Visualization Research Center

Bachelorarbeit

Affective Sonification

Nicolas Caycedo

Course of Study: Softwaretechnik

Examiner: Prof. Dr. Daniel Weiskopf

Supervisor: Sita Vriend, Markus Wieland, Rene Cutura, Frank Heyen, Niklas Rönnberg

Commenced: January 2, 2023

Completed: July 2, 2023

Abstract

Defining emotions is still a highly debated topic, with new information being constantly brought to light. External stimuli often have an impact on how we feel and can sway our mood in ways that are not yet fully understood. Emotions and affect have in turn an impact on how we perform cognitively and how we interpret other stimuli. Extensive research has been done in the influence that musics and sounds have over emotion. In the field of Sonification, data is often represented abstractly with sounds that are not immediately recognizable. We set out to design a user study and its accompanying tools to better understand how the auditory aspect of an audiovisual data representation affects the recipient. During the course of our project we develop a software application integrating elements of participatory design to help lay people create sound representing data on a graph. We perform a pilot study and lay the foundation of a fully-fledged user study to be made in the future. This document describes our design process, challenges, and considerations during the planning of the user study and its corresponding software.

Contents

List of Figures

1 Introduction

While there is evidence of music having an influence on people's emotions [JV08], comparatively little research has been done on this aspect of sonification. Music is an important part of the sonification field, but it does not encompass all the whole of it. It is reasonable then to extrapolate that a sonification consisting of music alone has the same effects on a person's mood. Musical sonification has been successfully used to convey emotion, albeit in exchange for accuracy of data representation [Rön21]. We want to expand this assumption not only to the rest of the field, but also in a more natural setting.

Sonification is often found accompanied by a visualization of data or in situations where it would not be possible or practical to convey information through a visualization [Figure 1.2]. We focus here on the first scenario. The question is then, how does sound affect mood and emotion? We design a user study in which an audiovisual sonification is presented to a lay person, to whom we refer as user or participant. The reason for choosing someone with little to no experience with sonification is to help mitigate bias from previous knowledge. In particular, preconceptions such as which chords or sound characteristics are perceived as "happy" or "sad". The visual component contains information relevant enough to be relatable to the user, but neutral as to not have a significant impact on the user's affect. Our hope is that the sonification aspect of the experience will have the largest influence on the user's emotion and thus help us determine how effective sonification is for this kind of scenario.

A variety of software is available for music production and sound design, which is often used in sonification. These tools, however, do not offer the possibility to design a sonification experiences oriented towards emotions [Neu19]. In other words, because we lack the knowledge about which characteristics of sonification produces which emotion, there is no trivial way to select an emotion first and have the tool help the designer work towards that goal. We aim to provide developers of existing and future tools with this knowledge.

Another aspect of current options for sound design software is the complexity. A full-fledged sonification can become very complex depending on the task's requirements. Besides this, for a platform to successfully offer the vast variety of sounds and options to individually manipulate aspects of them for compositing it needs to have an equally complex interface, for example Figure 1.1. Users in our study have ideally no previous experience with such tools. This introduces a new limitation: it is unrealistic to expect the users to become experts during the study while still avoiding these biases that come with experience in the field. We created a new tool that finds a compromise between the complexity of a sound design suit and the facility required for lay people to be able to create the sound component of our audiovisual experience.

Finally, we take a Participatory Design approach to conducting the study. Users have a say in how the components of the interface are described. This helps not only them to get more comfortable in a foreign environment in which they have no previous experience, but also researchers to get a better understanding on how the users are perceiving the test. We provide the participants with

Figure 1.1: An example of a Digital Audio Workstation's graphical interface. Taken from rekkerd.org

the necessary tools to design an audiovisual sonification for themselves, but let them decide how these tools are best presented to them. This combination of user-centric and participatory design increases the participant's engagement and makes them feel more comfortable with the software [JKKS20] [MMP+16].

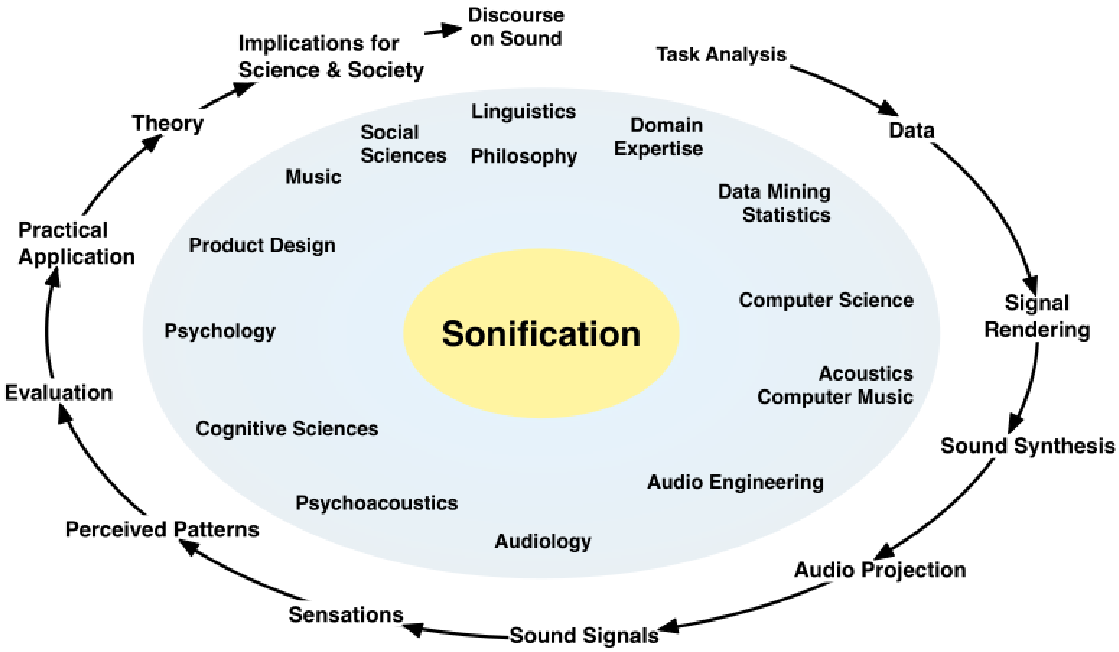

Figure 1.2: "The interdisciplinary circle of sonification and auditory display"[NHH+11]

2 State of the art

The field of Affective Sonification is still largely understudied. To better understand it, it is imperative to have at least some understanding of affective psychology and sonification itself. While the field of sonification is relatively new, we already have a good understanding of how important is the role of audio in everyday life. It is difficult to create a study that successfully encompasses all possible scenarios where a person is likely to find data represented by audio. Auditory displays and auditory icons, for example, are prevalent elements in modern lifestyle.

Regarding the psycological aspect of our research, the model of emotions determines the boundaries withing which we can interpret the collected data. Despite the long history of psychology in science and the study of emotions, there is currently no universal consensus in this regard [AVT+10]. Affective psychology is an active field in which novel theories are being proposed supporting or opposing previous research. We have careful consideration when picking a model that better suits our requirements and provides the users with the liberty of expressing themselves using natural language to better describe their perspective.

