# Instrument Timbres and Pitch Estimation in Polyphonic Music 

A Thesis<br>Presented to<br>The Academic Faculty

by
Beatus Dominik Loeffler

In Partial Fulfillment<br>of the Requirements for the Degree<br>Master of Science in Electrical and Computer Engineering

Georgia Institute of Technology
May 2006

# Instrument Timbres and Pitch Estimation in Polyphonic Music 

Approved by:

Professor Chin-Hui Lee
School of Electrical and Computer Engineering
Georgia Institute of Technology, Advisor

Professor Aaron D. Lanterman
School of Electrical and Computer Engineering
Georgia Institute of Technology
Professor David V. Anderson
School of Electrical and Computer Engineering Georgia Institute of Technology

Date Approved: April 10, 2006

Michael and Waltraud Löffler

## ACKNOWLEDGEMENTS

I want to thank my advisor Dr. Chin-Hui Lee for his guidance and patience, Dr. Lanterman and Dr. Anderson for additional input, Dr. Rüdiger Naumann-Etienne and his foundation for generously sponsoring my studies throughout the two past years, and Carlton Parker of the World Student Fund for his administrative and organizational work. Without their support, this work would not have been possible.

## TABLE OF CONTENTS

DEDICATION ..... iii
ACKNOWLEDGEMENTS ..... iv
LIST OF TABLES ..... viii
LIST OF FIGURES ..... ix
SUMMARY ..... x
I INTRODUCTION ..... 1
1.1 Organization and Goals of this Project ..... 1
1.2 Pitch ..... 1
1.2.1 Definition ..... 1
1.2.2 Harmonics, Overtones, and Partials ..... 2
1.2.3 Pitch Spacing in Western Music ..... 4
1.2.4 Musical Scales and the Weber-Fechner-Law of Perception ..... 5
1.2.5 Motivation for Pitch Estimation ..... 6
1.3 Timbre of Instruments ..... 6
1.3.1 The Three Phases of a Tone ..... 7
1.3.2 Pitch Dependency of Timbre ..... 8
1.3.3 Characteristics of some Specific Instruments ..... 9
1.4 Human Auditory Perception of Pitch ..... 12
1.4.1 $\quad F_{0}$ Extraction Only ..... 12
1.4.2 Minimum Difference of Adjacent Harmonics ..... 12
1.4.3 Highest Common Factor ..... 13
II OVERVIEW OF PITCH ESTIMATION APPROACHES ..... 15
2.1 Spectral Smoothing Using "Specmurt Anasylis" (sic) ..... 15
2.2 An Expectation Maximization for a Constrained GMM Using an Informa- tion Criterion ..... 16
2.3 A General Iterative Multipitch Estimation Algorithm ..... 16
2.4 Partial Searching Algorithm with a Tree Search and Dynamic Programming ..... 17
2.5 Independent Subspace Analysis, Prior Subspace Analysis, Generalized Prior Subspace Analysis ..... 17
III A GMM APPROACH TO PITCH ESTIMATION ..... 19
3.1 General Fundamentals of Expectation Maximization Algorithms ..... 19
3.2 Fundamentals of Gaussian Mixture Models (GMM) ..... 21
3.3 An EM Algorithm for GMM Parameter Estimation ..... 22
3.4 Missing Data in EM/GMM Analysis of Monophonic Music (Number of Components) ..... 22
3.5 Reinterpretation of the Magnitude Spectrum as a Spectral Density ..... 24
3.6 Modification of the EM Algorithm to Accept the "Frequency Density" whose Parameters Are to Be Estimated ..... 25
3.7 Missing Data in EM/GMM Analysis of Polyphonic Music (Number of Tones, Number of Components) ..... 26
IV HEURISTICS FOR INITIAL ESTIMATES ..... 27
4.1 Limiting the Frequency Range ("Cut-Off") ..... 27
4.2 Leveraging the Spacing of the Harmonics ..... 29
4.2.1 Peak-Picking Function ..... 29
4.2.2 Peak Filtering Heuristic ..... 30
4.2.3 Core Heuristic: Basic Idea ..... 30
4.2.4 Interpolation for "Missing Peaks" ..... 32
4.3 Dynamic Determination of Number of Mixtures ..... 32
4.4 Extension to Polyphonic Cases ..... 32
4.4.1 Maximum Number of Polyphonic Tones ..... 32
4.4.2 Iterative Scheme ..... 33
4.4.3 Constrained EM Algorithm ..... 34
4.5 Chapter Review ..... 35
V PRELIMINARY RESULTS AND ANALYSIS ..... 36
5.1 Implementation Remarks: Single-Pitch Prototype ..... 36
5.1.1 Results for Single Pitch Cases ..... 38
5.1.2 Potential Difficulties in Adopting GMM For Pitch Estimation ..... 38
5.2 Discussion of the Final Implementation ..... 39
5.2.1 Summary of Issues ..... 44
VI CONCLUSIONS AND DIRECTIONS FOR FUTURE RESEARCH ..... 46
6.1 Conclusions ..... 46
6.2 Directions for Future Work ..... 46
6.2.1 Account for Harmonic Overlapping of Pitches ..... 46
6.2.2 Follow Klapuri Closely ..... 46
6.2.3 Confidence Score for Missing Harmonic Interpolation ..... 47
6.2.4 Tune Some of the EM Numerical Parameters such as Termination ..... 47
6.2.5 Set Up an Instrument Identifier on the Separated Sources ..... 47
APPENDIX A - PITCH TABLE ..... 48
APPENDIX B - CODE LISTINGS ..... 49
REFERENCES ..... 79

