. 239–242. ACM, New York, 2002. doi: 10.1145/641007.641054 3

[36] A. Karwankar, E. Ruggiero, Z. Lipkin, M. K. Iyer, S. Brugel, P. Khatiwada

et al. *uCue: An Interactive Musical Interface to Enhance Formative Listening Experiences for Children with ASD*, pp. 340–357. ACM IDC '25, New York, 2025. doi: 10.1145/3713043.3727053 2, 3

[37] J. Katagiri. The effect of background music and song texts on the emotional understanding of children with autism. *Journal of music therapy*, 46(1):15–31, 2009. doi: 10.1093/jmt/46.1.15 2

[38] D. Keim, G. Andrienko, J.-D. Fekete, C. Görg, J. Kohlhammer, and G. Melançon. Visual analytics: Definition, process, and challenges. In *Information visualization: Human-centered issues and perspectives*, pp. 154–175. Springer, 2008. doi: 10.1007/978-3-540-70956-5_7 2

[39] O. Khyzhna and K. Shafranska. Music therapy as an important element in shaping communication competences in children with autism spectrum disorder. *Journal of History Culture and Art Research*, 9(3):106–114, 2020. doi: 10.7596/taksad.v9i3.2823 2

[40] J. Kim, T. Wigram, and C. Gold. The effects of improvisational music therapy on joint attention behaviors in autistic children: a randomized controlled study. *Journal of autism and developmental disorders*, 38:1758–1766, 2008. doi: 10.1007/s10803-008-0566-6 3

[41] A. B. LaGasse. Effects of a music therapy group intervention on enhancing social skills in children with autism. *Journal of music therapy*, 51(3):250–275, 2014. doi: 10.1093/jmt/thu012 2

[42] A. B. LaGasse. Social outcomes in children with autism spectrum disorder: a review of music therapy outcomes. *Patient related outcome measures*, pp. 23–32, 2017. doi: 10.2147/PROM.S106267 2

[43] H. Lam, E. Bertini, P. Isenberg, C. Plaisant, and S. Carpendale. Empirical studies in information visualization: Seven scenarios. *IEEE Trans. Visual Comput. Graphics*, 18(9):1520–1536, 2012. doi: 10.1109/TVCG.2011.279 9

[44] J. Lazar, J. H. Feng, and H. Hochheiser. *Research methods in human-computer interaction*. Morgan Kaufmann, 2017. doi: 10.1016/B978-0-12-805390-4.09987-8 4

[45] H. B. Lima, C. G. R. D. Santos, and B. S. Meiguins. A survey of music visualization techniques. *ACM Comput. Surv.*, 54(7), art. no. 143, 29 pp., July 2021. doi: 10.1145/3461835 3, 9

[46] M. Mauriello, D. Stevens, E. Ruggiero, and A. Karwankar. Interactive musical interface to enhance formative listening experiences for children with autism spectrum disorder, Aug. 21 2025. US Patent App. 19/201,395. doi: patent/US20250265035A1/en 3

[47] M. L. Mauriello, B. Shneiderman, F. Du, S. Malik, and C. Plaisant. Simplifying overviews of temporal event sequences. In *Proceedings of the 2016 CHI EA on Human Factors in Computing Systems*, CHI EA '16, pp. 2217–2224. ACM, New York, 2016. doi: 10.1145/2851581.2892440 2

[48] H. Mayer-Benarous, X. Benarous, F. Vonthron, and D. Cohen. Music therapy for children with autistic spectrum disorder and/or other neurodevelopmental disorders: a systematic review. *Frontiers in psychiatry*, 12:435, 2021. doi: 10.3389/fpsyt.2021.643234 2

[49] S. McAdams. Musical timbre perception. *The psychology of music*, 3, 2013. doi: 10.1016/B978-0-12-381460-9.00002-X 2

[50] K. S. McFerran and D. Rickson. Community music therapy in schools: Realigning with the needs of contemporary students, staff and systems. *International Journal of Community Music*, 7(1):75–92, 2014. doi: 10.1386/ijcm.7.1.75_1 2

[51] L. Michel, C. Ricou, F. Bonnet-Brilhault, E. Houy-Durand, and M. Latinus. Sounds pleasantness ratings in autism: Interaction between social information and acoustical noise level. *Journal of Autism and Developmental Disorders*, 54(6):2148–2157, 2024. doi: 10.1007/s10803-023-05989-6 2, 9

[52] M. Miller, D. Fürst, H. Hauptmann, D. A. Keim, and M. El-Assady. Augmenting digital sheet music through visual analytics. In *Computer Graphics Forum*, vol. 41, pp. 301–316. Wiley Online Library, 2022. doi: 10.1111/cgf.14436 3

[53] M. Monroe, R. Lan, H. Lee, C. Plaisant, and B. Shneiderman. Temporal event sequence simplification. *IEEE Trans. Visual Comput. Graphics*, 19(12):2227–2236, 2013. doi: 10.1109/TVCG.2013.200 2

[54] M. J. Muller and S. Kuhn. Participatory design. *Commun. ACM*, 36(6):24–28, June 1993. doi: 10.1145/153571.255960 3

