Sonic Serendipity: Embracing Discovery in File Finder-Based Improvisation

Austin Franklin
Mälardalen University
Västerås, Sweden
austin.franklin@mdu.se
The Royal College of Music
Stockholm, Sweden
austin.franklin@kmh.se

Rikard Lindell
Mälardalen University
Västerås, Sweden
rikard.lindell@mdu.se
Dalarna Univerity
Falun, Sweden
rli@du.se

Henrik Frisk
The Royal College of Music
Stockholm, Sweden
henrik.frisk@kmh.se

ABSTRACT

This paper describes the design and development of The Unfinder, a prototypical interface that allows users to search and improvise music with audio files from a large local repository using music information retrieval based on content and metadata. The research is part of an ongoing project (IRESAP) concerned with tools incorporating current music information retrieval strategies that both support artistic practices and have utility outside of a performance setting. In The Unfinder, we aim to exploit the balance between accurately and reliably retrieving audio material from a file search system, and the potential for failure in the system to do so. Our prior research is used to frame design choices which are measured with a user study used to evaluate the interface. In the study we observed nine users of varying musical backgrounds playing with the instrument while taking notes of their utterances and ideas. The transcriptions of the user's comments were analyzed using a thematic analysis method and five (5) themes were identified: perception, parameterization, identity, agency, and imaginaries. These themes indicate that the interface design is promising for artistic output, the use of a single feature for searching does not have much perceptual relevance, and the chosen features are useful for discovering audio files within serendipitous musical situations.

Author Keywords

Music Information Retrieval, File Organization, Interaction Design, Improvisation, Aesthetics

CCS Concepts

•**Human-centered computing** → **Interactive systems and tools;** •**Applied computing** → *Performing arts;* •**Information systems** → *Personalization;*

1. INTRODUCTION

A computer's interface is designed for creating, storing, editing, and finding text files. However, electronic musicians often store thousands of audio files in local repositories. Searching for audio files is difficult since structuring audio material is a personal task that can rely on aspects such as the tool used for creation, the date and time of creation, content-based features, and metadata. Additionally, the relationship between materials and the user's memory can change over time: materials can mature. Furthermore, a user can have many different needs with the material in the context of an artistic practice. These may include finding, creating, uncovering, controlling, or changing, to name a few. Despite this growing need, electronic filing and finding practices in computer operating systems have not changed much since 1995 [13].

In the paper Finding and Reminding: File Organization from the Desktop, Barreau and Nardi suggest that "the way information is used is a primary determinant of how it will be organized, stored, and retrieved" [2]. They write that the user's practice has a bigger impact on strategies than the design of the system. At the same time, Cascone suggests that the medium is no longer the message; rather, specific tools themselves have become the message since they support the growing number of artistic practices and aesthetics [6]. This understanding of tools connects with Berleants concept of aesthetic sensibility as it emphasizes engagement with the environment for shaping our experience while fostering perceptual awareness that is "developed, guided, and focused" [3]. As a tool can be considered an environment, a user's aesthetics play a vital role in shaping their practices with the tool. Furthermore, we believe that current music information retrieval methods exhibit potential for shaping musical experience and supporting aesthetic needs. They are also useful for cataloging media content in large repositories [1].

Early post-digital aesthetics were concerned with failure in digital systems [6]. Priest describes failure as "a logic wherein the absolute failure to satisfy an already determined purpose coincides with the success to satisfy an unintentional objective" [18]. Failure is often seen as having a purely indeterminate nature and is a fact of life for interactive systems as well as a resource for aesthetic and improvisational performance. Ravasio distinguishes between two types of failure in improvised musical performance: first- and second-order [19]. First-order failures happen for many reasons, including sheer bad luck – a drumstick breaks or a string snaps, for example. This, in our case, constitutes failure on behalf of the tool itself, or the inability of the tool to do what the user expects. Second-order failures occur within the context of a musical idiom or aesthetic, when an improviser abandons the "well-trodden path" and begins experimenting with the musical material. For example, if an improviser makes an attempt to change musical style but is unsuccessful, this constitutes a second-order failure.

Hazard et al. propose that creative and performative in-

NIME'24, 4–6 September, Utrecht, The Netherlands.