## LIST OF TABLES

1 Common pitch spacing ..... 5
2 HCF example ..... 13

## LIST OF FIGURES

1 Harmonics of strings [26] ..... 3
2 Harmonics of an oboe C4 tone, DFT magnitude ..... 4
3 Onset of clarinet and sax, from [12] p. 205 ..... 7
4 LTAS of the viola's F4 and A4 tones ..... 8
5 LTAS of the cello's F4 tone ..... 9
6 Pitch overlap in the string instrument family ..... 10
$7 \quad$ Spectrum of the oboe's C4 tone (no LTAS) ..... 11
8 The piano's amplitude over time ..... 11
9 General multipich estimation principle by [17] ..... 17
10 Many components arising from low pitch (bass trombone, $\mathrm{F} \# 1$ ) ..... 23
11 Only one component arising from high pitch (flute, B6) ..... 24
12 Illustration of the behavior of the EM algorithm for Gaussian Mixture Models ..... 26
13 Reducibility of the spectrum ..... 28
14 Spectrogram result without spectral cutoff ..... 28
15 Spectrogram result with spectral cutoff ..... 29
16 Summary of cutoff ranges ..... 30
17 GUI screenshot of single-pitch prototype ..... 37
18 GUI screenshot of the final implementation ..... 40
19 Heuristic working well in a high-pitch example ..... 41
20 Excess pitches in medium-pitch example ..... 42
21 Excess pitches in a medium-pitch example - without spectral smoothing ..... 43
22 Excess pitches in medium-pitch example - escaping the heuristic ..... 44

## SUMMARY

In the past decade, the availability of digitally encoded, downloadable music has increased dramatically, pushed mainly by the release of the now famous MP3 compression format (Fraunhofer-Gesellschaft, 1994). Online sales of music in the US doubled in 2005, according to a recent news article [2], while the number of files exchanged on P2P platforms is much higher, but hard to estimate.

The existing and coming informational flood in digital music prompts the need for sophisticated content-based information retrieval. Query-by-Humming is a prototypical technique aimed at locating pieces of music by melody; automatic annotation algorithms seek to enable finer search criteria, such as instruments, genre, or meter. Score transcription systems strive for an abstract, compressed form of a piece of music understandable by composers and musicians. Much research still has to be performed to achieve these goals.

This thesis connects essential knowledge about music and human auditory perception with signal processing algorithms to solve the specific problem of pitch estimation. The designed algorithm obtains an estimate of the magnitude spectrum via STFT and models the harmonic structure of each pitch contained in the magnitude spectrum with Gaussian density mixtures, whose parameters are subsequently estimated via an ExpectationMaximization (EM) algorithm. Heuristics for EM initialization are formulated mathematically.

The system is implemented in MATLAB, featuring a GUI that provides for visual (spectrogram) and numerical (console) verification of results. The algorithm is tested using an array of data ranging from single to triple superposed instrument recordings. Its advantages and limitations are discussed, and a brief outlook over potential future research is given.

## CHAPTER I

## INTRODUCTION

### 1.1 Organization and Goals of this Project

This introductory chapter explains some fundamental knowledge and frequently used terms about pitch, timbre and human perception of pitch. Some special cases of instrument timbres and pitch perceptional peculiarities are treated in more detail to demonstrate the breadth of the field of study.

Chapter 2, Overview of Pitch Estimation Approaches, gives a brief overview of some previous approaches to pitch estimation in the simpler monophonic context. Also, a general algorithm scheme for developing polyphonic pitch estimators is enumerated.

Chapter 3, A GMM Approach to Pitch Estimation, introduces key mathematical concepts used in a subsequent implementation of a pitch estimation system, first for a monophonic case, then for the more general polyphonic case. A Gaussian Mixture Model (GMM) is used to model the harmonics in the spectral magnitude. Model parameters are estimated via an Expectation Maximization algorithm.

The subsequent chapter, Development of Heuristics for Initial Estimates, discusses ways our EM algorithm can be initialized in the case of pitch estimation, both for single and multiple pitch cases. Issues related to running the algorithms on real data are discussed, and further heuristics are developed.

Finally, Chapters 5 and 6 discuss problems with and limitations of our current implementation and propose various future enhancements.

### 1.2 Pitch

### 1.2.1 Definition

The term pitch describes a psychoacoustic sensation of the auditory system and is therefore subjective. It should not be confused with frequency or note, although it has intimate
relationships with those terms.
A particular note on the keyboard of a piano can be matched to a pure sinusoidal single frequency by a listener, but the frequency spectrum of the signal produced by the piano will carry more than that sinusoidal frequency. Instead, it will show a more complex mixture of different frequencies with different weights, enabling the listener to clearly discern between instruments (Chapter 1.3). Psychoacoustic studies try to discover how the human hearing system matches these two quite different signals.

Frequency denotes a precisely defined analytic term, such as the number of vibrations of a string per second or the resulting vibration of atmospheric pressure carrying the signal to the listener's ear. The term will often be used in the context of the DFT or spectrograms, which contain DFT analyses.

For the purpose of this thesis, a note is simply the standardized name of a pitch, such as $\mathrm{A} \#$, or the graphical representation of it on the staff. It can thus be considered a synonym for pitch, if the definition relating the two is known. This leads to the topic of pitch standards, and there is a long history of such. Today, the most commonly used pitch standard is ISO 16 [14] (Appendix).

In ISO 16 , the A above middle C (A4) must be tuned to 440 Hz . This note is the center of attention in pitch standards, and all other pitches are typically tuned with respect to it.

### 1.2.2 Harmonics, Overtones, and Partials

As mentioned before, a pitch may be considered more complex than a single frequency; it is a mixture of sinusoids of different frequencies with different weights that are perceived to be equivalent to a single frequency. These "other components" are the so-called harmonics, which make the tone richer and more interesting to the listener. In fact, in music, we almost never deal with pure tones, that is, pitches of single frequency. These sound bland and boring, and the possibility of electronically producing a pure tone has not been around for long, considering the entire history of music. This mixture of frequencies is a kind of "signature" or "fingerprint" called timbre, which is dependent on the instrument's physical specifications (length, diameter, width, girth, specific shape, density of the materials used,


Figure 1: Harmonics of strings [26]
etc). We can think of timbre as the "color" or "quality" of sound.
For an example of timbre, when a string vibrates, it does not vibrate like only a single sinusoid. It will vibrate with a basic frequency $\omega_{0}$, but it will also vibrate at integer multiples of that basic frequency.

