

Developing Visualization Software for Musicians

by

Michael Brooks

May 2010

Abstract

This paper describes the initial stages of the design, construction, and evaluation of a software program that presents real-time music visualization to musicians as a form of feedback. I identify the musical features of pitch, loudness, and timbre, all of central importance to musicians, as good targets for visualization. I then describe several audio analysis techniques that can be used to approximately extract these features. The needs and preferences of potential users with regard to musical features and the specifics of the visual presentations were investigated through a user study involving the testing of a prototype. I explain the goals of the study and the architecture and specifics of the prototype. Finally I summarize the results of the study and propose future directions of research.

Acknowledgements

I would like to thank the following people:

- the 38 participants who completed my survey and especially the 4 who agreed to spend an hour of their semesters doing an interview.
- my advisor, Richard Salter, for keeping me on track over the course of this year.
- Nancy Darling for providing valuable feedback as I wrote the survey questions and applied for IRB approval.
- the Conservatory Deans and the organizers of the College of Arts and Sciences Orchestra for agreeing to distribute the survey to the students on their mailing lists.
- my trumpet-playing housemate for test-driving the survey before I sent it out to hundreds of people.
- my piano and flute teachers whose effort over the years inspired me to undertake this project in the first place.
- my wife, Katie Kuksenok, for encouraging me to take the project in directions I would never have thought of, for providing advice on analyzing some of the data, and for finding lots of nitpicky problems with this paper.
- the encouragement of countless other people who I spoke to about this project.

Contents

Abstract	ii
Acknowledgements	iii
1 Introduction	1
1.1 Why do musicians need software?	1
1.2 What kind of software do musicians need?	2
1.3 Background research	2
1.4 Paper overview	3
2 Musical feature extraction	4
2.1 Musical features	4
2.1.1 Dynamics	4
2.1.2 Pitch	5
2.1.3 Timbre	5
2.1.4 Tempo	6
2.2 Audio analysis and feature extraction	6
2.2.1 Sound pressure level	6
2.2.2 Frequency Spectrum	7
2.2.3 Spectral Centroid	7
2.2.4 Special Normalized Autocorrelation and Pitch	9
3 Constructing the software system	11
3.1 Requirements	11

3.2	The first prototype	12
3.3	The second prototype	12
3.3.1	Engine design and dataflow	13
3.3.2	The visualizations	14
4	The user study	18
4.1	Design and implementation	18
4.1.1	Survey overview	19
4.1.2	Interview overview	20
4.2	Results of the study	20
4.2.1	Survey population details	20
4.2.2	Interview population details	22
4.3	Results and Conclusions	22
4.3.1	Musical feature interest	22
4.3.2	Waterfall visualizations	27
4.3.3	3-Line visualizations	27
4.3.4	Spectrum Ring visualizations	28
4.4	Broad implications for design	29
5	Future work	30
5.1	More flexible software	30
5.2	More detailed user study	31
5.3	Final thoughts	31
	Bibliography	32
	A Survey Questions	34
	B Interview Questions	45

Chapter 1

Introduction

1.1 Why do musicians need software?

The learning and performance of music is an art that is complex and difficult to master. Musicians must learn a variety of skills from bizarre notations, to fine-grained physical control, to artistic and context-appropriate interpretation and expression. In actually playing music, musicians typically have to concentrate on all of these subtle issues at once. They must produce the desired sound, which can involve controlling the hands, the face, the lungs, and the overall posture simultaneously. They must produce the desired notes, which involves concentrating on either reading an information-rich score or recollecting a previously learned score, and then translating into certain physical actions. They must also control broader aspects of the music like dynamics, phrases, and articulations in order to produce the desired effect. All of this must happen in time with the desired musical tempo, and these are just the difficulties that are inherent to the performance of music, ignoring all of those that might arise from instrument, player or piece-specific characteristics.

Beginner music students have a lot of difficulty because in order to gradually introduce to the student all of these facets of musicianship, a teacher will have the student play very basic music that focuses on individual skills. Among other reasons, this is one important cause for the student becoming bored with practicing.

Another problem that musicians sometimes have is a lack of day-to-day feedback that is more objective than the musician's own ears. On a weekly basis, the musician might

get valuable criticism from an instructor, but on the other days he is generally left to his own devices. This presents difficulties for several reasons. First, the musician might be experiencing intermittent problems that don't manifest while he is with the teacher. Second, there is only so much that can be covered during a weekly lesson. Other problems are left to the student to work on. Third, the student may not be optimally suited to identifying and correcting whatever problems he has.

I believe that each of these three problems can be addressed by equipping students with software that provides objective feedback through music visualizations. This project investigates how visualization software could best be designed to help music students.

1.2 What kind of software do musicians need?

Even within the scope of music visualization software, there is a huge variety of approaches to visualization design. In order to narrow it down, the needs of musicians need to be more carefully considered in order to select the more effective of those designs. In a user-centered design process, these design ideas would be presented to users, who might simulate interaction with the prototypes with the prototypes. After analyzing the results of such a study, the best designs could be selected and recombined into new prototypes. This process would then repeat, and eventually a complete software product would be created.

In order to better understand the needs of users and to identify what types of designs musicians might prefer for this kind of software, I conducted a single user study involving a survey and interviews where the participants evaluated several prototype visualizations.

1.3 Background research

There have been relatively few research efforts applying real-time visualization to music, especially for the purpose of improving music education. Johnston et al provide an overview of some recent work that encourages reflection in music making through software tools [8]. Callaghan and Wilson et al have published on the results of thorough studies evaluating the use of visual feedback in singing education in [1], [14], and [15]. There has also been

some past work in constructing particular visualizations of musical features. Ferguson et al describe a visualization developed to combine elements of intonation, timbre, and loudness into a single visualization object in [6]. Dixon et al describe a method of visualizing tempo and loudness as a two dimensional plot in [3] and [4].

Visualizing real-time information about music rests on the ability to extract specific features from digital audio recordings. A major source of information in that regard has been the work of McLeod on efficient and accurate real-time pitch extraction and his program Tartini, described in [10] and [9]. Peeters has provided a comprehensive list of audio features intended to be used in music similarity and classification projects [11]. Tzanetakis et al, in their work on automatic genre classification, summarize a small set of features particularly useful for extracting timbre [12]. Finally, Dixon describes three recent tools that showcase feature extraction techniques focused on rhythm and tempo, loudness, and automatic transcription [2].

1.4 Paper overview

Chapter 2 first describes background about several specific facets of music that I was interested in providing to musicians via visualizations. It also describes appropriate methods for extracting these features from digital audio. In Chapter 3 I explain the architecture and construction of the prototypes that I build over the course of this project, leading up to the user study. Chapter 4 explains the details of the study, and summarizes the results. Chapter 5 draws conclusions and proposes future directions of research.

Chapter 2

Musical feature extraction

2.1 Musical features

Music has a multitude of characteristics that are recognized by musicians. For example, there is pitch, tempo, dynamics, tone, articulation, mood, and a variety of other features. Each of these is important in describing the music. Some of these are more universal to different types of music others, making them better targets for this type of software. Moreover, many of them, such as mood, are so subjective they would be impossible to measure.

I will describe three features that apply generally to most music and most instruments. While not possible to define rigorously, they are relatively well defined and so are excellent candidates for automatic detection. Because of their fundamental importance in most music, they should be of interest to most musicians. Ferguson describes research in a similar vein to this project that identified similar features as good subjects for visualization [6].

2.1.1 Dynamics

Dynamics is the musical term for the loudness of a sound. Dynamic changes are notated in modern music with one of a set of symbols, which musicians usually regard as relative loudness levels. The indicated loudness is supposed to fit within the dynamic range of the piece as a whole.

I predict that dynamics will be of interest to musicians because it is one of the most

noticeable parameters of music to an audience, and yet I believe that in many cases, especially among students, it receives relatively little attention. That is, musicians might inadvertently change their dynamic level without realizing it. Visualization software might help to correct behaviors like this.

2.1.2 Pitch

Pitch is that feature of music which is described as “high” or “low.” For most instruments it can be measured with a high degree of objectivity, and it is related to the lowest prominent frequencies that are present in a sound.

Accidental pitch inaccuracy is an unavoidable problem for most wind and string instruments and is something musicians deal with constantly. Providing visualizations that reflect pitch could therefore help musicians. This is not a new idea, certainly, since most musicians own and sometimes use an electronic tuner.