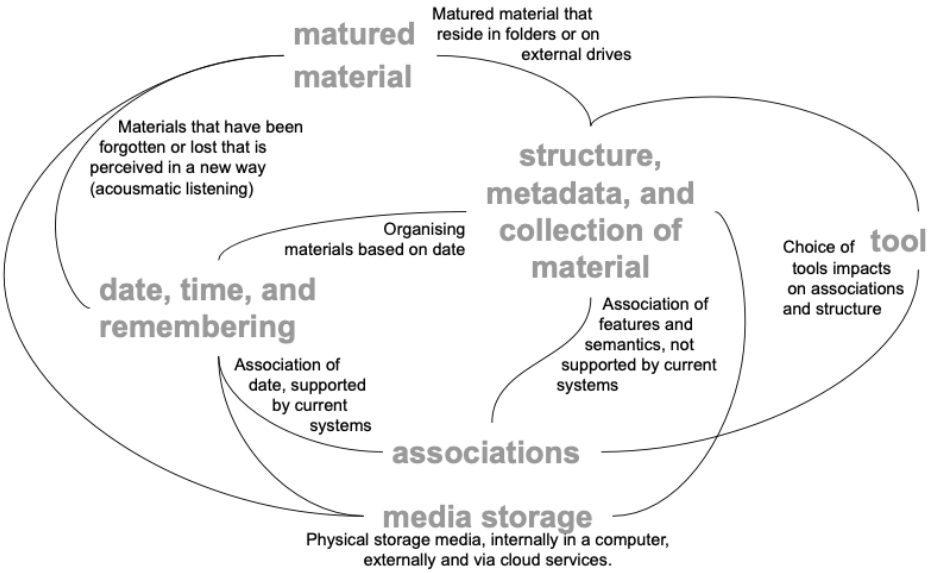

Figure 1: A relationship of themes based on how electroacoustic composers organize media material

terfaces should be designed to enable aesthetic failures and introduce a taxonomy that compares human approaches to failure with approaches to capable systems, revealing new creative strategies [10]. The taxonomy is a form of design knowledge that includes five (5) points: 1) taming the system; 2) gaming the system; 3) negotiating (with) the system; 4) serving the system; and 5) riding the system. In practice, improvisers will often seek a balance between these points, for example, negotiating with the system before giving up and serving the system, or moving between them as the need arises. When riding the system, the improviser "gives themselves up to both failure and the system, accepting its consequences and improvising in response." With this serendipitous approach, the system becomes a mechanism for discovering and accessing music even in cases where the user does not explicitly know what he or she is actually looking for. Alongside the aesthetics of failure, Hamilton's characterization of the "aesthetics of imperfection" [9] in an improvised performance is described as a "constant striving for new contingencies to respond to." In other words, failure is not merely tolerated in an improvised performance but is, in some cases, a valuable asset.

Based on the concepts presented here, the research question can be thought of as two-part: How may a local file organization and retrieval tool for media operate in the context of today's methods and practices, and what design features support these methods and practices? In the context of current music information retrieval methods, achieving certain aesthetic and system goals (failure, for example) is by no means a simple task. It may involve many things but should, based on Hazard et al.s' proposal [10], rely on indeterminacy that may be achieved as a result of inaccurate feature extraction or a low or inaccurate degree of perceptual correlation between analysis and audio files. But it could equally well rely on an analysis that has a far more precise resolution than human perception. With The Unfinder, we aim to exploit the balance between these two aspects: accurately and reliably retrieving audio material from the system, and the potential for failure in the system and what that failure may lead to as part of a broader artistic practice.

2. BACKGROUND

This research presents a new tool for organization and retrieval of audio content called The Unfinder and the findings of a user study on this tool. Design choices for The Unfinder are based on a previous pilot study that sought to understand how electroacoustic composers organize media material [13]. In this study, participants were asked to keep a journal during the preparation of an improvisation and reflect on how the material is being organized. Using Braun and Clarkes' thematic analysis method [4], six (6) themes were generated from this data: 1) storage media; 2) date, time, and remembering; 3) matured material; 4) structure, metadata, and collection of material; 5) associations; and 6) tool. Then, a model of the themes (Figure 1) was built showing the relationships between them.