Expressed in the frequency domain,

$$
\begin{equation*}
\omega_{m}=m \cdot \omega_{0} . \tag{1}
\end{equation*}
$$

The same principle can be observed with the air pressure in a trumpet or other winds, or the human voice. The above description is a simplification; the partials of real instruments have amplitudes and frequencies that change with time. In particular, they may not line up at exact integer multiples, and even the fundamental may change frequency over time due to a vibrato.

Figure 2 shows a Discrete Fourier Transform of a short (64ms) piece of an oboe's C4 pitch. Recall that the Fourier transform of a single sinusoid is

$$
\begin{equation*}
F\left[\sin \left(\omega_{c} t\right)\right]=\frac{1}{2} j\left(\delta\left(\omega+\omega_{c}\right)-\delta\left(\omega-\omega_{c}\right)\right) . \tag{2}
\end{equation*}
$$

This makes it clear that if we show only the positive frequency range and if we take the magnitude, we will see a single delta peak in the spectrum. We should discover something similar for the DFT magnitude of a musical pitch. Each harmonic should yield a spike located near an integer multiple of the fundamental (typically the spike with the lowest frequency).


Figure 2: Harmonics of an oboe C4 tone, DFT magnitude

Figure 2 does show significant peaks. Of course, they are not true delta peaks, since the analysis is over a finite window of data and real instruments cannot be perfectly periodic. We also see that the peaks of the harmonics are located at or near integer multiples of the fundamental.

The term overtone is a general term for referring to all the frequency components besides the fundamental. Overtones will be categorized into harmonics that are integer multiples of the fundamental, and partials, which consist of the rest.

This nomenclature is not standard; most authors consider all harmonics to be partials (but not all partials to be harmonics). Our convention allows us to avoid the cumbersome expression "nonharmonic partials." (see Figure 2)

### 1.2.3 Pitch Spacing in Western Music

The following table shows the pitches of a chromatic scale from C 4 to C 5 in western music.
This scheme of musical tuning is called a 12-tone equal temperament system, in which an octave is subdivided into a series of equal pitch ratios. The pitch is not linearly spaced across the octave (the last row shows the numbers for a hypothetical linear spacing) since humans perceive frequencies logarithmically. Mathematically speaking, the pitches of the scale follow a geometric sequence.

Table 1: Common pitch spacing

| Note | C 4 | $\mathrm{C} \# 4$ | D 4 | $\mathrm{D} \# 4$ | E 4 | F 4 | $\mathrm{~F} \# 4$ | G 4 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| Pitch/Hz | 261.6 | 277.2 | 293.7 | 311.1 | 329.6 | 349.2 | 370 | 392 |
| Lin. | 261.6 | 283.4 | 305.2 | 327 | 348.8 | 370.1 | 392.5 | 414.3 |
| Note | $\mathrm{G} \# 4$ | A 4 | $\mathrm{~A} \# 4$ | B 4 | C 5 |  |  |  |
| Pitch/Hz | 415.3 | 440 | 466.2 | 493.9 | 523.3 |  |  |  |
| Lin. | 436.1 | 457.9 | 479.7 | 501.5 | 523.3 |  |  |  |

### 1.2.4 Musical Scales and the Weber-Fechner-Law of Perception

Human perception, in general, is not associated with a linear scale. A general law describing human response to physical stimulus was discovered by Ernst Heinrich Weber and Gustav Theodor Fechner in 1860 [6]. A later extension was made by Stanley Smith Stevens in 1957 [27].

Stimulus $(S)$ and perception $(p)$ can be related by the differential equation

$$
\begin{equation*}
d p=k \cdot \frac{d S}{S} \tag{3}
\end{equation*}
$$

which is solved by

$$
\begin{equation*}
p=k \cdot \ln \left(\frac{S}{S_{0}}\right)=k \cdot \ln (2) \cdot \log _{2}\left(\frac{S}{S_{0}}\right)=\tilde{k} \cdot \log _{2}\left(\frac{S}{S_{0}}\right), \tag{4}
\end{equation*}
$$

where $\tilde{k}$ is an experimentally determined constant, and $S_{0}$ denotes a stimulus threshold below which nothing is perceived. Vice versa, this means that if we scale the perception linearly, such as with the musical scale, the stimulus will be a certain exponential function. In the case of pitch and the table above, it is

$$
\begin{equation*}
S=S_{0} \cdot 2^{p / \tilde{k}}=S_{0} \cdot 2^{p / 12} \tag{5}
\end{equation*}
$$

To verify this, set $S_{0}=261.6$ and let $p=1,2, \ldots 13$. The spacing of the pitches are important in note transcription systems, in which a nearest-neighbor classification based on this result might be used.

A measure frequently appearing in the literature is the cent, which equals one hundredth of an equal tempered semitone. This is equivalent to one twelve-hundredth of an octave, since there are twelve semitones to the octave. Thus, if the frequency $f_{2}$ is said to be a cent
higher than $f_{1}$, then

$$
\begin{equation*}
\frac{f_{1}}{f_{2}}=2^{1 / 1200} \approx 1.00057779 \tag{6}
\end{equation*}
$$

Humans can notice a difference in pitch of about 5 to 6 cents, which is an astoundingly high degree of resolution.

Despite this result for the pitch scale, a deviation can be observed for very high frequencies such as the $7^{t h}$ and $8^{t h}$ octaves and higher. Pitches perceived and performed by absolute pitch possessors and flutists, respectively, exceed the theoretical value, and do so increasingly the higher the note. An experimental study and discussion of physiological origins can be found in [22]. For the purpose of this project, this phenomenon shall not be taken into account

### 1.2.5 Motivation for Pitch Estimation

Some applications of pitch estimation include

- Melody note detection and query-by-humming (information retrieval)
- Score transcription systems / conversion to MIDI
- Music visualization
- Polyphonic sound separation, e.g. for instrument identification


### 1.3 Timbre of Instruments

The general ideas of pitch, timbre and the relationship of the harmonics have been presented previously. This section makes some more general remarks and reviews specific characteristics of instrument timbres.