2.1.3 Timbre

In a typical musical sound, there is a frequency called F_0 , the fundamental frequency. In an idealized sound model, there would also be an infinite number of additional frequencies present, at integer multiples of F_0 . So, an A oscillating at 440 Hz also has frequencies (called harmonics) 880 Hz, 1320 Hz, 1760 Hz, etc. These harmonic frequencies tend to decrease in amplitude as they get higher.

In actual sounds the placement of these harmonics varies slightly from the ideal model. In addition, the amount of each of these frequencies can vary a great deal. These variations in the harmonic series for the note give rise to a large part of what is collectively called timbre, or those characteristics of sound which are not explained by pitch and loudness. Visualizing timbre could be helpful because it might allow the musician to realize when they are not producing the timbre they want to produce.

2.1.4 Tempo

Tempo describes how fast the piece is played, and is usually measured in beats-per-minute. Tempo is important in that it greatly affects the audience's and the performer's experience of a piece of music. Musicians often have trouble achieving consistency in their tempo, especially when so many other factors demand their attention.

2.2 Audio analysis and feature extraction

The above musical features are all at least partly subjective, ill-defined, and disagreed upon. Two musicians probably will not have precisely the same conception of any of them. Something like timbre, in particular, is difficult to define [5]. Using the algorithms listed here, various features can be extracted from raw audio, and these features have been found to approximately represent the facets of music above.

2.2.1 Sound pressure level

This feature is calculated directly from the digital samples of the audio. For a segment of audio, the amplitudes are squared and the average is taken. This value is then converted into decibels by dividing by a reference value, taking the base 10 logarithm, and multiplying by 10 [13, pp 41-45]. Many more sophisticated techniques exist, but this value is a very rough approximation of the perceived relative loudness of a sound.

The mean squared amplitude of a sound segment x of length N is calculated with:

$$P_{msa} = \frac{1}{N} \sum_{i=0}^N x_i^2$$

Given a baseline P_0 , which I calculate from a segment of silence, the loudness L is given by the following formula:

$$L = 10 \log \left(\frac{P_{msa}}{P_0} \right)$$

I use this number to measure and represent the musician's dynamics. The output of this value for several seconds of a Beethoven Piano Sonata is shown in Figure 2.1.




Figure 2.1: The output of the loudness calculation over time.

2.2.2 Frequency Spectrum

The frequency spectrum is calculated by taking the Fourier transform of a segment of audio. Due to the Nyquist sampling theorem, the maximum frequency that can be encoded in the audio input I record (sampled at 44100 Hz) is 22050 Hz. The frequency space from 0 to 22050 is divided up into evenly spaced frequency bands, and for each of these bands f , the frequency component X of the sound x in that band is calculated directly according to

$$X_f = \sum_{t=0}^N x_t e^{-2\pi i f t / N}$$

However, since this must be calculated for N individual bins, rendering the running time quadratic, a more efficient indirect algorithm is used, known as the Fast Fourier Transform [9]. The lower half of the frequency spectrum for a trumpet note is shown in Figure 2.2.

2.2.3 Spectral Centroid

The spectral centroid is the weighted average of all the frequencies in the spectrum, weighted by how much that frequency is present in the sound. For a frequency spectrum X with N bins, the spectral centroid F_c is given by

$$F_c = \frac{\sum_{i=0}^N i X_i}{\sum_{i=0}^N X_i}$$




Figure 2.2: The frequency spectrum for a trumpet note. The upper half of the spectrum has been omitted.

Previous research efforts, while disagreeing on many aspects of how best to encode timbre, have often agreed that this is a very significant part of timbre [7]. Since it is higher when there are more high frequencies present, it can be said to represent the sound brightness. The spectral centroid for a series of vowels is shown in Figure 2.3.




Figure 2.3: The spectral centroid for a series of sung vowels.

2.2.4 Special Normalized Autocorrelation and Pitch

This is a variant from a large family of pitch detection approaches. The autocorrelation family is concerned with determining the most probable period of the audio segment. This period then determines a frequency, and hence a pitch. The audio is multiplied against itself at increasing offsets, or periods. Intuitively, the sums of these multiplications represent how well the wave lines up with itself at that period. The problem is then that the period with the highest correlation needs to be selected. Because the regular autocorrelation does not have well-defined normalization, this makes writing a peak-picking algorithm difficult.

The advantage of this variant is that it normalizes the autocorrelation, and also increases its accuracy. It was developed by Philip McLeod and is defined and justified in detail in [9]. The formula for the SNAC function n' for a particular period τ is

$$n'(\tau) = \frac{2r'(\tau)}{m'(\tau)}$$
$$r'(\tau) = \sum_{j=0}^{N-1-\tau} x_j x_{j+\tau}$$
$$m'(\tau) = \sum_{j=0}^{N-1-\tau} x_j^2 + x_{j+\tau}^2$$

The r' and m' terms are the cross terms and square terms, separated from the regular autocorrelation function. An example SNAC function is pictured in Figure 2.4.

The pitch can then be calculated by picking the best peak from the SNAC function. There is currently no guaranteed method of choosing the best peak, but McLeod's approach, which achieves workable accuracy, is to select the peak within a certain threshold of the maximum. This peak identifies a period which can be converted into a frequency and then to a pitch. The output of pitch over time is shown in Figure 2.5.

It should be noted that this only works with any accuracy at all when only a single pitch is present in the sound. Multiple pitches cause it to behave unpredictably. Moreover, as McLeod notes, this algorithm also sometimes makes octave-errors due to picking the wrong peak of the SNAC function.




Figure 2.4: The SNAC function for a trumpet note. The fundamental peak is circled.




Figure 2.5: This shows pitch output over time.

Chapter 3

Constructing the software system

In order to determine if, why, and how musicians would like to use visualization software, I ran a user study in which the participants tested a prototype that I had created. This chapter describes the architecture and construction of that software system.

3.1 Requirements

There were several key requirements that I needed the software to satisfy.

First, the system needed to be platform independent to make sure that the participants could test the software on whatever system they were most familiar with.

Second, the system needed to run in real time. My initial intuition was that a real time visualization program would be most helpful, since it would allow instant adjustment of the musician's playing in reaction to visual feedback.

Third, the system needed to be easy to develop because it was going to be fairly large. Additionally, I wanted a system that would support the eventual instantiation of a large library of visualizations all working with the same underlying engine. This was due to the fact that I would only be running one study and would want to allow students to see as much diversity as possible.

Having considered these three priorities, I selected Java as the language for the program. It has the great advantage of being platform independent and supporting rapid development. However, it also has a reputation for poor performance. This never became a problem in

my application, but it was one argument for preferring C++ (a more typical choice for audio-handling software).

As part of the decision to use Java, I also committed to using the Java Sound API, which worked adequately. After testing a number of alternatives, I selected the JTransforms library available from <http://sites.google.com/site/piotrwendykier/software/jtransforms>. This library proved to be much faster than the other library I tested, the implementation in the Apache Commons Math package.

3.2 The first prototype

During the first semester of working on this project, I developed a very early prototype exhibiting many of the analytical features I had been working on.

This prototype had the primary purpose of serving as a testing platform while I ironed out the many bugs in the pitch detection process. There were many changes to the system since its purpose changed depending on what particular calculation I was trying to fix at the time. However, at its most complete, it had the following interface:

Shown here is a plot of the SNAC feature in the upper left. The upper right shows pitch-vs-time. The lower left shows the pitches as well, but segmented into notes. On the lower right is a plot using the note segmentations to plot note length.

Both the SNAC and pitch extraction calculations became more sophisticated after this prototype was produced, but even in this rudimentary state the pitch plot proved to be fairly compelling. The note segmentation algorithm was very imprecise, but it too worked fairly well. Finally, the note-length visualization was promising, but I did not end up migrating it to the next prototype.