For the present study we focused on the associations, date, time, and remembering, and the matured materials themes. These themes are about various associations between media, files, and other aspects of artistic production. As we are interested in both perceptual characteristics and design considerations, we have largely explored query-by-example systems [20]. There are several existing tools for organizing audio content such as FluCoMa [22], MIRLCRep [23], SoundTorch [11], StockSynth [21], the Minimal-Impact Personal Archive [7], and the Freesound.org API [8], to name a few. While these tools are focused on retrieval as their central thesis, we believe, like Barreau and Nardi [2], that the users' practice is a vital component that provides design considerations for this problem. Content-based systems are difficult to build due to the fact that perception varies with age and culture, among other things [14]. Traditional methods for finding audio files (e.g., text queries, directory lists, etc.) are also time consuming and tedious, and automatic annotation is currently not mature enough [5]. The design of the system, however – how the user interacts with the materials – also impacts the efficacy of the tool.

3. THE UNFINDER

The Unfinder (Figure 2) affords the ability to 1) quickly recall audio files during a performance, 2) sort them using some feature or metadata attribute, and 3) find audio

files outside of a performance setting for use in other applications. The found audio files, up to four, begin playing automatically. To use The Unfinder, a user first selects an audio file from a previously loaded folder containing audio material. From here, the waveform of the audio file is displayed with the overlaid feature data for that file. The user then uses a window (the low and high range in Figure 2) to select a portion they deem sonically important. When the window changes, a similarity measure is calculated using the absolute difference of the mean of the windowed area. The four most similar audio files are returned to the user and can be performed using a variety of additional controls and effects.

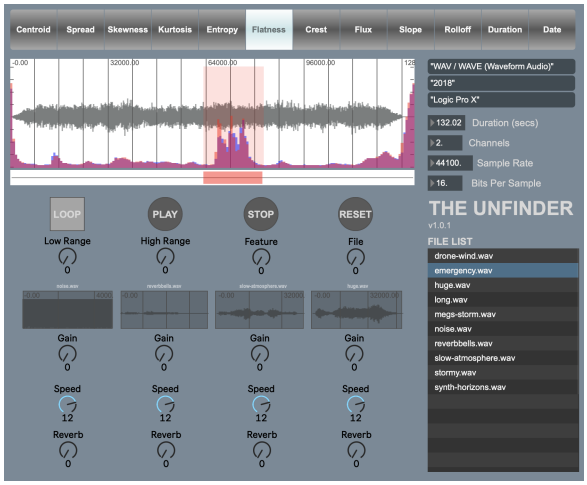

Figure 2: The Unfinder prototypical digital interface

The interface is designed to be controlled using a 4x4 grid of knobs. In particular, the EN16 controller by Intech Studios (Figure 3) was used. Two of the knobs select the file and feature while another two select the low and high position of the window used for the search. The remaining knobs control volume, playback speed and pitch, and reverb for each of the 4 returned channels.

The analysis is done using FFmpeg and FFprobe to gather spectral stats and metadata attributes. All spectral features are calculated using an FFT frame size of 65536 samples which are then written into .json files. In Max, the *foldobject* loads all .json files and the *dict* object parses the files to obtain values for keys needed for calculation. The interface provides the following ten (10) sonic features: centroid, spread, skewness, kurtosis, entropy, flatness, crest, flux, slope, and rolloff; and the following two (2) metadata attributes: duration and date. Duration and date return audio files closest in length specified by the window in milliseconds or closest in date using the creation date from the files' metadata. Additionally, other metadata is displayed, such as the `encoded_by` key, which pertains to the tool used to encode the file, such as Reaper or Pro tools for example. The sonic features relate to the previous study's associations theme, and the duration and date metadata attributes were inspired by the date, time, and remembering and the matured materials themes of the previous study.

3.1 Designing for Failure

The Unfinder allows users to curate their own repositories as they interact with the system (Figure 4). Many audio analysis tools contain pre- and post-processing methods for filtering sections of audio that skew analysis, typically using an amplitude threshold to remove silences or rate of change

Figure 3: The EN16 controller by Intech Studios

to detect and segment transients. When a user selects a feature, only that feature is used to calculate the similarity. When the window changes, the similarity is recalculated, returning new audio files. At the same time, the window for the audio file is stored for that feature and file. If it becomes one of the four most similar audio files in a new search, only that portion of the file is returned to the user for playback. This allows a user to adjust the specificity level within the interface by removing beginning and ending silences and isolating sections and makes performing or browsing a much more personal task.