In particular, timbres differ mostly in the weight and number of the harmonics, and differ to a much lesser extent in the basic shape of the peaks (this fact will be exploited later in our mathematical modeling). The spectral magnitude of most acoustic instruments approximately follows a $1 / f$ curve above a certain frequency, so most of the significant timbral parameters are found in the lower frequency bands. This fact will also be exploited later.


Figure 3: Onset of clarinet and sax, from [12] p. 205 fig:peakCutoffs

### 1.3.1 The Three Phases of a Tone

So far, we have assumed that timbre is constant over time. The three temporal phases of a tone (onset/attack, steady state, offset/release/decay) each help define the overall timbre of an instrument. The attack particularly enriches the musical listening experience and seems to reveal important physical parameters not contained in the other phases. The decay, on the other hand, is far less important as far as timbre recognition is concerned, since the timbral characteristics can be clearly heard during the (often much longer) steady-state phase, which inherently always precedes the offset phase.

A typical spectral magnitude of a $64 m s$ frame of a monophonic piece in the steady state is shown in the introduction, but for better analysis of this phase, it is useful to look at time-averaged spectra to be certain of observing general tendencies (often called long-term average spectra, or LTAS). Any noise present in the signal will then approximately average to zero.

For the onset phase on the other hand, LTAS would be detrimental, since we have an explicit time dependency of the spectral features. This can be easily demonstrated by some spectrograms, as in Figure 3.

The two instruments in Figure 3 not only differ in number of harmonics and their weighting (the latter is not visible here), but especially in time onset structure of harmonics.

These two instruments, clarinet and saxophone, could easily be differentiated from an organ. As also shown in [12], a principal 8' open flue organ stop has a second harmonic that begins much earlier than the fundamental (about 35 ms ).

### 1.3.2 Pitch Dependency of Timbre

Generally, the timbral parameters in all phases are pitch-dependent. This is one reason why a reliable pitch estimation system is a prerequisite to instrument identification. Figure 4 shows the notes A4 and F4 performed by the same viola in the same environment, on different strings.


Figure 4: LTAS of the viola's F4 and A4 tones

Although the pitches are quite close, the weights of the harmonics vary considerably between these two tones. This is not a special case; it seems impossible to create a single weighting profile for an instrument across all its pitches.

Figure 5 shows the LTAS of an F4 from a cello. Of course, telling instruments apart within a family such as the strings (violin, viola, cello, bass) is challenging. This is further complicated by the range overlap between different instruments, as shown for the string family in Figure 6.

Furthermore, there is not just an overlap between instruments but also strings on the same instrument (see also Figure 6). This would have to be considered when training a model for an instrument, dramatically increasing the number of samples needed.


Figure 5: LTAS of the cello's F4 tone

### 1.3.3 Characteristics of some Specific Instruments

Many instruments have distinctive characteristics. For instance

- Figure 7 show that the peak "feet" have a distinctive shape. One side typically has a sharp upward or downward turn, but on the other side of the valley the change is slight.
- Plucked instruments (such as guitar, strings, harp) and the piano have a particular short-time energy or amplitude envelope. Figure 8 shows the time samples of a piano note, thus indicating the energy envelope.
- The trumpet has strong and regular harmonics up through high frequencies.
- Overblowing of wind instruments or organ pipes allows another octave (or octave and perfect fifth for stopped pipes) to be reached, that otherwise would not be available to the instrument. The timbre changes audibly.
- Stopped organ pipes sound an octave lower than open pipes at the same length, but only the odd-numbered harmonics are present. Conversely, an open pipe and a closed pipe at half the length produce the same pitch, but their timbres are quite different because of the lack of even-numbered harmonics in the signal of the closed pipe ([12] p.176).


Figure 6: Pitch overlap in the string instrument family


Figure 7: Spectrum of the oboe's C4 tone (no LTAS)


Figure 8: The piano's amplitude over time

### 1.4 Human Auditory Perception of Pitch

Perceptional studies show that the auditory system is not straightforward. It is still not clear exactly how the auditory cortex, a part of the brain, processes information or how training is performed. However, indications point towards a tonotopical organization of the auditory cortex, that is, certain cells are sensitive to specific frequencies ([23, 30]). Spectrally rich pitches are easier to discriminate than those of sinusoidal tones [28]. Further, it appears that the brain separates encoding of spectrally complex information from temporally complex information [28]. In a similar way it may separate phonetic content from musical content [29].

However the brain processes auditory information internally, some important results may be derived using a black-box approach in which the listener is asked to give subjective perceptual judgement on specially prepared sound sources. The first experiments of this kind were performed by Helmholtz [32], who also formulated the place theory of pitch perception, which will be discussed in more detail here. The counterpart, the temporal theory of pitch perception, will be skipped because it does not directly relate to our spectrum-based pitch estimation algorithm. Note that neither of the theories alone can sufficiently model actual perception in all cases.

### 1.4.1 $\quad F_{0}$ Extraction Only

The most straightforward pitch determining method extracts only the fundamental, i.e. the first harmonic. The problem is that some sounds may be missing the fundamental, yet the perceived pitch is still the same as if it were present. Examples of sounds of this category are overblown flutes and organs. This fact is also used in small loudspeakers that cannot reproduce lower frequencies well.

### 1.4.2 Minimum Difference of Adjacent Harmonics

The pitch perception of a sound with a missing harmonic suggests that the brain looks at the spacing between adjacent peaks for pitch determination, so the same idea is used here.

$$
\begin{equation*}
\text { pitch }=\min _{j \in 1 . . M-1} f_{j+1}-f_{j} . \tag{7}
\end{equation*}
$$

While this method works for navigating around the simple case where only the fundamental is missing, different mixtures can be created for which this method does not work either. An example is presented in Table 2, in which there are only three harmonics at rather high frequencies relative to their distances. exist. The pitch extracted by the minimumdifference method would clearly be 200 Hz , but this does not match interpretation of a human subject.