3.3 The second prototype

At the start of the second semester of this project, I started the development of a totally new prototype. It made use of some of the GUI code from the first prototype and some of the calculation code that I had finally succeeded in perfecting. However, the system had a

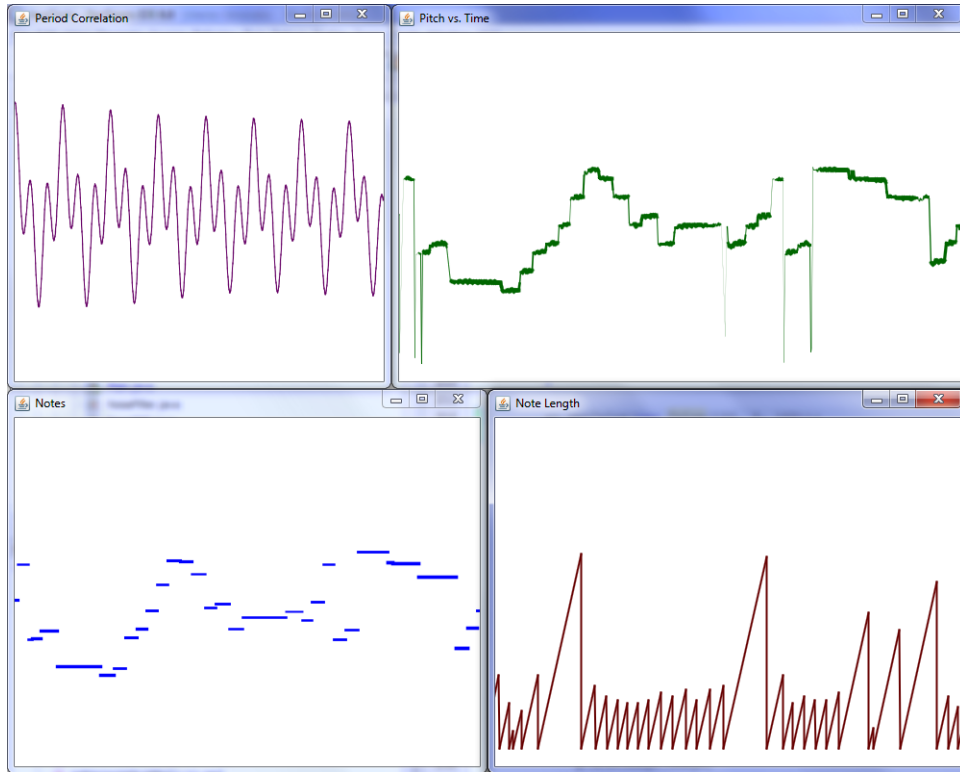


Figure 3.1: The interface for the 1st prototype

completely different architecture and was intended to support development into a reasonably complete package.

There were several reasons for scrapping the first prototype. For one, it was not extensible any more. It had suffered so many haphazard modifications that it was disorganized. As an indication of this, I had implemented three separate main methods over the course of its development, each one being phased out in turn for the third.

The new system featured a fairly complex data-handling scheme by which audio could be produced at a source module and then distributed to whichever systems required.

3.3.1 Engine design and dataflow

There are three primary components:

The Audio Input Module wraps the Java Sound API and also provides notifications to the rest of the system when audio is available. It produces buffers of audio and stores them in a shared location. It then spawns a new thread in the Audio Analysis engine to carry out most

of the calculations, as well as to save the completed calculations in a globally available queue. The User Interface Manager runs its own, independent thread which periodically polls the queue and extracts any new data. It then selects from among the calculations performed by the Analysis Engine and provides the necessary data to the various Visualization objects to generate their graphics.

The system architecture is illustrated in Figure 3.2. Because the chunks of audio produced by this system are 2048 in length and the sampling rate is 44100 Hz, every step of this data path has to run in under 1/20 of a second.

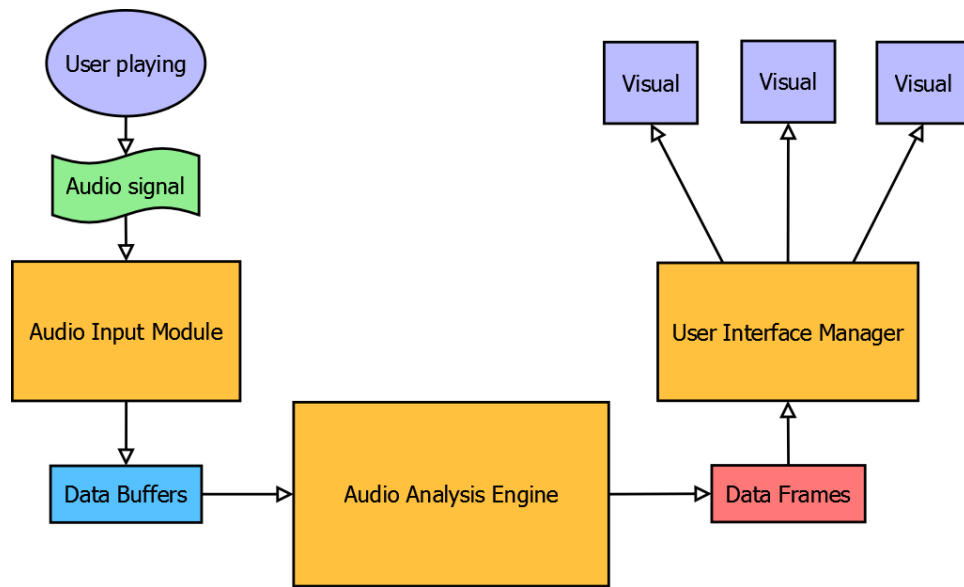


Figure 3.2: The architecture for the second prototype, showing data flow

3.3.2 The visualizations

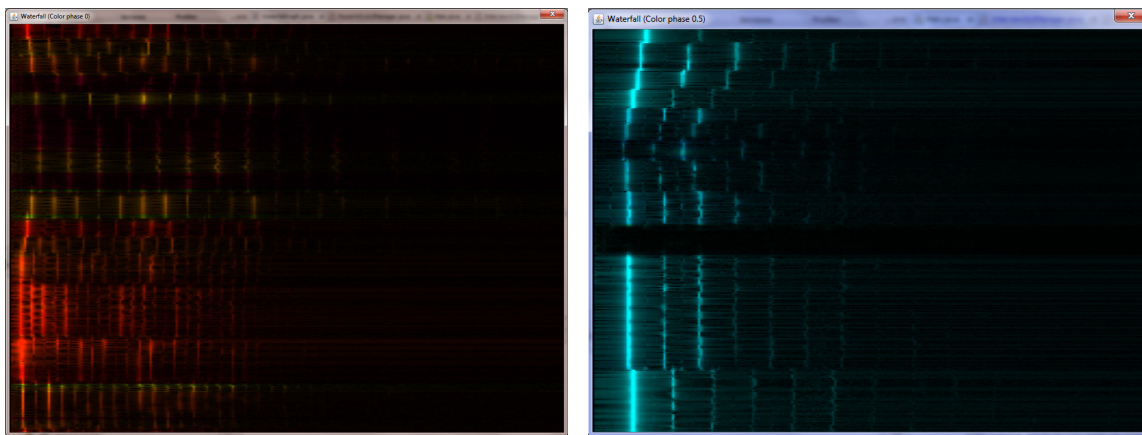
This second prototype, which I used for the study, presents six different visualizations. There are three primary types of visualization presented, each with two minor variants.

3.3.2.1 Waterfall visualizations

The first two that I presented to users were the Waterfall visualizations. These primarily show the frequency spectrum, timbre, and loudness of the music. Their visuals are generated from

the frequency spectrum directly, and are colored using the spectral centroid, approximately representing brightness.

As new data frames arrive, their frequency spectrum is drawn across the very top row of pixels on the visualization, and the existing image is pushed downwards by one pixel height. The resulting animation resembles a waterfall. Figure 3.3a illustrates the red-orange-yellow variant of this visualization, which was presented to users first. The blue-green-purple variant, shown in Figure 3.3b, is simply a re-colored version of the first.



(a) Red Waterfall variant

(b) Blue Waterfall variant

Figure 3.3: Both variants of the Waterfall visualization

3.3.2.2 3-Line visualizations

These visualizations, like the Waterfall visualizations, record the last several seconds and maintain them on the screen until they are eclipsed by new data. Unlike the waterfall visuals, though, these plots present their data as simple timelines. In order to enhance the visibility of the timelines, I had them color in the area under each line. Furthermore, each line is colored differently to make them easier to distinguish.

The first variant, plotting Pitch, Loudness, and Brightness on its three lines, is pictured in Figure 3.4a. It pulls the pitch estimation, loudness, and brightness (as the spectral centroid) and uses those to generate its plots. While the range on the pitch and brightness plots is fixed, the loudness plot automatically scales its upper bound to fit whatever loudness data

it has been provided thus far.

The second variant is only different in that it replaces Pitch with Intonation (Figure 3.4b). The Intonation plot compares the pitch against the absolute 12-tone equal-tempered scale and identifies which note is closest to the current pitch. It then plots the difference between the pitch played and the identified closest note. This shows how sharp or flat the player is.