Rather than automatic segmentation or pre-filtering, we believe this curation method strikes the balance between retrieval and the potential for failure by relying on a user's practice with the material. It allows them to tame the system and negotiate with the system, or to approach the system in a personalized manner, as one would with an acoustic instrument [10]. For example, setting the window too wide or narrow may cause crucial spectral data to become noticed or ignored by the system, which would skew the similarity measure. Moreover, if a user makes sonic connections between different files when setting the window that are not understood by the given feature, then this too will return unexpected and surprising results. The intention is for users to communicate their desires to the system and accept the results if the desires are not properly met. Finally, using features which do not possess much perceptual correlation (i.e., objective algorithms [24]) as well as metadata attributes can be an interesting and rewarding artistic challenge in an improvisational setting.

4. USER STUDY

The design of The Unfinder and its implementation of music information retrieval tools were evaluated in a user study. We carried out the study at the Royal College of Music in Stockholm with nine users of varying musical backgrounds and experience levels, ranging from music enthusiasts to professional music composers, producers, researchers, and students, thus all participants had some level of musical training and proficiency. In particular, we identified the following three (3) research questions as most pertinent:

1. Does the design encourage artistic engagement?
2. Do the chosen features have perceptual meaning?
3. Are the chosen features useful for finding sounds in a repository based on a currently selected sound?

These questions were identified based on the primary conceptual and design features of The Unfinder discussed above. Although they were not asked explicitly, they were central

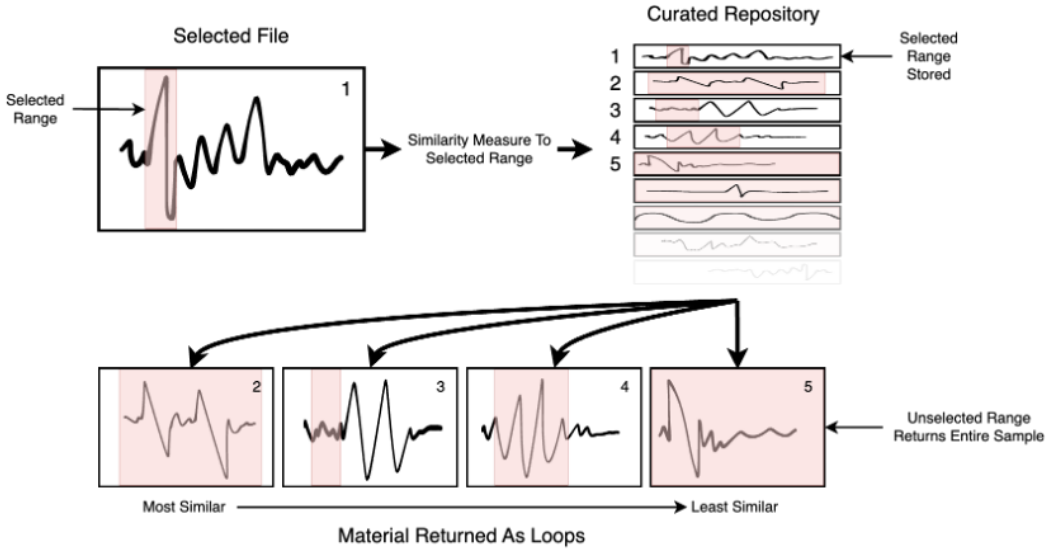

Figure 4: The personal curation process in The Unfinder

to our investigation of each users' experience. Each participant was given a brief demonstration of the basic controls, after which we instructed them to engage with The Unfinder for approximately 15 minutes. During this time we recorded handwritten notes of their comments and other notable moments from their engagement and asked questions in a semi-structured interview format. These notes reflect comments on three (3) research questions relating to The Unfinder's design. The notes were typed up immediately after the study was completed in order not to lose recollection or context in the interpretation of the data. The transcriptions were analyzed using the thematic analysis method [4].

The materials used for this study were limited to 256 audio files with a total volume of 671 megabytes from our own repositories. Although they have inherent aesthetic value from our own artistic practices, the files were not curated and could equally well represent a backup SD card or a disc.

4.1 Findings

Thematic analysis is a pragmatic method to guide and organize the interpretation of data performed in six (6) steps: (1) become familiar with the data, (2) generate codes, (3) search for themes, (4) review themes, (5) define themes, and (6) write-up [15]. While transcribing the data, we made some initial codes. The codes were further developed and formed into themes. Each researcher worked independently and then discussed and merged our analyses once a consensus was reached. The following revised themes describe recurring motifs throughout users' engagement with the interface: perception, parameterization, identity, agency, and imaginaries.