### 1.4.3 Highest Common Factor

Suppose we have $M$ harmonics present. To overcome the shortcomings of the preceding methods, the highest common factor approach divides each center frequency of a harmonic by a sequence of integers $1 . . N$. $N$ should be sufficiently larger tham $M$. Among this set of numbers, we find the $M$ that are closest to each other so that we will pick one number uniquely from each set associated with a certain harmonic. The pitch is then computed by simply taking the arithmetic mean of these $M$ numbers. Let $\pi_{1} \ldots \pi_{M}$ denote the $M$ closest values. The pitch estimate is then

$$
\begin{equation*}
\text { pitch }=\frac{1}{M} \sum_{i=1}^{M} \pi_{i} \tag{8}
\end{equation*}
$$

In Table 2, there are three harmonics, so $M=3$. The three bold numbers in the table correspond to the three closest ones, prohibiting two or more from the same row.

Table 2: HCF example

| div by $\rightarrow$ <br> harmonic $\downarrow$ | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 1 | 1040 | 520 | 346.7 | 260 | $\mathbf{2 0 8}$ | 173.3 | 148.6 | 130 |
| 2 | 1240 | 620 | 413.3 | 310 | 248 | $\mathbf{2 0 6 . 7}$ | 177.1 | 155 |
| 3 | 1440 | 720 | 480 | 360 | 288 | 240.0 | $\mathbf{2 0 5 . 7}$ | 180 |

The pitch is then determined to be

$$
\begin{equation*}
\text { pitch }=\frac{208+206.7+205.7}{3}=206.8 \approx 207 \tag{9}
\end{equation*}
$$

which corresponds exactly to a human subject's matching of the test tone with a pure frequency.

The reason for looking at these special cases and discussing human perception is that we would like our pitch estimator to work similarly to our brain, in the hope that its results would match what we hear. The heuristics in Chapter 4 do not take all the above discussed cases into account. However, the reader should keep in mind that estimating a pitch is generally a complex task, and that for a more sophisticated system developed in the future, perceptional aspects might assume a higher priority.

## CHAPTER II

## OVERVIEW OF PITCH ESTIMATION APPROACHES

To our knowledge any approach to the problem of pitch estimation can be categorized into one of two general categories: short-time frame based, or spectrogram based. The former looks at isolated signal slices and the latter uses the whole spectrogram of the entire piece. While the short-time slice approach could be viewed as just a special case of the spectrogram approach, the distinction makes sense if we look at the mathematical formulations. Furthermore, the different approaches are often uses in two application contexts, namely audio streaming and entire-file-based analysis.

The first four sections present four approaches for the time-slice technique. The last section summarizes some spectrogram-based techniques. This overview is not meant to be exhaustive, but it should show some breadth and highlight important general concepts.

### 2.1 Spectral Smoothing Using "Specmurt Anasylis" (sic)

In principle, the "specmurt anasylis" approach [24] finds the pitch by reducing the set of harmonics to a single peak by a signal filtering technique. It is meant to produce piano-roll visualizations, which are similar to spectrograms. The piano-roll output has a linear pitch frequency axis and is obtained using a wavelet transform (might also use a CQT). The pitch is modeled as a single delta peak in the frequency domain, and the resulting spectrum $v(x)$ is obtained by the convolution of this delta distribution $u(x)$ (or in the polyphonic case, superposed deltas) with a "common harmonic structure" $h(x)$ :

$$
\begin{equation*}
v(x)=h(x) * u(x) \tag{10}
\end{equation*}
$$

The left side is a known (signal spectrum). We must make an educated guess about the common harmonic structure $h(x)$ (basically, the timbre information), which is crucial to this technique's performance. All that remains is a "simple" deconvolution of the above equation to obtain the pitch estimate(s), which must be picked from the remaining peak(s).

The authors of [24] do not go this far, though, and are satisfied with the visual output and a comparison to the MIDI equivalent. To convert the piano-rolls of their described system to a musical score, a pitch-spelling algorithm would need to be attached (a comparison of such algorithms is given in [20]).

The follow-up paper [25] discusses how to optimize the shape of the common harmonic structure.

### 2.2 An Expectation Maximization for a Constrained GMM Using an Information Criterion

This approach, given in [15], [16], models each "harmonic structure" (set of harmonics for a particular pitch) with a constrained Gaussian Mixture Model, and is thus similar to the formulation in Chapter 3. The difference is that [15] uses a log-frequency transform and an alternate maximization part which does not only constrain the means, as we do here, but the weights as well. The Akaike information criterion (AIC) is used for estimating the number of concurrent pitches, in this paper denoted as $K$. A high parameter (in the range of $K=12$ ) is initially assumed, and then iteratively reduced. The minimizer of the AIC function across these values will determine the best guess of the true value of $K$. Thus, this implementation requires the reestimation of the parameters via EM as many as 12 times. The authors of [15] have not included any analytical or experimental remarks about computational complexity.

### 2.3 A General Iterative Multipitch Estimation Algorithm

The paper titled "Robust Multipitch Estimation for the Analysis and Manipulation of Polyphonic Musical Signals" [17], introduces a general scheme for polyphonic pitch estimation and describes a particular approach to pitch estimation based on Short-Time Fourier Transform (STFT). Bandwise processing of DFT frames is used to come up with bandwise pitch estimates that are then combined into a single score.

The first part of the general scheme, predominant pitch estimation, finds the pitch of the most prominent sound standing out from the background interference of other harmonics and noisy sounds. In the second part, the spectrum of the detected sound is estimated and


Figure 9: General multipich estimation principle by [17]
subtracted from the mixture. The estimation and subtraction steps are then repeated for the residual signal, until a predefined maximum number of pitches has been reached, or no new pitch estimates can be found confidently.

The follow-up paper [31] takes the outputs of the multipitch estimator described in [17] and improves on their accuracy in the least-squares sense while retaining the structure of the sounds. Smooth linear models are constructed for all significant sounds; the parameters of the models are obtained with a least-squares solution.