2.1 Sonification

To the best of our knowledge, no prior research has been conducted to investigate the likelihood of encountering a specific application of sonification. We make the empirical assumption that an audiovisual representation of data is a scenario both common and simple to understand for a lay person. The study revolves around the scenario to be experienced and complemented by the user. This contributes greatly to the participant's sense of comfort, because they will be placed in a familiar situation despite not having experience with the details of how to create such a scenario.

The emotional charge of a sound has been shown to be linked to experiences and the ability to recognize a familiar sound [AVT+10]. While some physical characteristics, e.g. loudness, still contribute to emotional judgement, sound designers often use unrelated sounds to emulate their target. This works well when the goal of the designer is to reproduce a sound effect with practical resources. Taking sonification as the auditory counterpart of visualization, which often implements abstract figures instead of real-life representations, we are interested in the emotional capacity of an abstract sonification.

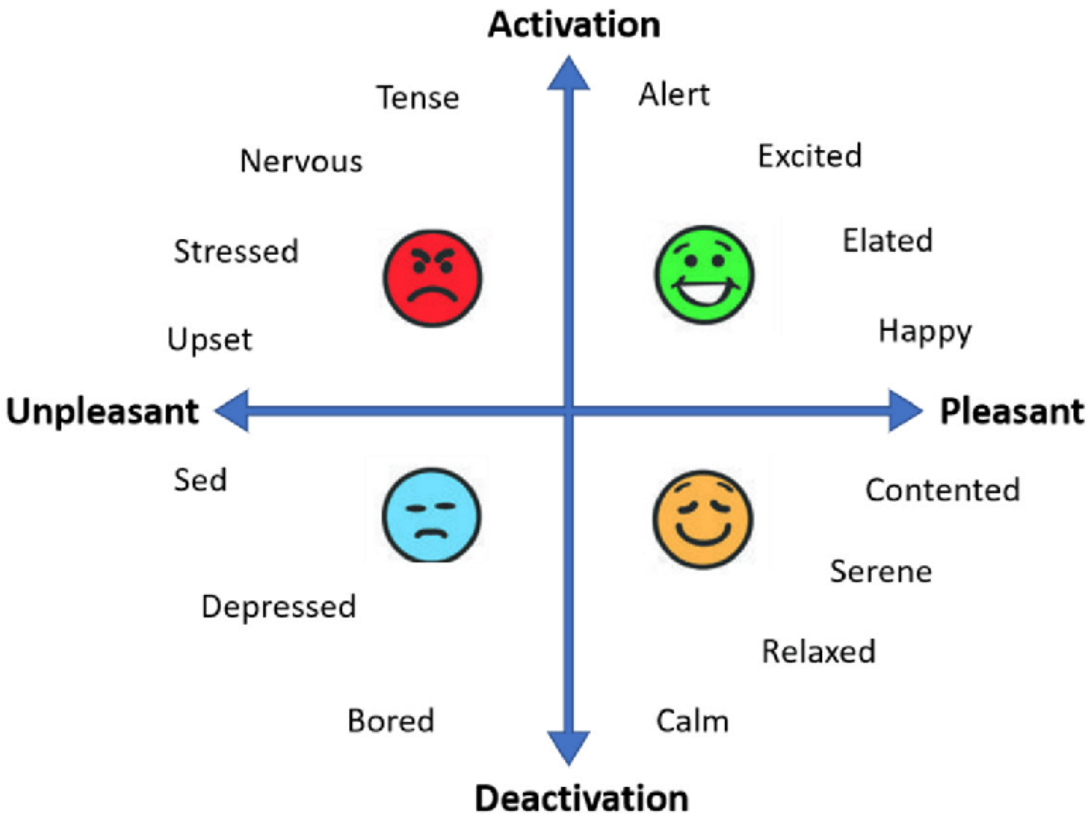

Figure 2.1: Circumplex Model of Affect

2.2 Emotional model

Emotions are a subject of constant study and rapid change in understanding. Even today there are competing theories with varying levels of compatibility between each other. The emotional model used during the user study can have a large impact in the resulting data. The reason for this is that the participants are expected to describe their own perceived emotions. This model will determine the limits of expression available to the users when reporting the data.

The Circumplex Model of Affect [Figure 2.1] is often used in computer science and related fields when researching its psychological aspects. We use this model for the design of our user study for the wide range of expression that it allows the users. Its low complexity in terms of dimensionality allows the participants to easily identify their experiences within the model. The researchers can map this to a numerical scale for analysis.

While the domain of this model is more granular and flexible than others, it is arguably more difficult for a lay person to translate their emotional reactions to a two-axis chart. We pair this model then with open questions for the participants to use natural language to express themselves. This is to attempt to mitigate misunderstandings between the model and a more natural expression of emotions. When discrepancies arise between the two, we favor the emotional model for uniformity of data representation.

3 User Study

The goal of the tool we designed is to better help us understand the perceived influence of sonification experiences on the emotions of the user. Emotion have been shown to play an important role in the cognition of human beings [Phe06]. In this context, we consider it relevant to know how the sonification aspect of an experience impacts the affect of the user. Sonification contributes to the overall experience. We believe that the emotional aspect of its contribution is non-negligible. With this in mind, we attempt to find how big this contribution is for users. We design a qualitative and exploratory user study to assist us in our research.

The tool was designed around a specific user-study. An in-depth explanation of the system itself is discussed in Chapter 4. While there is evidence of music's effect on emotions, comparatively little efforts have been made to study the same phenomenon in sonification. Our first goal with the study is to determine the self-reported significance of sonification during an audiovisual experience. For this purpose the study is divided into two segments: a familiarization stage and an audiovisual sonification stage. The user then reports their perception throughout the study and helps researchers understand how the different aspects the setting affect their experience.

3.1 Familiarization Stage

The starting point of the application provides the user with an interface lacking any labels or hints on what each component does. The user is then tasked to experiment with the tool and try to determine the functionality. There are blank spaces next to each component, allowing the user to take notes or write labels and hints for themselves. At this point, the goal of the user is to use natural language to help themselves better be able to produce desired results with the system.

3.2 Audiovisual Sonification Stage

While data in itself might set the default mood for the experience, we try to mitigate this by using data generally perceived as neutral. More specifically, we pick data for the average temperature over one calendar year in a region familiar with the user [Figure 3.1]. Other data taken into consideration is rain tendencies and air pollution. Rain is a good candidate and the study can easily be adapted for such a data set. On the other hand, air pollution is more likely to have a negative impact on the user's mood from the start and was therefore disqualified. The data is presented to the user in the form of an animated visualization lasting six (6) seconds.