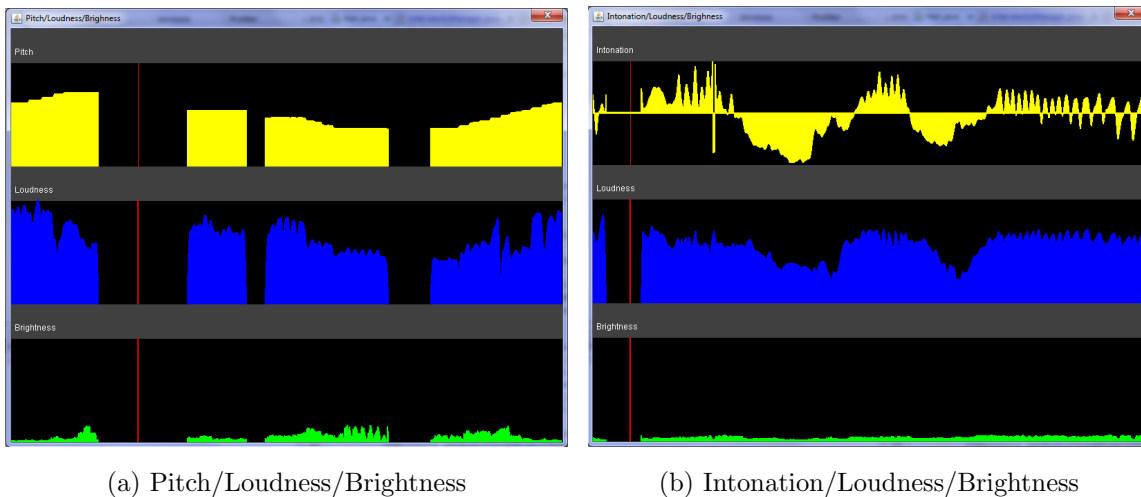


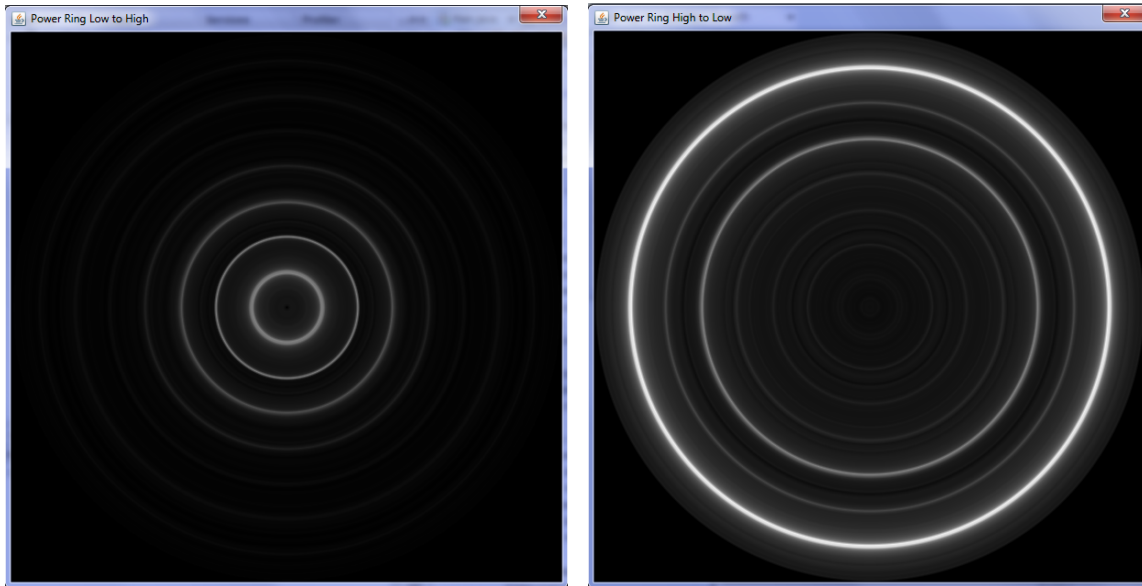
Figure 3.4: The 3-line visualizations

3.3.2.3 Spectrum Ring visualizations

These final two visualizations take a different approach from the other two types. Instead of displaying a few seconds of audio, they show snapshots of information about the current audio frame. These visualizations are simple in that the only feature they show is the frequency spectrum. They plot the frequency spectrum by sweeping it out in a circle around the center. Unlike the previous two types of visualization, this one is drawn in grayscale instead of with colors.

The first variant of this type plots the spectrum with the low frequencies in the middle and the high frequencies on the outer rim. It is pictured in Figure 3.5a. One effect of drawing the frequencies in this way is that the higher frequencies dominate the visualization because

they are drawn more prominently. Because the high frequencies of a sound may fluctuate in unpredictable ways, this visualization can appear chaotic. That is why I decided to include the second visualization of this type, pictured in Figure 3.5b. This plots the low frequencies on the outside and the high frequencies on the inside, minimizing their effect on the energy of the visualization.



(a) Low on inside, High on outside

(b) High on inside, Low on outside

Figure 3.5: The Spectrum Ring visualizations

Chapter 4

The user study

This chapter describes the design of the study that I used to evaluate potential interest in music visualization software among musicians and to identify the specifics of that interest. I also discuss the summarize and discuss the results of the study.

4.1 Design and implementation

The goals of the study were to answer two important questions:

- What feature are musicians interested in visualizing?
- How should those visualizations look?

Out of the several kinds of information my prototype can provide about music, including pitch, loudness, and timbre, which do musicians find most helpful? Which are most fun to see? I attempted to ask this question in two ways. In the survey, I included questions where the participants had to assign levels of interest to each of four possible musical features. Subsequently, in the interviews, I was careful to record which specific information in the visualizations the participant responded to.

For the second question, I also used both survey and interview data. Part of the survey asked about prior music software use, and while the main purpose of these questions was to gauge level-of-experience in the more general category of music software, I was also looking for

information about what kinds of software had been successful in getting used by musicians. Then, in the interviews, I focused on what specific qualities of the visualization's appearance rendered it more or less useful than the others, and which qualities made it more or less fun to watch.

I will describe in more detail how both the surveys and interviews were put together in the following two sections.

4.1.1 Survey overview

The survey had five sections containing questions about demographics, technology use, music experience, use of software in music, and preferences about visualizations. In addition to option-based questions, there were several open-ended questions throughout the survey that allowed the student to expand their answers to previous questions. At the end of each survey I asked the student to voluntarily submit their contact information if they would be willing to do an interview.

The survey, hosted by SurveyMonkey, was distributed via email to two populations of musicians. One group to receive the email was the entire population of students enrolled in the Oberlin Conservatory of Music. This email was distributed to the students through the Conservatory administrators, who agreed to help in this capacity. The second group of students contacted was the Oberlin College Arts and Sciences Orchestra members, a group made up of College (non-Conservatory) students that meet weekly to play orchestral repertoire leading up to semesterly concerts. The main organizers of this group agreed to distribute my survey to their mailing list. Thirty-eight people in total completed to the survey.

The objective in targeting these two groups was to collect the viewpoints of both extremely serious, pre-professional musicians and musicians who, while usually well beyond intermediate level on their instrument, play mainly for the fun of it. The size of the Conservatory mailing list that received my recruitment email was just over 600 students. The number of members on the College Orchestra list was around 50 students.

4.1.2 Interview overview

I recruited the interview group from those students who indicated they were willing on the survey. I managed to arrange meetings with four of the students I contacted.

The interviews that I conducted covered several topics. I began the interviews by asking for some more detail about the student's use of computers in their practicing. Since most people on the survey had indicated that they rarely brought their computers with them to practice, I wanted to find out why this might be.

Though some of the specific questions might vary from student to student, I also asked everyone to identify the thing they found most difficult about playing music in general, and also the thing they found most difficult about playing their instrument. My goal here was both to see if students with different goals had different reactions to the prototype, and also simply to get the student to start thinking critically.

Before showing them the prototype, I explained some more about what the objectives and approaches of my research were, and asked them for some ideas about what kinds of goals they might have for this software and what they imagined it might look like.

That out of the way, I began showing the student each visualization in turn. The student would typically play for a couple of minutes, and give several remarks as they occurred to him. I asked about each visualization's usefulness and about ways each could be improved.