### 2.4 Partial Searching Algorithm with a Tree Search and Dynamic Programming

The algorithm presented in [34] assumes that the fundamentals for each pitch are given. It recognizes that a particular significant peak can be "shared" by the harmonics of two different pitches, and comes up with the best guess for a solution. It tries to find these respective harmonics via a "relaxed harmonics condition," using a particularly defined tree search algorithm and dynamic programming. A peak-picking heuristic, or some other method of coming up with the prospective harmonic locations on the frequency axis, can be used prior to running the algorithm. The results don't look overly promising at present when the number of pitches is not known in advance. However, to our present knowledge, this approach should be investigated further.

### 2.5 Independent Subspace Analysis, Prior Subspace Analysis, Generalized Prior Subspace Analysis

Independent Subspace Analysis (ISA; [4]) considerably extends Independent Component Analysis (ICA; [13]). First, the one-dimensional signal is projected onto a two-dimensional space using the STFT or Constant-Q Transform [3]. Then, the ICA constraint of requiring
as many mixture observation signals as there are sources is relaxed. Further, dynamic independent components are allowed to account for the nonstationarity of the signal. Sources are tracked by the similarities of dynamic components over small time steps.

Prior Subspace Analysis (PSA; [10], [8]) requires an individual prior subspace for each note of an instrument. Since this is a large collection for most instruments, the first application of PSA to music analysis was an application to drum sounds.

The magnitude spectrogram of the signal is modeled as the superposition of $l$ unknown spectrograms $Y_{j}$, and each of the spectrograms can be represented as the outer product of an invariant frequency basis function and an invariant amplitude basis function:

$$
\begin{equation*}
\mathbf{Y}=\sum_{j=1}^{l} Y_{j}=\sum_{j=1}^{l} a_{j} s_{j} . \tag{11}
\end{equation*}
$$

PSA assumes that there are known frequency basis functions or prior subspaces $a_{p r}$ that are good approximations to the actual subspaces (these must be found using some domain knowledge). Substituting for the $a_{j}$ with these prior subspaces yields

$$
\begin{equation*}
\mathbf{Y}=\sum_{j=1}^{l} a_{p r} s_{j}=A_{p r} s \tag{12}
\end{equation*}
$$

which is solved in PSA by premultiplying the overall spectrogram with the pseudo-inverse of the frequency basis functions, to get an initial estimate for $s$. ICA is used to improve these estimates.

In Generalized Prior Subspace Analysis, given in [9], Equation 11 is essentially reinterpreted in such a way that the prior subspaces are viewed as an undercomplete signal dictionary and that the data is sparse in nature, which results in a new version of the PSA algorithm [9].

## CHAPTER III

## A GMM APPROACH TO PITCH ESTIMATION

This chapter provides the necessary theoretical tools for the pitch estimation approach presented in this document. Our algorithm merely takes spectral information into account, via STFT.

In general, the DFT serves as a good basis for frequency analysis, since its properties are well-explored and implementations using the Fast Fourier Transform (FFT) are efficient. However, other transformations such as the Modulation Spectrogram [11] (originally thought of speech), the Constant-Q-Transform (CQT; [3]), or wavelet transformations make sense [24]. The CQT is mainly motivated by the way it scales the frequency axis: it is linear in the pitches of the chromatic scale discussed in 1.2.3. Generally, psychoacoustic results of pitch perception question the efficiency of the DFT as the basis of the further analysis, but a more detailed study of the suitability of other transforms is outside of the scope of this thesis.

The following sections discuss in detail the mathematical modeling of the harmonics via Gaussian Mixture Models and the Expectation Maximization algorithm designed to estimate the parameters of the model, eventually leading to a set of frequencies (ideally) corresponding to the true center frequencies of the harmonics.

### 3.1 General Fundamentals of Expectation Maximization Algorithms

The Expectation Maximization (EM) technique is a general method for solving maximum likelihood (ML) estimation problems with incomplete data. It was first formally formulated by Dempster, Laird, and Rubin [18] and extended for superimposed signals by Feder and Weinstein [7].

Assume the probability density function of an observable random variable $\mathbf{Y}$ is $f_{\mathbf{Y}}(y ; \Theta)$.

In the EM approach, we hypothesize a "complete data" random variable $\mathbf{X}$, which is nonobservable. They two kinds of data are related by a noninvertible and unknown function $H$,

$$
\begin{equation*}
H(\mathbf{X})=\mathbf{Y} \tag{13}
\end{equation*}
$$

Therefore, the density for $\mathbf{X}$ can be expressed by

$$
\begin{equation*}
f_{\mathbf{X}}(x ; \Theta)=f_{\mathbf{X} \mid \mathbf{Y}=y}(x ; \Theta) \cdot f_{\mathbf{Y}}(y ; \Theta) \tag{14}
\end{equation*}
$$

over the complete support of $H$. Rearranging and taking the logarithm of both sides yields

$$
\begin{equation*}
\log \left(f_{\mathbf{Y}}(y ; \Theta)\right)=\log \left(f_{\mathbf{X}}(x ; \Theta)\right)-\log \left(f_{\mathbf{X} \mid \mathbf{Y}=y}(x ; \Theta)\right) \tag{15}
\end{equation*}
$$

If we now take the conditional expectation with respect to $X$ given $\mathbf{Y}=y$ for a parameter value $\Theta^{\prime}$ on both sides, we are left with

$$
\begin{align*}
\log \left(f_{\mathbf{Y}}(y ; \Theta)\right)= & E\left[\log \left(f_{\mathbf{X}}(x ; \Theta)\right) \mid \mathbf{Y}=y ; \Theta^{\prime}\right]- \\
& E\left[\log \left(f_{\mathbf{X} \mid \mathbf{Y}=y}(x ; \Theta)\right) \mid \mathbf{Y}=y ; \Theta^{\prime}\right] \tag{16}
\end{align*}
$$