At this point, the participant is tasked with sonifying the presented data. The user interface does not change. The visualization is integrated into it and the participant can dynamically select and manipulate sounds to represent the data. This sonification consists of a starting and an ending

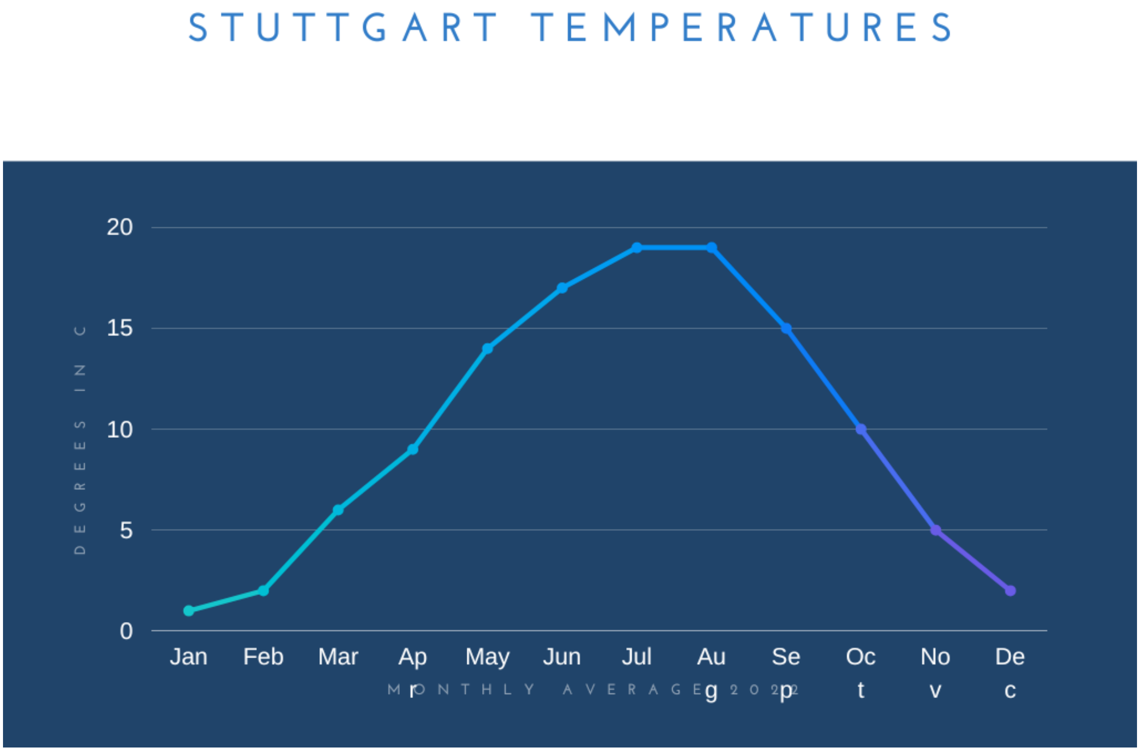

Figure 3.1: Visualization of average temperatures in the Stuttgart area in the year 2022.

sound. An interpolation is then automatically generated and matched in speed to the animation. The animated component helps the user visualize the relevant data that accompanies a playing sound. The participant can then make adjustments for their sonification to better match the visual aspect of the experience until they are satisfied with the result.

3.3 Design decisions

As discussed before, the data was selected for its neutrality. This helps to concentrate the study on the auditory aspect. Besides this it is important to avoid data that can cause distress or other negative emotions in particular to preserve the safety of the participants. Data too abstract, on the other hand, might make it difficult for the participants to relate and successfully design an auditory representation of the data. We believe striking this middle ground is key to get the best quality of data from the study. While unknown or fictional data might provide the neutrality of the visualization, it might also limit the investment of the participants.

The length of the animation (and accompanying sonification) was also a point of discussion. Previous research remarks the important role that the structural characteristics of music play in our emotional responses to it, among other aspects. To produce an effective sonification experience it is thus necessary to allow the participant to integrate some variation in the design. For this reason, we

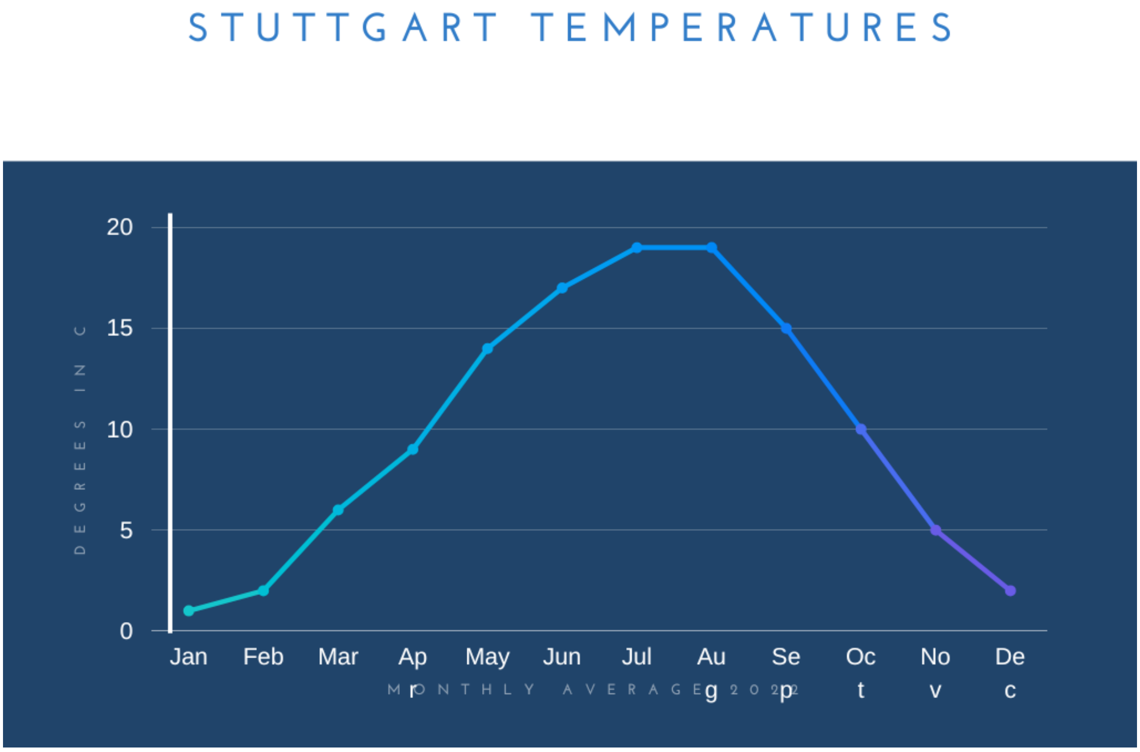

Figure 3.2: Animation start

decided to animate the visual representation of the data. We choose six seconds for our animation as a standard duration and provide the user with control over the speed of the overall experience. This way we aim to make the duration sensible for a short task, but allow for some variation to better achieve the desired result.

Two variants were considered for the animation, which are shown in Figure 3.2, Figure 3.3, and Figure 3.4. The variant marked as älternative"was first proposed, but the second variant is preferred because of visual clarity. It is easier to determine which point of the animation is currently playing and what data is yet to be represented during this loop.