4.1.1 Perception

Perception refers to the user's ability to make a connection between the feature and the sound. The data suggests that the tool's use of individual features do not support perceptual similarities between different audio files. For instance, user eight (U8), an electronic music composition student enquired how the four audio files were similar to the selected sound. "It's difficult to hear the similarity". In this case the user was using the default spectral centroid parameter. User three (U3), a novice with no experience with music technology, stated: "Centroid does not seem to be similar," indicating the discrepancy between similarity based on the selected feature and the perceptual similarity. From the observations and the utterances from the users it appears to be difficult for them to hear similarities between audio files determined to have similarity by The Unfinder.

4.1.2 Parameterization

Parameterization refers to content and metadata features and their ability to communicate sonic characteristics to users through description. It also includes interface design considerations that may impact the usability or ease-of-learning of the tool. Several users demonstrated difficulty connecting a sonic characteristic to the name of the feature. User four (U4), a composition student, asked "What is in flux?" while U8 asked "what is the flux value?" However, U8 only used the centroid feature and did not explore other search parameters. A different word for each feature may be better suited to communicate features' unique characteristics (i.e., brightness for centroid, for example). However, there are many factors, such as the specific content of the repository and how engaged a user was with multiple features and/or files, that could have contributed to this problem. U3 expressed clarity, stating: "Aha, when I change this [low and high range] it's completely different," while user seven (U7), another composition student, realized how to change the set of found audio files using several features such as skewness and duration.

4.1.3 Identity

Identity refers to The Unfinder's concept and its dual identity as a tool for retrieval and improvisation which made the initial confrontation with the tool somewhat difficult to understand. U3 asked: "Is it a composition tool or a search tool?" However, given the time constraint of the study, most users treated the tool as a musical instrument. U7 remarked "It's more of an instrument than a file-finder," which was echoed by U8's comment of the tool being "a joyful mixture of loops, one becomes immersed." Furthermore, user nine (U9) preferred using more distinct audio files, ones deemed less perceptually similar by The Unfinder (channels 3 and 4) for performance.

4.1.4 Agency

Agency refers to the affordance of the tool as well as how much of the artistic process is determined by the user versus the tool. U3 said: "It is a lot of fun but I do not know how much I am controlling it or if it just plays on its own." This indicates that for this particular user the tool appeared to have a life of its own. Later, U3 found out a relationship between actions in the system and the produced result, while user five (U5), a composition student, also had an "Aha!" moment after learning the mute function. Their distinct alterations of low and high positions for the parameters resulted in a distinct response of selected audio files. User six (U6), a composition student uttered: "Now I got this" while playing with the speed controller for the returned audio files. Although the connection between parameter similarity or perceived similarity was feeble we observed that the adept musicians and composers in this study were able to artistically express themselves with the tool. For instance, U9 played the volume faders musically, whereas the reverb was used more sparingly.

4.1.5 Imaginaries

The imaginaries theme emerged based on the user's imagination and refers to the tool and its potential integration into a user's workflow or to what impedes integration. Although having initial difficulty with the design, user two (U2), a recording engineer, suggested a potential design for creating music integrated with the sound information retrieval tool. U2 first stated: "[This is a] hybrid looper. [...] I think mostly about music in linear stems in a DAW. I would have to adapt to Ableton or this." At the same time U2 suggests that collections of audio files could be saved to a grid: "Save a set to a grid of pads [with a] quick access to make something." Similarly, user one (U1), a music researcher and audio expert, suggested a button to save states that enable you to play the selected states.

5. DISCUSSION

We set out to explore the design of a music information retrieval instrument that can both find audio files in personal media collections based on ten sonic features and two metadata attributes and at the same time encourage artistic engagement. The theme perception strongly indicates that users largely did not perceive perceptual similarities between a selected sound and the four audio files that the system found to be the most similar in the dimension of the selected sonic feature. This is most likely due to the design of the tool only using a single feature for comparison at a time, which does not create a strong enough perceptual correlation. We consider this a first-order failure, a limitation of the tool using a single feature for analysis and similarity.