This looks cumbersome, so for convenience we define

$$
\begin{align*}
& L(\Theta)=\log \left(f_{\mathbf{Y}}(y ; \Theta)\right) \\
& U\left(\Theta, \Theta^{\prime}\right)=E\left[\log \left(f_{\mathbf{X}}(x ; \Theta)\right) \mid \mathbf{Y}=y ; \Theta^{\prime}\right]  \tag{17}\\
& V\left(\Theta, \Theta^{\prime}\right)=E\left[\log \left(f_{\mathbf{X} \mid \mathbf{Y}=y}(x ; \Theta)\right) \mid \mathbf{Y}=y ; \Theta^{\prime}\right]
\end{align*}
$$

to get the more elegant expression

$$
\begin{equation*}
L(\Theta)=U\left(\Theta, \Theta^{\prime}\right)-V\left(\Theta, \Theta^{\prime}\right) \tag{18}
\end{equation*}
$$

This is a special way of writing the original loglikelihood that we want to maximize. Using Jensen's inequality on the $V$-term yields

$$
\begin{equation*}
V\left(\Theta, \Theta^{\prime}\right) \leq V\left(\Theta^{\prime}, \Theta^{\prime}\right) \tag{19}
\end{equation*}
$$

We infer that if $U\left(\Theta, \Theta^{\prime}\right)>U\left(\Theta^{\prime}, \Theta^{\prime}\right)$, then

$$
\begin{equation*}
L(\Theta)>L\left(\Theta^{\prime}\right) \tag{20}
\end{equation*}
$$

An EM algorithm starts from an arbitrary initialization point $\hat{\Theta}^{(0)}$ at time $t=0$. Subsequent parameter estimates are computed in the following manner:

- Expectation-step: Compute

$$
\begin{equation*}
U\left(\Theta, \hat{\Theta}^{(t)}\right) \tag{21}
\end{equation*}
$$

- Maximization-step:

$$
\begin{equation*}
\hat{\Theta}^{(t+1)}=\max _{\Theta}\left(U\left(\Theta, \hat{\Theta}^{(t)}\right)\right) . \tag{22}
\end{equation*}
$$

- The iteration is repeated, for instance, until

$$
\begin{equation*}
\left\|\hat{\Theta}^{(T)}-\hat{\Theta}^{(T+1)}\right\| \leq \varepsilon \tag{23}
\end{equation*}
$$

for some iteration step $T$ (other conditions such as a maximum number of iterations are also possible). Smaller epsilons result in more precise final estimates at the cost of more iteration steps needed to converge.

The formulation follows the scheme of a steepest-descent algorithm. Hence, EM algorithms will converge to a local maximum, and in general we do not know whether or not a particular found local maximum coincides with the global maximum we want. The final results are thus heavily dependent on the starting points, but the EM procedure itself does not define those.

### 3.2 Fundamentals of Gaussian Mixture Models (GMM)

Gaussian mixture models are intended to express a more general, multimodal probability density function (i.e., multiple peaks) as a superposition of single Gaussian distributions (called "mixtures"). This is convenient because the ML estimates of the parameters of Gaussians are analytically tractable and well-known. In most cases, the GMM formulation can be used as a good approximation to a real distribution, even if mixture components are not really Gaussian. The mathematical formulation is simply

$$
\begin{equation*}
f_{\mathbf{X}}(x ; \Theta)=\sum_{m=1}^{M} w_{m} \cdot N\left(x ; \Theta_{m}\right), \tag{24}
\end{equation*}
$$

where the $w_{m}$ are mixture-weights that must sum to unity, i.e. $\sum_{m=1}^{M} w_{m}=1$, and $\Theta_{m}=$ $\left[\sigma_{m}^{2}, \mu_{m}\right]$ denote the means and variances, as found in the Gaussian density function

$$
\begin{equation*}
N\left(x ; \sigma^{2}, \mu\right)=\frac{1}{\sqrt{2 \pi \sigma^{2}}} \cdot \exp \left(-\frac{(x-\mu)^{2}}{2 \sigma^{2}}\right) . \tag{25}
\end{equation*}
$$

Further,

$$
\begin{equation*}
\Theta=\bigcup_{m=1}^{M} \Theta_{m} \tag{26}
\end{equation*}
$$

and $M$, the number of mixture components, is assumed to be known for now.
With knowledge of the observable data $x_{i}$, the (unbiased) ML parameter estimates of a single Gaussian are [5]

$$
\begin{align*}
& \hat{\mu}=\frac{1}{N} \cdot \sum_{i=1}^{N} x_{i} \\
& \hat{\sigma}^{2}=\frac{1}{N-1} \cdot \sum_{i=1}^{N}\left(x_{i}-\mu\right)^{2} \tag{27}
\end{align*}
$$

### 3.3 An EM Algorithm for GMM Parameter Estimation

The general EM approach can now be adapted to work on GMM problems. The "missing data" are the means, variances, and weights (or prior member probabilities) of the mixture components, which must be estimated from a limited number of samples.

- The typical expectation step for a GMM [21] yields

$$
\begin{equation*}
\hat{P}\left(m \mid x_{i}\right)^{t}=\frac{\hat{w}_{m}^{t} \cdot N\left(x_{i} ; \hat{\sigma}_{m}^{2 t}, \hat{\mu}_{m}^{t}\right)}{\sum_{k=1}^{M} \hat{w}_{k}^{t} \cdot N\left(x_{i} ; \hat{\sigma}_{k}^{2 t}, \hat{\mu}_{k}^{t}\right)} \tag{28}
\end{equation*}
$$

which is the estimated probability (at iteration step t ) that a sample $x_{i}$ belongs to Gaussian component $m$.

- The maximization part gives

$$
\begin{align*}
& \hat{w}_{m}^{t+1}=\frac{1}{N} \sum_{i=1}^{N} \hat{P}\left(m \mid x_{i}\right)^{t} \\
& \hat{\mu}_{m}^{t+1}=\frac{1}{\sum_{i=1}^{N} \hat{P}\left(m \mid x_{i}\right)^{t}} \cdot \sum_{i=1}^{N} x_{i} \hat{P}\left(m \mid x_{i}\right)^{t}  \tag{29}\\
& \hat{\sigma}_{m}^{2 t+1}=\frac{1}{\sum_{i=1}^{N} \hat{P}\left(m \mid x_{i}\right)^{t}} \cdot \sum_{i=1}^{N}\left(x_{i}-\mu_{m}^{t}\right)^{2} \hat{P}\left(m \mid x_{i}\right)^{t}
\end{align*}
$$

- An instance of a termination condition for the iterative scheme is given above (general EM), and is also used in our implementation.