3.4 Interview questions

There are three moments when we deem relevant to perform a short interview with the participant. The first round of questions takes place before starting the study. These questions aim to determine the context of the study. With this demographic information the researchers should be able to identify important variables and set the cultural framework for the study.

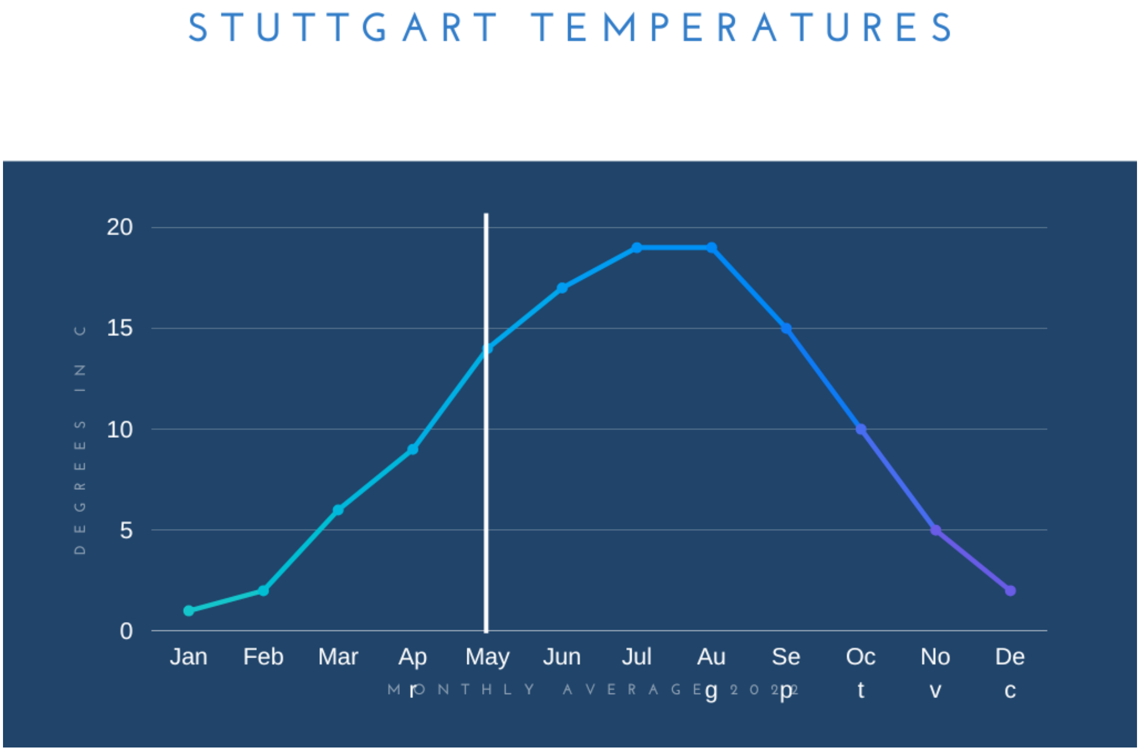

Figure 3.3: Middle point of the animation

After the participant has familiarized themselves with the application interface and labeled the tools comes a second round of questions. At this point in the study our main interest is to understand the rationale behind the labeling. The participant should explain how they interpreted the effects of each component of the tool. This is the first step on the participatory design, where the user attempts to improve the usability of the interface.

Finally, the Audiovisual Sonification stage, a final interview is performed. At this point, the participant has had the opportunity to actually use the tool in practice and get some results for the audiovisual experience. Our interest is then the perceived usability of the participant. Things like levels of satisfaction with the result and other practical feedback are enquired. This will not only be the main data points for later analysis, but also provides researchers with important information about how future research could be performed and how the tool could be adapted and improved.

Besides the interview, the whole process of the user study is recorded. Participants are encouraged throughout to speak out loud and describe their train of thought, as well as their experiences. During these recordings, reactions in real time to changes in the sounds can be collected. These information also helps better understand the decision-making process each participant goes through. By contrasting their spontaneous reactions and externalized dialogue with the answers during the interviews, we may achieve a more profound understanding on the process of designing the sonification.

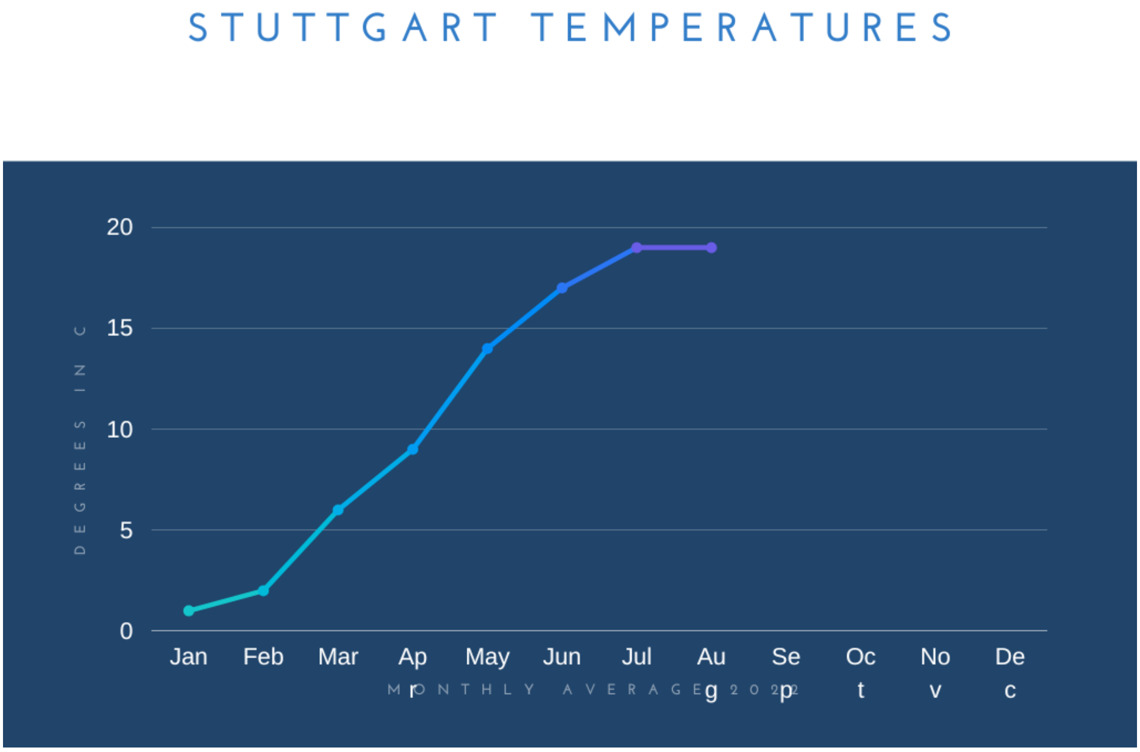

Figure 3.4: Middle point of the animation (alternative)

Second interview:

- Briefly describe how you chose the label for each component
- How difficult was it to choose the labels?
- Which component was the most difficult to identify?
- Which component was the easiest to identify?