4.2 Results of the study

4.2.1 Survey population details

Due to the relatively small size of my survey group, I was somewhat concerned about its diversity in a number of important dimensions. One particular dimension of interest is how many students came each of the three programs at Oberlin: the Conservatory, the College, and the Double-Degree program (enrolled in both). Because Conservatory musicians are expected to be more serious and advanced about their music study and College musicians to be more casual, the fact that the population is almost evenly divided among these three categories indicates a certain basic diversity of viewpoints. There were 11 College students,

13 Conservatory students, 13 Double degree student, and 1 student in neither program.

A second dimension that is important to my project is what instruments students play. Since different instruments have different sound characteristics and are difficult to play in different ways, it is likely that musicians playing different instruments have different goals and preferences about visualization software. There were 25 distinct instruments reported by students. Figure 4.1 shows how students were distributed among the instruments, divided by which program at Oberlin the students belong to.

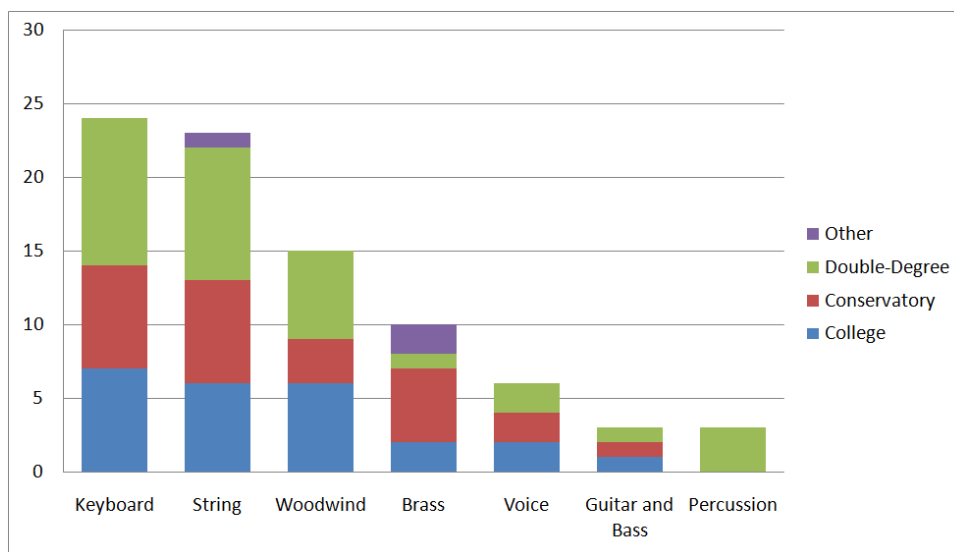


Figure 4.1: Instrument representation from the three programs.

As a more direct means of evaluating the student's commitment to studying music, I also asked how many days in a typical week they practiced their instrument. There was a reasonable variety of responses to this question. The details are shown in Figure 4.2.

Since the survey contained several open-ended questions on four main topics, I have provided the response quality for each of these topics in Figure 4.3. For each participant, the response rate in each category is represented by the height of that category on the figure. As the figure indicates, the questions about Music Software Use were well answered by a large group of students. The questions about Visualization Preferences were well answered by most of the students. The other categories only got a few responses, and so are not as useful in drawing general conclusions.




Figure 4.2: Practicing per week in the three programs.

4.2.2 Interview population details

The students that I interviewed were an oboist, two violists, and a saxophonist. Additionally they listed as secondary instruments piano, guitar, and electric guitar. One of the students was enrolled in the Conservatory, two were in the College, and one was a Double Degree student. Two of the students were not taking any private lessons on their primary instrument, and two had regular weekly lessons.

4.3 Results and Conclusions

4.3.1 Musical feature interest

Since the primary questions of my study were about what musical features musicians were interested in visualizing, one significant result that emerged from both the survey and interview data is that there was a lot of disagreement among musicians about this.

There were two main questions on the survey that asked about musical feature interest. On the first, I asked musicians to rate their interest in the four categories of “Pitch and Intonation,” “Tone and Timbre,” “Loudness and Dynamics,” and “Attacks and Releases,” which I guessed could be four features of interest. The participants rated each feature on

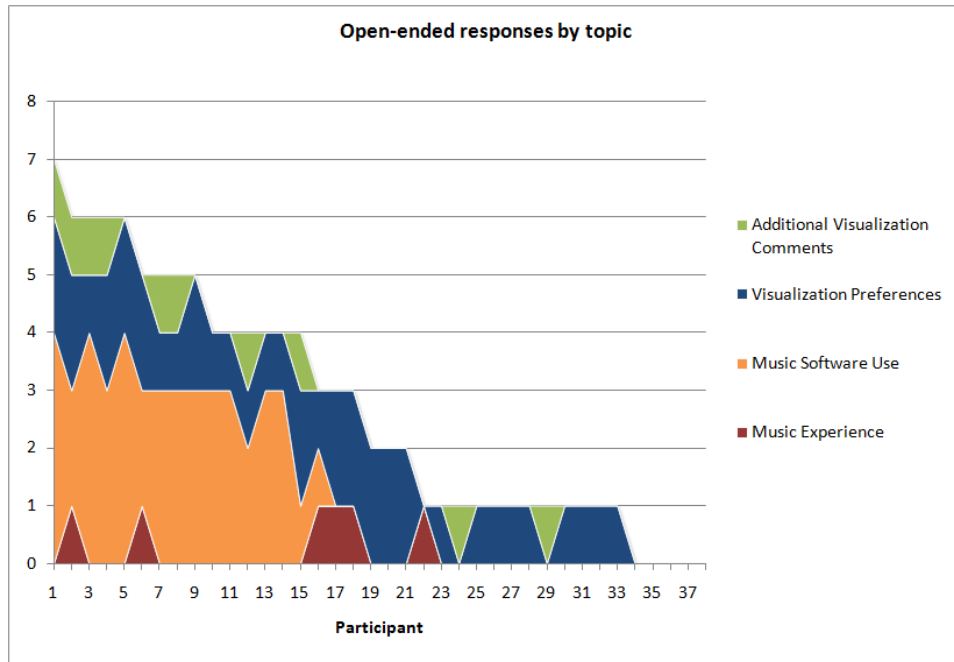


Figure 4.3: Response quality in four open-ended question topics

a scale of “Not interested,” “Somewhat interested,” and “Very interested.” The results are summarized in Figure 4.4. The three interest levels were converted to 0, 1, and 2, respectively. Then each student’s average interest level across the feature categories was subtracted from his interest level in each category to account for students that were overall more enthusiastic or less enthusiastic than others. The numbers in the figure show the average over the set of students of each of these adjusted preference levels in the four categories. The prominent lack of interest in Loudness and Dynamics, relative to the other three categories, was found to be significant with $p = 0.0048$. The higher interest in Tone and Timbre was found to be significant with $p = 0.055$.

This question was followed by an open-ended question asking the student to explain their answers further. While the previous feature interest question shows that Tone and Timbre are preferred over Loudness and Dynamics, the responses to this question revealed that there is a lot of disagreement in why certain features would be better to visualize than others.

The figure shows that 4 people wrote that Tone and Timbre was subjective. However, of these 4 people, two had a moderate interest in visualizing tone, one was in favor of visualizing tone, and one was strongly against it. The reasons that were given for this varied on the one

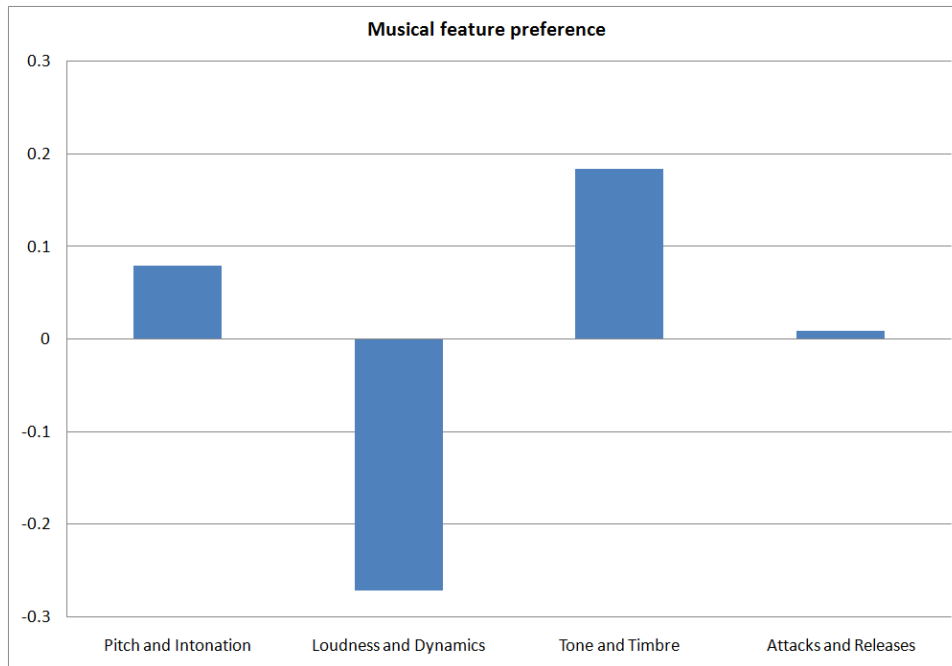


Figure 4.4: Adjusted interest in each feature category.

end from “Tone is a rather personal hearing thing, which is why I wouldn’t want a computer judging it” to “Timbre and Intonation are the hardest to be objective about.”