Previously, Malt and Jourdan have described the problem of using descriptors in music applications as three-fold: 1) The lack of knowledge of the relationships between descriptors and the pertinent perceptual characteristics of the sound for use in musical composition; 2) The fact that one descriptor is not always sufficient in order to characterize a complex sound; 3) The lack of a large choice of descriptors for artists to test and learn to use them [16]. While the first two points seem to hold true in this study, the large choice of descriptors and naming conventions used in The Unfinder likely contributed to users' confusion while trying to determine sonic similarity and learn the controls. It is unclear if this constitutes a failure on behalf of the tool, given the limited scope of the study and the length of time each user was engaged with The Unfinder. With more time and more personalized audio material this outcome may have been different.

The theme parameterization indicates that the features and their names are somewhat confusing, which also connects to Malt and Jourdans' first problem. Nevertheless, in the themes identity and agency we have found a strong indication that The Unfinder produces an artistically interesting result. Interestingly, the first-order failures of the tool with regards to the perception and parameterization themes appear to create second-order failures with respect to the identity and agency themes. Within the agency theme specifically, the tool often appears to make decisions on behalf of its user, abandoning the "well-trodden path" or changing direction in unexpected ways. In this context, The Unfinder enables serving and riding the system rather than taming or negotiating with it and turns the tool into more of an instrument for improvisation. We have come to consider the tool's curation method a strategy for taming the system and their input into the system as negotiating with the system rather than simply controlling it directly.

The theme imaginaries indicates Krippendorff's pragmatic validity because users engaged in the design, suggested improvements, and imagined how they would find the design useful in the context of their individual workflows [12]. It also suggests that prior experience with music technology, such as a digital audio workstation (DAW), informs the understanding of new designs and that reconsidering design choices from this perspective could greatly enhance the value of The Unfinder. Additional new design features could include new ways to visualize audio material selection by using a self organizing map (SOM), such as the one made available through FluCoMa [22]. Regardless of specific features or design improvements suggested by users, how can we make sure that The Unfinder maintains or even increases this serendipitous failure whilst decreasing the utilitarian failure? For example, a 'failure' parameter could be added that finds more or less correlative audio material when searching. However, this relies on explicit instruction from the user and therefore does not fall into the category of first-order failure. This question warrants further exploration and will be the subject of future study.

Users were unable to find specific sounds with The Unfinder because they lacked prior knowledge of the material. However, we confirmed users' artistic engagement with the tool. From an artistic point of view, dissimilarity is as interesting as similarity, which may sometimes warrant the use of a single feature or a less perceptually correlative one. On the other hand, the more capable an information retrieval system is, the less room for creatively stimulating failure potentially exists within it. The balance between finding what you want and what you did not know you wanted (or "Finding and Discovering" á la Barraeu and Nardi) is perhaps best understood as a complex interrelation between the themes described above and the users' aesthetic sensibilities.

6. CONCLUSIONS

The Unfinder is a digital interface that allows users to search and improvise with audio files from large local repositories using content and metadata-based approaches. The design of the tool is based on a previous study that identified users' practices. The Unfinder's development is outlined in this paper along with unique challenges related to personalization. We conducted a user study on The Unfinder to determine: 1) if the interface design is artistically engaging, 2) if the chosen features have perceptual meaning, and 3) if the chosen features are useful for finding sounds in a repos-

itory based on a currently selected sound. Our results indicate that the interface design encourages artistic engagement and the chosen features are useful for finding audio files within serendipitous musical situations. However, the use of a single feature for searching was not shown to have much perceptual correlation.

7. ACKNOWLEDGMENTS

Information Retrieval in Embedded Systems for Audiovisual Artistic Processes (IRESAP) is supported by The Knowledge Foundation (KKS). The first author developed The Unfinder tool and was instrumental in writing the first draft of the paper. All authors contributed to the design of The Unfinder and participated in the design of the user study. The first and second authors performed the user study and initial thematic analysis. All authors are accountable and were engaged in writing the submitted manuscript.

8. ETHICAL STANDARDS

The study did not handle any sensitive personal data. To use the data produced in the study, we have obtained the explicit consent from the participants. The protocol method was limited to hand written notes. The participants have certified voluntary participation, providing us the right to publish pseudonymized quotes, and the participants have certified awareness of the study's procedure. The data from the study were transcribed and stored pseudonymously. Quotations from our observations have been presented pseudonymously. We also fully comply with NIME Principles Code of Practice on Ethical Research [17].