Third interview:

- Do you feel that the tool allowed you to successfully creat the sonification you were aiming for?
- What controls are missing form the interface that limited your ability to create the sonification?
- Are there any labels you would change after using the software to create a sonification?
- Which emotional impact were you trying to convey with the sounds you picked?
- Do you feel like you achieved this goal?

4 Affective Sonification System

We designed a program to help us conduct the user study described previously. Commercial sound design tools, both software and hardware, employ advanced techniques which require previous experience and training. To a lay person, these interfaces may appear complicated and difficult to understand and use. Because we specifically select participants without any training or experience in this field, our interface needs to be simple enough to be understood in a short period of time during the study. The current interface can be seen in Figure 4.1. There is no time limit for any stage of the project, but we aim to finish each session within one hour. This means that each stage of interaction with the software should take approximately fifteen minutes.

To this end, the interface allows the participant to process one sound at a time. One sound is comprised of one or more tones which can, as a group, have their properties modified to achieve some variety in the outcome. The graphical interface is divided in three main horizontal sections. The top row is divided in two sections: musical notes and wave forms. One full octave of musical notes is available. Each note is represented by a checkbox that determines whether or not the note is active, i.e. sound is playing. No notes are activated by default when the program starts. Wave forms are represented by a group of radial buttons. This means that only one wave form may be active at a time, i.e. the wave form that all active notes take. Available wave forms are: sine (default), sawtooth, square, and triangle.

The middle row includes some controls for the sounds as well as space for the visualization used in the second stage of the study. The first section has two vertical sliders that control the amplitude and the frequency of the wave, respectively. These sliders' values range between 0 and 99, and are set to 49 by default. It is important to note that both sliders also affect all the active notes, for simplicity. Despite the minimum value of zero for the amplitude slider, it is not possible to completely mute the sound with it. This is to prevent functionality overlap with the checkboxes as this might confuse participants. The frequency slider allows the user to shift the sound by one octave up or down.

On the right section of the middle row there is room for the visualization as well as one additional slider. During the first stage of the study, the data and animation are hidden from view to reduce the amount of visual noise for the participant, as these do not play an important role during familiarization with the software. The bottom horizontal slider controls the speed of both the visualization and the sonification, which are synchronized to have the same duration. In the current iteration, this horizontal slider is not hidden but serves no function during the first stage, making it difficult to label. This design flaw can be addressed in future work with the application.

The bottom section of the interface is a control panel consisting of four buttons. Each button has a descriptive text and performs a simple action. From left to right:

- Save Labels: sets the label text for each component and prevents further editing. See Section 4.1.
- Save start: stores the current configuration for the starting point of the sonification.

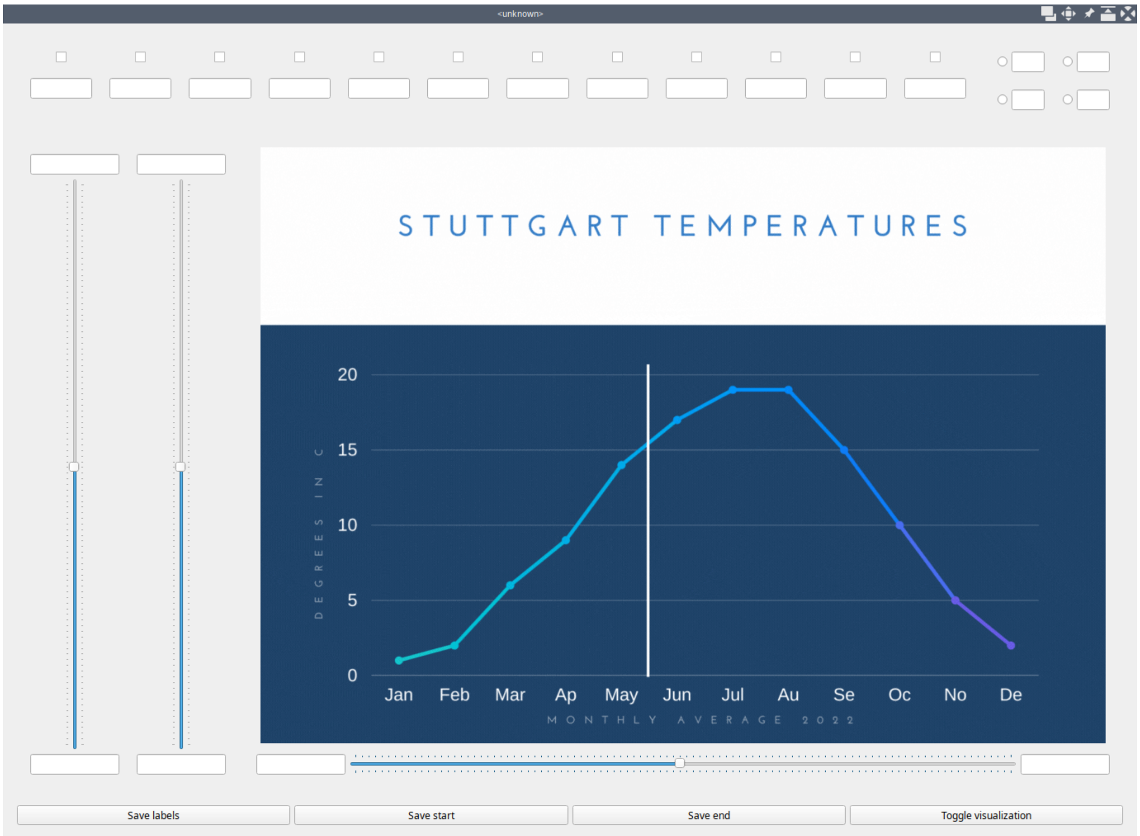

Figure 4.1: Current state of the graphical interface for the application.

- Save end: stores the current configuration for the end point of the sonification.
- Toggle visualization: brings the visualization into view if it is hidden, or removes it from view otherwise.

The start and end of the sonification refers to the sound configurations that will be interpolated to represent the data in the second stage of the study.

4.1 Participatory design

Each component of the interface involved in the sound design process has an empty text box next to it. In the text box, the participant can freely write any label they deem best to explain what is the functionality of the component. We only provide a brief introduction about how to interact with the application. From that point, we encourage the user to experiment with different aspects of the tool and to try to determine by themselves in what way each component is affecting the sound. During this experimentation phase, labels are to be filled using any kind of language that the user feels most comfortable with. While we encourage use of language descriptive of emotions or mood, the main goal remains for the user to use labels meaningful for themselves.

Similar to other design decisions taken for this research, the participatory design deviates from more traditional methods to find a compromise between complexity and usability. Because it is not trivial to translate a physical design into a digital user interface, we do not utilize tangible materials for our participatory design solution. Instead, we allow the user to dynamically alter the interface with which they will be interacting later. There is, of course, a limiting factor to the freedom of this customization. While it is technically possible to allow the participants to completely revamp the appearance of the application, the added complexity both for development and for the participation makes it unfeasible to achieve within the scope of our project.