[55] T. Munzner. Visualization analysis and design. In *Proceedings of the SIGGRAPH Courses '25*, art. no. 10, 2 pp. ACM, New York, 2025. doi: 10.1145/3721241.3733989 2

[56] D. M. Naranjo, J. R. Prieto, G. Moltó, and A. Calatrava. A visual dashboard to track learning analytics for educational cloud computing. *Sensors*, 19(13):2952, 2019. doi: 10.3390/s19132952 2

[57] J. Nielsen. *Usability Engineering*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1994. doi: doi/10.5555/2821575 4

[58] J. Nielsen and T. K. Landauer. A mathematical model of the finding of usability problems. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems*, CHI '93, pp. 206–213. ACM, New York, 1993. doi: 10.1145/169059.169166 9

[59] G. Norman. Likert scales, levels of measurement and the "laws" of statistics. *Advances in health sciences education*, 15(5):625–632, 2010. doi: 10.1007/s10459-010-9222-y 9

[60] C. North. Toward measuring visualization insight. *IEEE Computer Graphics and Applications*, 26(3):6–9, 2006. doi: 10.1109/MCG.2006.70 2

[61] G. Paine. Towards unified design guidelines for new interfaces for musical expression. *Organised Sound*, 14(2):142–155, 2009. doi: 10.1017/S1355771809000259 2

[62] P. Pirolli and S. Card. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of international conference on intelligence analysis*, vol. 5, pp. 2–4. McLean, VA, USA, 2005. 2

[63] A. E. Reschke-Hernández. History of music therapy treatment interventions for children with autism. *Journal of Music Therapy*, 48(2):169–207, 2011. doi: 10.1093/jmt/48.2.169 1

[64] N. SANTOS, G. BERNARDES, R. COTTA, N. COELHO, and A. BAGANHA. Assessing musical preferences of children on the autistic spectrum: Implications for therapy. 2024. doi: 10.13140/RG.2.2.17204.67207 2, 9

[65] A. Sarikaya, M. Correll, L. Bartram, M. Tory, and D. Fisher. What do we talk about when we talk about dashboards? *IEEE Trans. Visual Comput. Graphics*, 25(1):682–692, 2019. doi: 10.1109/TVCG.2018.2864903 2

[66] R. K. Sawyer and D. Henriksen. *Explaining creativity: The science of human innovation*. Oxford university press, 2024. doi: 10.1093/oso/9780197747537.001.0001 2

[67] M. Sedlmair, M. Meyer, and T. Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE Trans. Visual Comput. Graphics*, 18(12):2431–2440, 2012. doi: 10.1109/TVCG.2012.213 5

[68] B. Shneiderman. The eyes have it: a task by data type taxonomy for information visualizations. In *Proceedings 1996 IEEE Symposium on Visual Languages*, pp. 336–343, 1996. doi: 10.1109/VL.1996.545307 2

[69] B. Shneiderman. Designing for fun: how can we design user interfaces to be more fun? *Interactions*, 11(5):48–50, Sept. 2004. doi: 10.1145/1015530.1015552 4

[70] H. Smith. *Practice-led research, research-led practice in the creative arts*. Edinburgh University Press, 2009. doi: stable/10.3366/j.ctt1g0b594 2

[71] G. Tsiris, N. Spiro, O. Coggins, and A. Zubala. The impact areas questionnaire (iaq): a music therapy service evaluation tool. In *Voices: A World Forum for Music Therapy*, vol. 20. GAMUT-Grieg Academy Music Therapy Research Centre (NORCE & University of Bergen), 2020. doi: 10.15845/voices.v20i2.2816 5

[72] P. Vlachogianni and N. Tselios. Perceived usability evaluation of educational technology using the system usability scale (sus): A systematic review. *Journal of Research on Technology in Education*, 54(3):392–409, 2022. doi: 10.1080/15391523.2020.1867938 5

[73] J. Whipple. Music in intervention for children and adolescents with autism: A meta-analysis. *Journal of music therapy*, 41(2):90–106, 2004. doi: 10.1093/jmt/41.2.90 2

[74] T. Wigram. *Improvisation: Methods and techniques for music therapy clinicians, educators, and students*. Jessica Kingsley Publishers, 2004. doi: books?id=k9EPBQAAQBAJ 3

[75] T. Wigram and C. Gold. Music therapy in the assessment and treatment of autistic spectrum disorder: clinical application and research evidence. *Child: care, health and development*, 32(5):535–542, 2006. doi: 10.1111/j.1365-2214.2006.00615.x 2

[76] K. Wilkie, S. Holland, and P. Mulholland. Towards a participatory approach for interaction design based on conceptual metaphor theory: a case study from music interaction. In *Music and Human-Computer Interaction*, pp. 259–270. Springer, 2013. doi: 10.1007/978-1-4471-2990-5_15 3

[77] A. Xambó and G. Roma. Human–machine agencies in live coding for music performance. *Journal of New Music Research*, 53(1-2):33–46, 2024. doi: 10.1080/09298215.2024.2442355 2

[78] Y. Xing, M. T. Ortoleva, R. Borgo, and A. Abdul-Rahman. Collaborating across domains and roles: An interview study of visualization design practices. *IEEE Trans. Visual Comput. Graphics*, 32(1):571–581, 2026. doi: 10.1109/TVCG.2025.3634711 3