9. REFERENCES

[1] V. Akkermans, F. Font Corbera, J. Funollet, B. De Jong, G. Roma Trepat, S. Togias, and X. Serra. Freesound 2: An improved platform for sharing audio clips. In *Klapuri A, Leider C, editors. ISMIR 2011: Proceedings of the 12th International Society for Music Information Retrieval Conference; 2011 October 24-28; Miami, Florida (USA). Miami: University of Miami; 2011.* International Society for Music Information Retrieval (ISMIR), 2011.

[2] D. Barreau and B. A. Nardi. Finding and reminding: file organization from the desktop. *ACM SigChi Bulletin*, 27(3):39–43, 1995.

[3] A. Berleant. Aesthetic sensibility. *Ambiances. Environnement sensible, architecture et espace urbain*, 2015.

[4] V. Braun and V. Clarke. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101, 2006.

[5] P. Cano and M. Koppenberger. Automatic sound annotation. In *Proceedings of the 2004 14th IEEE Signal Processing Society Workshop Machine Learning for Signal Processing, 2004.*, pages 391–400. IEEE, 2004.

[6] K. Cascone. The aesthetics of failure: Post-digital tendencies in contemporary computer music. *Computer Music Journal - COMPUT MUSIC J*, 24:12–18, 12 2000.

[7] D. P. Ellis and K. Lee. Minimal-impact audio-based personal archives. In *Proceedings of the the 1st ACM workshop on Continuous archival and retrieval of personal experiences*, pages 39–47, 2004.

[8] F. Font, T. Brookes, G. Fazekas, M. Guerber, A. La Burthe, D. Plans, M. D. Plumbley, M. Shaashua, W. Wang, and X. Serra. Audio commons: bringing creative commons audio content to the creative industries. In *Audio engineering society conference: 61st international conference: audio for games*. Audio Engineering Society, 2016.

[9] A. Hamilton. The aesthetics of imperfection. *Philosophy*, 65(253):323–340, 1990.

[10] A. Hazzard, C. Greenhalgh, M. Kallionpaa, S. Benford, A. Veinberg, Z. Kanga, and A. McPherson. Failing with style: Designing for aesthetic failure in interactive performance. In *Proceedings of the 2019 chi conference on human factors in computing systems*, pages 1–14, 2019.

[11] S. Heise, M. Hlatky, and J. Loviscach. Soundtorch: Quick browsing in large audio collections. In *Audio Engineering Society Convention 125*. Audio Engineering Society, 2008.

[12] K. Krippendorff. *The semantic turn: A new foundation for design*. crc Press, 2005.

[13] R. Lindell and H. Frisk. The unfinder: Finding and reminding in electronic music. In *the 16th International Symposium on Computer Music Multidisciplinary Research*, 2023.

[14] M. P. Lynch, R. E. Eilers, D. K. Oller, and R. C. Urbano. Innateness, experience, and music perception. *Psychological Science*, 1(4):272–276, 1990.

[15] M. Maguire and B. Delahunt. Doing a thematic analysis: A practical, step-by-step guide for learning and teaching scholars. *All Ireland Journal of Higher Education*, 9(3), 2017.

[16] M. Malt and E. Jourdan. Zsa. descriptors: a library for real-time descriptors analysis. In *5th Sound and Music Computing Conference, Berlin, Germany*, pages 134–137, 2008.

[17] F. Morreale, N. Gold, C. Chevalier, and R. Masu. Nime principles & code of practice on ethical research. 2023.

[18] E. Priest. *Boring formless nonsense: Experimental music and the aesthetics of failure*. Bloomsbury Publishing USA, 2013.

[19] M. Ravasio. Improvisational mistakes. *British Journal of Aesthetics*, page ayad031, 2023.

[20] J. Serra. Audio content-based music retrieval.

[21] S. Streich and B. S. Ong. A music loop explorer system. In *ICMC*, 2008.

[22] P. A. Tremblay, O. Green, G. Roma, J. Bradbury, T. Moore, J. Hart, and A. Harker. Fluid corpus manipulation toolbox. 2022.

[23] A. Xambó, G. Roma, A. Lerch, M. Barthet, and G. Fazekas. Live repurposing of sounds: Mir explorations with personal and crowdsourced databases. In T. M. Luke Dahl, Douglas Bowman, editor, *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 364–369, Blacksburg, Virginia, USA, June 2018. Virginia Tech.

[24] Y. Yuan, C. Cronin, D. Mullensiefen, S. Fujii, and P. E. Savage. Perceptual and automated estimates of infringement in 40 music copyright cases. 2023.