A participatory design approach provides two main benefits. On the user's side, it allows for better understanding of the tool. It also promotes active participation and helps the researchers notice and talk through a lack of understanding of a particular aspect of the study. On the other hand, the labels themselves provide the researchers with another perspective on how the user is perceiving the change in sound. This is a key point to visit during the second interview.

4.2 Prototypes and design decisions

During conception and development of the user interface, the application went through several reviews and changes. The initial concept was to produce a system that would take an emotion as input and, based on previous research, generate an adequate sonification. This first goal was quickly scrapped and reworked into something more resembling our end product. There are no prototypes of this stage.

After further consideration, we integrate participatory design and a user central approach as goals for the finished tool. Based on that, we evaluate existing solutions for sound design as well as the possibility to come up with our own. Due to limiting factors of existing tools and their high complexity we settle for making the current system. Further details on the implementation may be found in Chapter 5.

The first prototype was presented in a colloquium and refactored largely based on the feedback received [Figure 4.2]. In summary, there was no space available to display the visualization and some additional functions were provided in the form of buttons and sliders. In particular, this initial prototype included a low-pass and a high-pass filter, additional to the components previously discussed. These effects can be effectively achieved by just using other tools already provided. Besides the redundancy, including them in the interface might produce situations in which certain components would behave differently when used in conjunction with the filters. This might confuse the participants and make it more difficult for them to properly use the tool and understand it.

The following are feature candidates that did not make it into the final design of the application.

- Including a frequency slider per musical note. While it could have provided a more granular control over the available frequencies to mix together for chords and progressions, the added complexity to the interface would have made it difficult to work with. Understanding what each slider does and coming up with a different label for each is not easily achieved in the time framework we were aiming for.

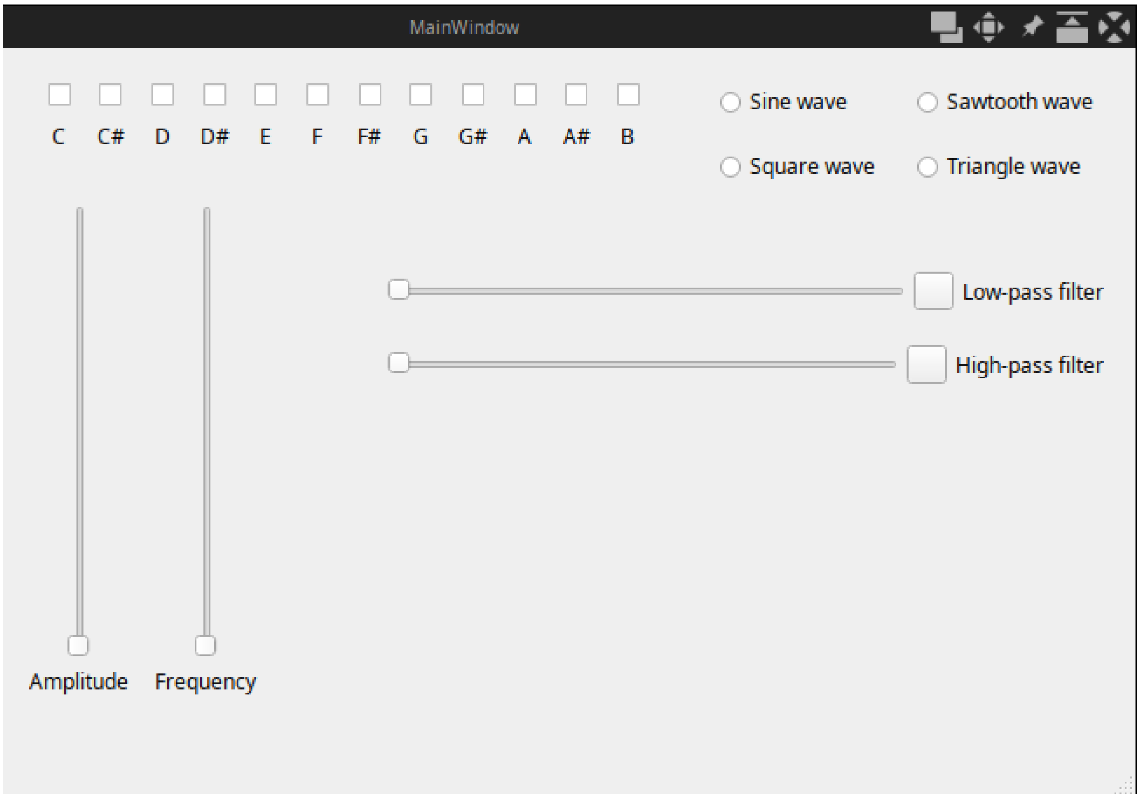

Figure 4.2: First prototype of the tool

- Having several tabs, each with its own sound configuration. This makes it possible to design more complex sonifications for a graph. For example, in a data set with twelve entries such as the one we use, the participant could create twelve steps for the interpolation and allow for a more interesting experience. However, due to the lack of experience of the participants it could be overwhelming. Design, and by extension sound design, is not a trivial activity. Asking an inexperienced participant to perform such a taxing task as sonifying several data points in a single short session could lead to diminishing returns.

One of the most significant changes to the interface is the inclusion of the start and ending points for the interpolation. While the idea of having several tabs could hurt the study, having a single chord or sound playing throughout the visualization is not enough to properly convey the data we are presenting. Furthermore, evoking emotions with a single monotone sound is not easily accomplished. Figure 4.3 shows the result of integrating the feedback received for the first prototype.