In the category of Pitch and Intonation, 2 people said that this feature was subjective. Both of these people were weighted towards not visualizing Pitch and Intonation, for the reasons that a computer might have difficulty with nonstandard scales and because pitch is simply too fuzzy to nail down. One additional person mentioned the potential difficulty that a computer could have with nonstandard tunings, writing that “it would always tune perfectly with a piano.”

On the other hand, 3 people mentioned that Pitch and Intonation are objective, and all three of these people were in favor of visualizing pitch. The reasons given by all three people were variations on “Pitch is just about the least subjective parameter listed and seems well suited for a computer.” The musicians surveyed have at least two very different interpretations of pitch. I could not identify any other factor that might contribute to this division. All five of these students were enrolled in the Conservatory, although some were Double Degree students. All of them also play stringed instruments. Computer use did differ between the two groups, but with such small numbers it is impossible to say how much this

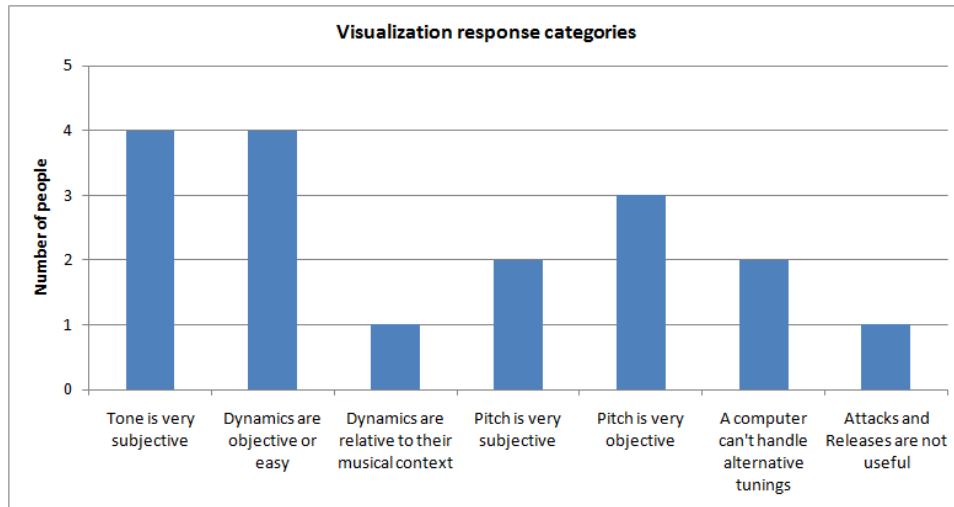


Figure 4.5: Reasons for visualization feature preferences.

could explain the difference.

The people in the “Dynamics are objective or easy” category were all against the visualization of dynamics. One person wrote that he was “Less interested in dynamics because that is something I think I can judge easily on my own” and several of the other three had similar reasoning.

One possible explanation for this variety of opinions about visualization features is that musicians that play different instruments have different feature priorities. Four students made specific mention of instrument specific needs as explaining their feature preferences. These students play the violin, piano, double bass, and French horn. The violinist pointed out that attacks would be “not as effective for a string player as a wind player.” The pianist wrote that “the way I approach the key is important. A lot of my practice time is devoted to deciding how I will attack the keys, how to come off the keys.” The bassist mentioned how useful it might be to observe changes in timbre as the bow is drawn across the strings. Finally the French horn player emphasized pitch and intonation as very important because “the French horn can have a tendency towards quite a bit of intonation problems.” Overall, students playing different instruments have identified different features as particularly useful for their instrument, perhaps implying that the type and content of an effective visualization is fundamentally related to the instrument of the musician.

It is intuitive that different musicians necessarily have different preferences, perhaps

due to general skill level or because of particular areas of technique that the musician has trouble with. Five musicians mentioned specifically their own personal difficulties and how those colored their feature choices. For example, one student wrote that “[Pitch/Intonation and Attacks/Releases] are least consistent in my playing,” and selected those as her top priorities.

While it might be obvious that different musicians playing different instruments would make different feature choices, it is interesting to see the degree of that difference. This conclusion was further supported by the interview data. Because the interviewees played 3 different instruments, I had the opportunity to observe how instrumental differences affected the player’s perception of the visualizations. Also, one of the interviewees agreed to also try playing the piano, which provided additional insight. On the other hand, because two of the students were violists, I could also see how instrument alone did not determine the student’s reaction.

At the start of the interviews, I asked each participant about personal difficulties they had with playing music, and with playing their instrument. The responses are summarized in Figure 4.6. The answers seem to implicate Timbre as the most popular category, agreeing with the survey results. On the other hand, when the participants were actually shown the visualizations, the responses were much more varied.

Instrument	Responses
Sax	<ol style="list-style-type: none"> 1. Getting the low range soft 2. Getting a not-too-brassy, blendable tone
Viola	<ol style="list-style-type: none"> 1. Maintaining tone and technique at the same time 2. Keeping the hands relaxed
Viola	<ol style="list-style-type: none"> 1. Fast passages – technique 2. Tone quality in the high range
Oboe	<ol style="list-style-type: none"> 1. Developing a flexible, versatile tone quality 2. Reed quality and related trouble

Figure 4.6: Interviewee instrument-specific problems.

The following sections summarize the interviewee responses to each of the three types of visualizations in the prototype.

4.3.2 Waterfall visualizations

The biggest problem with this visualization that became clear in the interviews was that there was insufficient contrast both between the dark background and the colors, and between the colors themselves. The red variant (see Figure 3.3) was identified by 3 of the 4 interviewees as more difficult to see than the blue variant, which appeared to be brighter. On the other hand, the color variation was more apparent on the red variant, which 2 interviewees noted.

The Waterfall visualizations were both identified by participants as providing helpful information on loudness, an unintended effect since loudness is only implicitly represented here through the frequency spectrum. One participant said that the overtones were very useful to visualize. Another said that for beginners it might be particularly useful, with respect to showing dynamics. The oboist mentioned that it might prove useful as a way of comparing the characteristics of different reeds, which he said was a difficulty particular to the oboe. Interestingly this type of visualization also turned out to work very well with piano music. One of the violists tried it out with a piano, and said that it reflected very well his note attacks, which a survey respondent identified as being of high importance to pianists.

4.3.3 3-Line visualizations

With these visualizations (see Figure 3.4), the overarching theme was that presenting 3 graphs simultaneously was not helpful. This was for two reasons. The first was that it was difficult to concentrate on three graphs at once, which 3 interviewees spoke about. The second was that on different instruments certain individual graphs did not work as well, and some worked better. For example, one violist found the pitch plot to be useful for noticing small inconsistencies in tone that caused the pitch estimation algorithm to give random results for a fraction of a second. On the other hand, the saxophonist felt that the pitch plot was not that useful, but identified loudness and brightness as interesting. While for some of the interviewees, the brightness plot exhibited little or unpredictable change, rendering more confusing than valuable, it reacted very well to the saxophone. Since the saxophonist was particularly interested in observing his timbral changes, he preferred this plot to the other

two, although he also liked the brightness plot. Unlike the saxophonist, the oboist did not experience any meaningful change in brightness while he played. Also, the pitch estimation algorithm appeared to react badly to the oboe’s sound, and had an uncharacteristically high number of octave errors. As a result, he found the loudness plot the most useful of the three.

The intonation plot on the second variant of this visualization was well-received by some players, but not by others. The saxophonist and one of the violists commented that it would be very useful, but the other violist said it would not be that useful for an advanced player. The oboist had particular difficulty in getting it to judge his pitch as accurate. This is probably because the oboe is a very pitch-stubborn instrument; there is a reason orchestras tune to the oboe and not the other way around. Two interviewees also mentioned the danger of growing to rely on an intonation meter, although this danger is present with regular tuners as well. Two students also said that it would be preferable if the intonation plot did not use a fixed scale to judge against, but rather was simply a fine-grained version of the pitch plot.