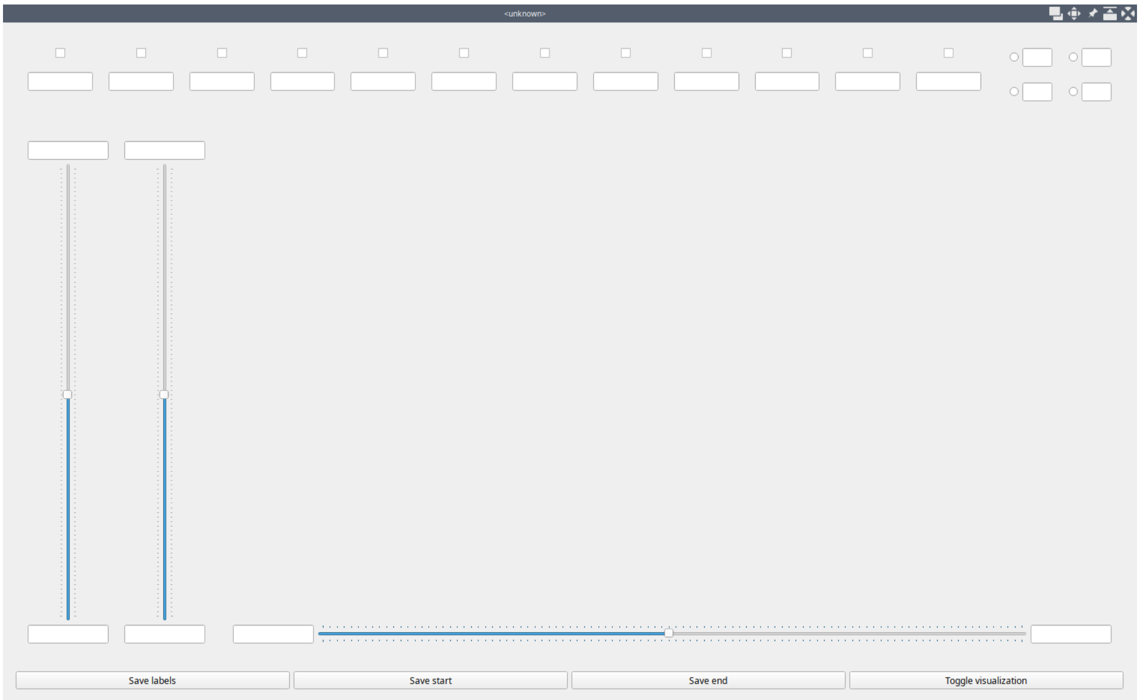

Figure 4.3: User interface after integrating feedback from colloquium

5 Implementation details

The system is comprised of two components: a SuperCollider server and a graphical interface made with Qt5 for Python. SuperCollider is a popular choice for working with sound design and music composition with code. The client-server nature of SC adds some complexity to the system, but the Open Sound Control protocol provides a friendly interface for interacting with the server. On the client side of things, we choose Python for its ease of use and widespread popularity [Sta]. The Qt framework provides bindings for Python, allowing us to create the graphical interface. Additional libraries are used for OSC.

5.1 SuperCollider

Most of the logic of the system takes place in SuperCollider. The server listens for messages sent using the OSC protocol and interprets the data using native methods to best produce the sound. We store all the music notes in a dictionary data structure, which is initially empty. A listener maps the received message to a music note that is generated based on global variables storing the other aspects, such as wave form or frequency. Using global variables in this way guarantees that all music notes are uniformly altered.

Three listeners are present to control the global variables. Sound aspects ‘ frequencyChange‘, ‘ amplitudeChange‘, and a control variable for each wave form are altered upon receiving an OSC message. After one of the variables is updated, the dictionary is iterated upon to set the new state of the system according to the new values. Two more dictionaries are available to store the start and end states of the sonification. This allows to perform the interpolation when the designer is finished working.

This design is a simple approach that gives SuperCollider full control over how to translate the numerical parameters from the graphical interface into the wave properties that it uses. Here we sacrifice data accuracy for sound accuracy. This is because SuperCollider is optimized for sound desing, but makes it more difficult to extract data. The server does not provide a way to retrieve information. In theory, we could send all the client-side data to the server and process everything on the server side. However, this would still would just mean giving up the advantages of using a general purpose programming language like Pyhton.

5.2 Architecture and Python

The client is in charge of drawing the graphical interface for the user to interact with. Here we have the most amount of raw data in the sense that it is the graphical interface what controls the SuperCollider server remotely. However, the data sent to the server is an abstract number with which the participants might be more familiar with. For example, amplitude and frequency are aspects that go from level 0 to 99 in steps of 1. This requires a mapping to the actual frequency and amplitude values used for the sound, but provides a known environment for the user in which they will feel more comfortable with.

A potential alternative to this approach is to use a mathematical framework such as ‘numpy‘ to perform the calculations on the client side and then send a message to the server with the values post processing. There is the potential to introduce bugs more easily with this approach, as math libraries are not specifically designed to work with sound and it relies on the programmer to properly map and manipulate the user input for SuperCollider. We aim to keep down complexity of the system for ease of changes and updates that it may require in the future for other studies.

5.3 Execution of implementation

At the start of implementation we used the tool QtDesigner to generate the code for the interface, as seen in Figure 5.1. This proved to be troublesome as the automatically generated code included some bugs difficult to get around of without modifying it directly. Of course, this is a further maintenance problem because automatically generated code would simply be rewritten and is difficult to document. By the end of the project, we rewrite the Qt code manually from the ground up to provide a better implementation with a more sensible level of abstraction that makes the codebase both maintainable and expandable.

The initial control window displays two buttons: a "labels"button which opens the pre-labeled interface for reference, and a "boxes"button which renders the unlabeled interface for the participant to fill in [Figure 5.2]. Internally, only one instance of the interface window is made and reused to both save memory and centralize the control and communication with SuperCollider. This single instance coupled with the internal structure of the window allows for granular control over the server by sending OSC messages. The Qt framework provides a robust event-driven architecture that we leverage to not just send messages when the user interacts with the interface, but also when closing the windows or setting a starting or ending point. During the first implementation with QtDesigner this was difficult to achieve and SuperCollider would continue to play undesired sounds at several points in time.

The SuperCollider server needs to be running. There are issues when trying to execute the server as a standalone program, so it is not possible at the time of writing to compile the whole system into a single script. Instead, the researcher needs to open a SuperCollider language interpreter (usually through the SuperColliderIDE, but it can also be done with third party tools) and start the server manually from there. This will execute the server with some initial configuration that allows it to function properly. Once the server is ready, two blocks of code need to be registered with it. The first one is a ‘SynthDef‘ used to control each music note and its properties individually. The second block encompasses the OSC listeners as well as control and update functions to keep track of the dictionaries and actually instance the ‘Synths‘ within the system.

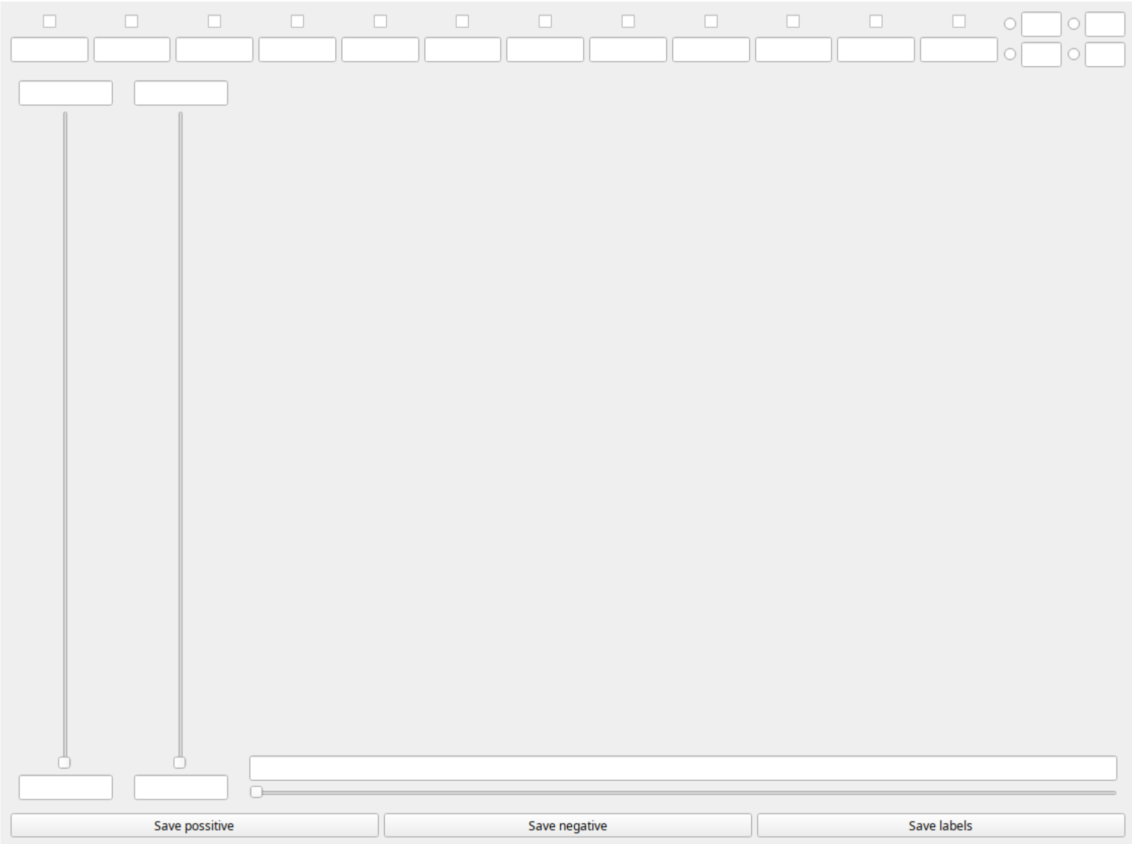

Figure 5.1: User interface generated automatically with QtDesigner

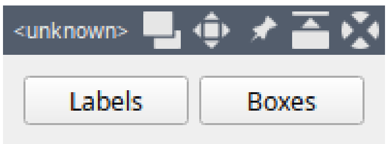

Figure 5.2: Main window of the application

6 Study pilot

In the course of the project we performed one pilot study. During the study, the user reported to have no previous experience with sound design nor a musical background. For the first stage of the experience we observe that the user used practical language to label the components of the interface. The user did not express strong emotional association with the musical notes. Using the frequency slider dynamically was described as entertaining, but once set to a static value it had little impact on the user's impression. Changing the wave form had the strongest effect on the perception of the sound. In particular, the square wave would make it initially uncomfortable, although this effect could be mitigated by lowering the overall frequency.

At the time of the pilot, the interpolation had not yet been implemented. However, the participant was tasked with creating a start and end for a potential sonification. Despite not being able to listen to the final result, the user expressed satisfaction with the result for each end point. As part of their feedback, they mentioned that the application lacks the capability to keep track of various sounds simultaneously and potentially using more than two steps for the interpolation. This is an important piece of information that can be expanded upon before performing a fully fledged user study. It is important to keep the complexity reasonably low, but allowing for some more liberty during the sonification step can be very valuable for the study. Because of this limitation, the participant described the experience as uncomfortable due to the constant monotone sound being reproduced during the design phase.

7 Conclusion and Outlook

Affective Sonification is still a young field with plenty of exploration to be done. We take a user centric approach to the research itself and try to circumvent biases embedded in traditional education of sound design and musical training. Our interests lie in a more natural way to express an emotional experience originating from audiovisual representations of data. We provide an outline for a potential user study along with the tools to conduct it. Our development strategy allows for further improvements and additions to the software to encompass a wider set of scenarios to be tested.

It is important to remark the role that cultural context plays in such a study. The software development as well as the design of the study and the pilot all took place in Germany, mainly with European people and people belonging to western cultures. An initial run of the study would most likely take a small sample of people with a similar cultural background, but further research can be applied to other cultures or a more intercultural context. Findings in situations where sonifications are abstract and independent from well known sounds with strong emotional associations could provide a better insight into what drives affect in relation to sound, or doesn't. In conclusion, we look forward to an iterative development of the tooling and implementation in at least one user study as described in this document.

Bibliography

[AVT+10] E. Asutay, D. Västfjäll, A. Tajadura-Jimenez, A. Genell, P. Bergman, M. Kleiner. "Emoacoustics: a study of the physical and psychological dimensions of emotional sound design". In: *3rd International Workshop on Perceptual Quality of Systems (PQS 2010)*. ISCA, Sept. 2010. DOI: 10.21437/pqs.2010-5 (cit. on p. 13).

[JKKS20] S. Jänicke, P. Kaur, P. Kuzmicki, J. Schmidt. "Participatory Visualization Design as an Approach to Minimize the Gap between Research and Application". In: *VisGap - The Gap between Visualization Research and Visualization Software* (2020). DOI: 10.2312/VISGAP.20201108 (cit. on p. 10).

[JV08] P. N. Juslin, D. Västfjäll. "Emotional responses to music: The need to consider underlying mechanisms". In: *Behavioral and Brain Sciences* 31.5 (Oct. 2008), pp. 559–575. DOI: 10.1017/s0140525x08005293 (cit. on p. 9).

[MMP+16] O. Metatla, F. Martin, A. Parkinson, N. Bryan-Kinns, T. Stockman, A. Tanaka. "Audio-haptic interfaces for digital audio workstations". In: *Journal on Multimodal User Interfaces* 10.3 (May 2016), pp. 247–258. DOI: 10.1007/s12193-016-0217-8 (cit. on p. 10).

[Neu19] J. G. Neuhoff. "Is Sonification Doomed to Fail?" In: *Proceedings of the 25th International Conference on Auditory Display (ICAD 2019)*. Department of Computer and Information Sciences, Northumbria University, June 2019. DOI: 10.21785/icad2019.069 (cit. on p. 9).

[NHH+11] J. G. Neuhoff, A. Hunt, T. Hermann, T. Hermann, A. Hunt, J. G. Neuhoff. *The sonification handbook*. Logos Verlag, 2011, p. 564. ISBN: 9783832528195 (cit. on p. 11).

[Phe06] E. A. Phelps. "Emotion and Cognition: Insights from Studies of the Human Amygdala". In: *Annual Review of Psychology* 57.1 (Jan. 2006), pp. 27–53. DOI: 10.1146/annurev.psych.56.091103.070234 (cit. on p. 15).

[Rön21] N. Rönnberg. "Sonification for Conveying Data and Emotion". In: *Audio Mostly 2021*. ACM, Sept. 2021. DOI: 10.1145/3478384.3478387 (cit. on p. 9).

[Sta] StackOverflow. *Developer Survey 2023*. URL: https://survey.stackoverflow.co/2023/?mod=djemCIO#most-popular-technologies-language (cit. on p. 27).

All links were last followed on July 02, 2023

Declaration

I hereby declare that the work presented in this thesis is entirely my own and that I did not use any other sources and references than the listed ones. I have marked all direct or indirect statements from other sources contained therein as quotations. Neither this work nor significant parts of it were part of another examination procedure. I have not published this work in whole or in part before. The electronic copy is consistent with all submitted copies.

place, date, signature