Since different graphs were preferred by different people, an appropriate solution would be to allow every user to select which graph or pair of graphs they wished to see. Another point that was raised for this visualization in particular was the importance of adding a playback and review feature that would free the musician from watching the visualization while playing.

4.3.4 Spectrum Ring visualizations

This visualization was well received by all four musicians. One in particular was attracted to the black-and-white coloring of this visualization, saying it had “charm” and reminded him of a speaker. While the oboist did not find either of these particularly useful (though he did say it looked cool), the violists and the saxophonist commented on its ability to convey articulation, loudness, and the spectral information.

On the other hand, the violists found it to be too chaotic. One of the violists also said that it tended to reach a peak level of intensity far before he, playing his instrument, did. That is to say he would play louder and louder, and the visualization would eventually stop growing brighter. He therefore suggested enhancing the expressiveness of the visualization, although he liked its representation of soft sounds.

Overall, the Low-to-High variant of this visualization was preferred (see Figure 3.5a). The High-to-Low version was thought to be both less intuitive and dull to look at.

4.4 Broad implications for design

Two themes emerged from the combination of interview and survey data which would be of central importance for any effective educational music visualization software:

- One visualization does not fit all. Personal goals as well as instrument-specific sound characteristics must be supported by a highly flexible and customizable library of visualizations.
- Encouraging self-reliance rather than dependence on the software must at all times be a goal of the software. Effective software must engage the musician in multiple ways that build the musician's understanding and control of their sound. Musicians will be wary about adopting any software that attempts to replace their own ears.

Chapter 5

Future work

There are two main lessons that can be taken from the results of this project. The first is that because of the great variety of musicians, any software intended to be used by a diverse of them must be extremely flexible. The second is that a more detailed study would be required to fully understand the needs of musicians in the area of visualization software.

5.1 More flexible software

Even with a small study like this one, musicians showed that they would have very different requirements for visualization software. While there was some agreement about specific aspects of the prototype, the interviewees also felt very differently about some of the features. For example, given that the saxophonist really liked the Brightness plot but that the oboist did not, a wide library of features would need to be available for the musician to choose from.

Maybe the ideal system would be a mix-and-match type setup where the player could select a set of musical features to extract and then specify exactly how these get mapped to visual elements using some set of parameterized presets. This would allow adjustment of variables like the plotted range of values, which would be a problem if both very high and very low instruments tried to use the software. Neither of their pitch ranges, being very different, could be optimally represented on a fixed-range plot.

An important customization to add, as shown by the reaction to the Waterfall visu-

alizations, would be the ability to choose color schemes. Some players might prefer blue, and some red, for whatever reason. Some might prefer a black and white coloring, as the saxophonist did with the Spectrum Ring visualization.

My prototype could only show real-time visualizations, so it did not allow the musician to record himself and then go back and review the visuals after the fact. None of the interviewees like this, and commented on how it was impossible to concentrate on both the visualization and playing at the same time. This feature would be an important one to add in any further study.

5.2 More detailed user study

There were two aspects of my study that I would expand in order to gain more complete answers to the questions my research was asking. The first is simply the size of the study. The 38 musicians who completed the survey were enough to get some interesting results, but since only some of these people answered the open-ended questions, the usefulness of those was not as great as it would be in a larger study. Also, since the interviews were so helpful in getting detailed information about what musicians thought about the software I had written, many more interviews should be conducted.

The second way in which the study should be increased in size is to develop a wider variety of prototype visualizations. The musicians only saw three types of designs in this project, and given the fact that their preferences didn't converge on any of these, I think that as many different designs as possible need to be tested by users in order to really find out what kinds of designs are most useful.

5.3 Final thoughts

Because there is not a popular precedent for software such as this, continuing work on this project would be a potentially valuable contribution to music education. For signal processing and visualization research, this work represents some of the first steps in developing a compelling technical problem with exciting potential for impact and application.

Bibliography

- [1] J. Callaghan, W. Thorpe, and J. van Doorn. The science of singing and seeing. *Proceedings of the Conference of Interdisciplinary Musicology, Graz, Austria*, pages 1–10, 2004.
- [2] S. Dixon. Analysis of musical content in digital audio. *Graphics and Multimedia: Applications, Problems, and*, pages 1–30, 2004.
- [3] S. Dixon, W. Goebel, and G. Widmer. The performance worm: Real time visualisation of expression based on langners tempo-loudness animation. *Proceedings of the International Computer Music Conference 2002*, 2002.
- [4] S. Dixon, W. Goebel, and G. Widmer. The "Air Worm": An Interface for Real-Time Manipulation of Expressive Music Performance. *Proceedings of the International Computer Music Conference 2005*, 2005.
- [5] R. Erickson. *Sound structure in music*. University of California Press, Berkley, 1975.
- [6] S. Ferguson, A. Moere, and D. Cabrera. Seeing sound: Real-time sound visualisation in visual feedback loops used for training musicians. *IEEE Proceedings of the International Conference on Information Visualization*, 2005.
- [7] P. Herrera-Boyer, G. Peeters, and S. Dubnov. Automatic classification of musical instrument sounds. *Journal of New Music Research*, 32(1):321, 2003.
- [8] A. Johnston, S. Amitani, and E. Edmonds. Amplifying reflective thinking in musical performance. *Proceedings of the 5th conference on Creativity & Cognition*, 2005.

- [9] P. McLeod. *Fast, Accurate Pitch Detection Tools for Music Analysis*. Phd thesis, University of Otago, 2008.
- [10] P. McLeod and G. Wyvill. A smarter way to find pitch. In *Proceedings of the International Computer Music Conference*, page 138141, 2005.
- [11] G. Peeters. A large set of audio features for sound description (similarity and classification) in the CUIDADO project, 2004.
- [12] G. Tzanetakis and P. Cook. Musical genre classification of audio signals. *IEEE Transactions on speech and audio processing*, 10(5):293302, 2002.
- [13] J. Watkinson. *An introduction to digital audio*. Focal Press, Oxford, 2nd edition, 2002.
- [14] P. Wilson. Does real-time visual feedback improve pitch accuracy in singing? *ses.library.usyd.edu.au*, 2006.
- [15] P. Wilson, K. Lee, J. Callaghan, and C. Thorpe. Learning to sing in tune: Does real-time visual feedback help? *CIM07: 3rd Conference on Interdisciplinary Musicology, Tallinn, Estonia*, 2(1):1519, 2007.

Appendix A

Survey Questions

(Questions begin on the next page)

Visualization Software for Musicians

Survey Consent

Thank you for viewing my survey!

Purpose: For my Oberlin College honors project, I am in the process of writing a new program intended to help musicians get objective feedback in a new way as they practice. As part of the design process, it is critical that musicians provide input and feedback. In completing this survey, you will have an impact on the software's design. This will help to make sure the final product is as useful and interesting to musicians as possible.

Confidentiality: The information collected by the survey is anonymous. If you agree, you will be asked at the end of the survey to provide your contact information so that you can be recruited for a follow-up interview. Your contact information will be kept confidential, should you choose to provide it.

Benefits: There are no direct benefits to you for participating in the study. As a musician, you might benefit from the software product that is to be based on the results of this study.

Participation: Participation in the study is voluntary. Refusal to participate involves no penalty. You are free to end the survey and withdraw from the study at any time with no penalty.

The survey takes about 15 minutes to complete.

If you have questions about this research, please contact:

Primary researcher:

Michael Brooks

OCMR 2920

Phone: 440-865-2437

Advised by Professor Richard Salter, Computer Science Department

Institutional Review Board Chair:

Lynne Bianchi

College of Arts and Sciences

Office of the Dean

Phone: 440-775-8410

1. Do you agree to participate in this study?

- I agree to participate
- I do not agree to participate

You must be 18 or older to participate in this study.

2. Are you 18 or older?

- Yes, I am 18 or older
- No, I am younger than 18

Demographic Information

Visualization Software for Musicians

This section collects some background information about you.

3. What year were you born?

4. What is your gender?

5. What is your ethnicity?

6. What is your occupation?

- High school student
- College student
- Teacher
- Professional musician
- Other (please specify)

7. If applicable, what schools are you enrolled in or affiliated with?

- Oberlin College of Arts and Sciences
- Oberlin Conservatory of Music
- Oberlin High School
- Other (please specify)

8. If you are a student, what is your year in school?

Visualization Software for Musicians

9. If you are, will be, or have been in college, what are your main fields of interest?

- | | | |
|---|---|---|
| <input type="checkbox"/> African American Studies | <input type="checkbox"/> East Asian Studies | <input type="checkbox"/> Latin |
| <input type="checkbox"/> Anthropology | <input type="checkbox"/> East European Studies | <input type="checkbox"/> Latin American Studies |
| <input type="checkbox"/> Archeological Studies | <input type="checkbox"/> Economics | <input type="checkbox"/> Law and Society |
| <input type="checkbox"/> Art | <input type="checkbox"/> Education | <input type="checkbox"/> Mathematics |
| <input type="checkbox"/> Astronomy | <input type="checkbox"/> Engineering | <input type="checkbox"/> Multicultural Studies |
| <input type="checkbox"/> Athletics and Phys Ed | <input type="checkbox"/> English | <input type="checkbox"/> Music Program |
| <input type="checkbox"/> Biochemistry | <input type="checkbox"/> Environmental Studies | <input type="checkbox"/> Neuroscience |
| <input type="checkbox"/> Biology | <input type="checkbox"/> Film | <input type="checkbox"/> Philosophy |
| <input type="checkbox"/> Chemistry | <input type="checkbox"/> French | <input type="checkbox"/> Physics |
| <input type="checkbox"/> Chinese | <input type="checkbox"/> Gender, Sexuality and Feminist Studies | <input type="checkbox"/> Politics |
| <input type="checkbox"/> Cinema Studies | <input type="checkbox"/> Geology | <input type="checkbox"/> Psychology |
| <input type="checkbox"/> Classics | <input type="checkbox"/> German | <input type="checkbox"/> Religion |
| <input type="checkbox"/> Cognitive Science | <input type="checkbox"/> Greek | <input type="checkbox"/> Rhetoric and Composition |
| <input type="checkbox"/> Comparative American Studies | <input type="checkbox"/> Hispanic Studies | <input type="checkbox"/> Russian |
| <input type="checkbox"/> Comparative Literature | <input type="checkbox"/> History | <input type="checkbox"/> Sociology |
| <input type="checkbox"/> Computer Science | <input type="checkbox"/> International Studies | <input type="checkbox"/> Theater and Dance |
| <input type="checkbox"/> Creative Writing | <input type="checkbox"/> Italian | <input type="checkbox"/> Third World Studies |
| <input type="checkbox"/> Creativity & Leadership | <input type="checkbox"/> Japanese | <input type="checkbox"/> Undecided |
| <input type="checkbox"/> Dance | <input type="checkbox"/> Jewish Studies | |
| <input type="checkbox"/> Other (please specify) | | |

Use of Technology

This section is about your personal familiarity with technology.

10. How old were you when you first used a computer?

Visualization Software for Musicians

11. Do you currently own a computer?

- Yes
- No

12. How long ago was your computer purchased?

- I don't own a computer
- Less than 1 year ago
- 1 to 3 years ago
- 3 to 5 years ago
- More than 5 years ago
- I don't know

13. Does your computer have a built in microphone?

- I don't own a computer
- Yes
- No
- I don't know

14. Do you have an external microphone that can be used with your computer?

- I don't own a computer
- Yes
- No
- I don't know

15. Have you ever used a microphone with a computer?

For example, to record yourself or for Skype and other voice-chat software.

- Yes
- No
- I don't know

Visualization Software for Musicians

16. About how many hours per day do you spend on each of the following computer-based activities?

	0 hours	less than 1 hour	1 to 3 hours	3 to 5 hours	more than 5 hours
interacting with a computer, for any reason	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
communicating (email, IM, etc.)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
using social networking websites (Facebook, etc.)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
browsing the web	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
listening to music	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
playing games	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
writing computer programs	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Musical Background

This section asks about your experience with music.

17. Have you ever studied a musical instrument or voice?

Yes

No

18. How old were you when you began to study music?

19. What is your primary musical instrument, including voice?

20. For your primary instrument, how many years of experience do you have in each of the following categories?

regular study and practice

regular private lessons

playing with a musical group or ensemble

Visualization Software for Musicians

21. If you play any other instruments (including voice), list them and for how many years you have studied them.

Example:

Tuba, 3 years

22. About how many days per week do you practice music?

- 7 days per week
- 5 or 6 days per week
- 3 or 4 days per week
- 1 or 2 days per week
- 0 days per week

23. About how often do you have music lessons?

- More than once a week
- Once a week
- Once a month
- Less than once a month
- I do not take music lessons

24. If you take music lessons, characterize your teacher's musical experience? (Check all that apply)

- Professional musician
- Professional music teacher
- Amateur musician
- Music student
- Other (please specify)

25. If you teach music students, how many students are you currently teaching?

Visualization Software for Musicians

26. How often do you record yourself playing music in order to study your own playing?

- More than once a week
- More than once a month
- Less than once a month
- Never

27. Please provide any additional comments you have at this point.

Software for Musicians

This section covers your experience with software for studying music.

28. About how many days per week do you have a laptop in the same room with you while you practice music?

- Six or seven times a week
- Two to five times a week
- Once a week or less
- I never have a laptop with me

29. Have you ever used computer software designed to help with your study of music?

For example, this includes programs that help with memorization, tuning and note errors, and practicing.

- Yes
- No

30. Please list any such computer programs you have used to aid your study of music.

Helpful instructions:

The next two questions ask you to criticize any music software from the list you entered in the above question.

Visualization Software for Musicians

31. Choose a program from the list you entered.

Describe one aspect of it that you found particularly helpful, effective, or appealing, and explain why.

32. Choose a program from the list you entered.

Describe one aspect of the program that hindered its usefulness, was badly designed, or was annoying, and explain why.

33. How often do you use any of this software while you practice music?

- 6 or 7 times per week
- 3 to 5 times per week
- 1 to 2 times per week
- 0 times per week
- I have never used any music software

34. How often do you use this software during your music lessons?

- At every lesson
- Every 2 or 3 lessons
- Every 4 lessons or less often
- I have never used software during a music lesson

35. Please provide any additional comments you have at this point.

Real-time Music Software

This section asks about your ideas about real-time music visualization.

Visualization Software for Musicians

Almost done!

Important explanation:

Imagine that you could have a software program running on your laptop in front of you while you practiced. It could record your playing and run calculations as you play. For example, it could calculate pitch, loudness, and various other musical qualities. The calculations would be used to generate a moving image on your screen that would change to reflect changes in your playing.

Several of the following questions are open-ended, so please answer them as completely as you can.

36. For each of the following aspects of your playing, how interested would you be in visually observing how it changes over time while you play?

	Not interested	Somewhat interested	Very interested
Loudness and dynamics	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Tone or timbral characteristics	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Attacks and releases	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Pitch and intonation	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Please explain your choices for each of the above items:

37. Please describe any other specific aspects of your playing that you would want to see visualized on a computer as you play.

38. Some software already exists that performs some of the functions described above. If you have ever used a program that seems related, please list it below.

39. Please provide any additional comments, questions, or suggestions you have at this point.

Interview Request

Visualization Software for Musicians

Some of the people who completed this survey will be contacted to arrange in-person interviews. Interviews will cover in greater depth the same topics as in the survey. It will also involve testing an early version of my software, and helping improve the design. In order to be interviewed, you must provide a means of contacting you. This information will be kept private and will never be used for any other purpose.

40. If you wish to be considered for such an interview, please provide your email or phone number so that a meeting time can be arranged.

Email Address:

Phone Number:

Survey Complete

Thank you for taking the time to complete this survey! Your answers are very much appreciated.

Appendix B

Interview Questions

There was some variation in these questions from interview to interview.

1. What kind of computer do you have? Do you like it? What are your favorite features?
2. You say you never bring your laptop with you when you practice. Why not?
3. Where do you practice?
4. Have you ever heard of any other students using music software?
5. You say you never record yourself. Why not?
6. What do you find hardest about playing music?
7. Hardest about playing the viola?

Now explain the project in more detail.

The program uses a microphone. As time passes, it analyzes the musician's sound. It provides a real-time visualization that reflects the sound. Two main types of visualization: history and snapshot.

1. What kinds of purposes do you think this software might serve?
2. Before seeing what I have, what are you imagining this software does? (Maybe draw a picture?)

Explain the possible project goals:

1. Provide objective feedback during individual practice
2. Help make practicing interesting

1. What do you think about these goals?

Switch to visualization. Repeat these questions for each visualization.

1. Play your instrument for a couple of minutes. Tell me what you're thinking as you play. How do you interpret the display?

Explain what is going on in the visualization.

1. How could the visualization better convey its information?
2. How do you think this visualization might be useful? Can you see yourself using it?