### 3.4 Missing Data in EM/GMM Analysis of Monophonic Music (Number of Components)

Now that we have an iterative scheme for estimating the mixture parameters, let us take a look back at the pitch estimation problem. The fundamental and the harmonics are spread


Figure 10: Many components arising from low pitch (bass trombone, F\#1)
in approximately equal distance from each other across the magnitude-spectrum. A lower pitch results in these peaks being closer to each other, and higher pitches force them farther apart.

Given the limited frequency interval of typically $0 . .8 k H z$ (for a 16 kHz sampling frequency), we will see a different number of mixture components if we look at different pitches. Figures 10 and 11 demonstrate this. A low pitch has many components (bass trombone $\mathrm{F} \# 1$ in Figure 10), whereas a high pitch has few components (flute B6 in Figure 11).

Naturally, if the assumed number of components $M$ does not coincide with the actual number, our estimation is not optimal, even with the best initial parameter guess. There are a number of reasonably well-explored algorithms to determine $M$ (see [1]), but they are not low in complexity. In testing of the EM at this single pitch stage we experimented with different choices for $M$ and found that for many cases, $15 \leq M \leq 20$ works best (although still far from the desired quality). More discussion about this is given in Chapter 5.1.


Figure 11: Only one component arising from high pitch (flute, B6)

### 3.5 Reinterpretation of the Magnitude Spectrum as a Spectral Density

So far, an EM algorithm has been formulated to work on monophonic cases using a GMM formulation. However, we assume that we have samples $x_{i}$ from the unknown distribution. In our problem of pitch estimation, we are not supplied with these samples. The definition of the (discrete) Fourier Transform does not provide any meaning for some "sample" to form a frequency density, simply because a single sample does not carry an inherent frequency. Indeed, the DFT is deterministic in nature.

The EM presented above can be reformulated to merely work with the input of the density (strictly speaking, it is a probability mass function, or pmf, the discrete equivalent to the probability density) of the samples. In the implementation, this density is the magnitudespectrum for a single frame of the sound file. Note that we are using a probabilistic artifice to solve what is basically a deterministic function fitting problem.

### 3.6 Modification of the EM Algorithm to Accept the "Frequency Density"whose Parameters Are to Be Estimated

Let the "frequency density" input to the EM algorithm be denoted by $\phi\left(\omega_{i}\right)$, where the support points $\omega_{i}$ are the frequencies at which the DFT for a sound frame was computed (it is mandatory that $\sum_{i=1}^{\Omega} \omega_{i}=1$ in order to be a valid density).

- Expectation-step

$$
\begin{equation*}
\Gamma\left(m, \omega_{k}\right)^{t}=\frac{w_{m}^{t} \cdot N\left(\omega_{i} ; \hat{\sigma}_{m}^{2 t}, \hat{\mu}_{m}^{t}\right)}{\sum_{k=1}^{M} w_{k}^{t} \cdot N\left(\omega_{i} ; \hat{\sigma}_{k}^{2 t}, \hat{\mu}_{k}^{t}\right)} . \tag{30}
\end{equation*}
$$

This is the curve of a single Gaussian component at its current parameters, normalized over the sum over all mixtures. Comparing this to the former equation for $P$, this expression gives a likelihood measure for each abscissa value of the density graph, whereas before, probabilities were computed for each such value, even if it occurred multiple times. The "number of occurrences" is accounted for by the multiplication with $\phi\left(\omega_{k}\right)$ in the maximization step, where now real positive values can and will also be assumed.

- Maximization-step:

$$
\begin{align*}
& \hat{w}_{m}^{t+1}=\sum_{k=1}^{\Omega} \phi\left(\omega_{k}\right) \cdot \Gamma\left(m, \omega_{k}\right)^{t}, \\
& \hat{\mu}_{m}^{t+1}=\frac{1}{\hat{w}_{m}^{t+1}} \cdot \sum_{k=1}^{\Omega} \omega_{k} \cdot \phi\left(\omega_{k}\right) \cdot \Gamma\left(m, \omega_{k}\right)^{t},  \tag{31}\\
& \hat{\sigma}_{m}^{2 t+1}=\frac{1}{\hat{w}_{m}^{t+1}} \cdot \sum_{k=1}^{\Omega}\left(\omega_{k}-\hat{\mu}_{m}^{t}\right)^{2} \cdot \phi\left(\omega_{k}\right) \cdot \Gamma\left(m, \omega_{k}\right)^{t} .
\end{align*}
$$

This formulation enables a better understanding of the nature of this EM algorithm: multiplication of each mixture with the observed data will move it towards the next peak and eventually converge, yet depending on the initial estimates, some peaks may be "out of reach," especially if the amplitudes between the peaks are very close to zero. Figure 12 demonstrates this visually.


Figure 12: Illustration of the behavior of the EM algorithm for Gaussian Mixture Models

### 3.7 Missing Data in EM/GMM Analysis of Polyphonic Music (Number of Tones, Number of Components)

Progressing now to the polyphonic case, the formula for GMM introduced above would still be able to describe any multitone mix, yet the membership of each Gaussian component to a particular tone is not expressed. This necessitates a GMM formulation for the polyphonic case:

$$
\begin{equation*}
f_{X}(x ; \Theta)=\sum_{p=1}^{P} \sum_{m=1}^{M(p)} w_{m}^{p} \cdot N\left(x ; \Theta_{m}^{p}\right) . \tag{32}
\end{equation*}
$$

The number of components $M$ is now dependent on the pitch they belong to, and here denoted $M(p) . P$ is the number of overlaid pitches. The parameters $\Theta_{m}^{p}=\left[\sigma_{m}^{2 p}, \mu_{m}^{p}\right]$ are now clearly separated according to pitch membership, and the means for a particular $p$ should be thought of as, at least ideally, equally spaced from each other in the frequency domain. The problem of finding a good guess for $P$ will be addressed in Chapter 4.

## CHAPTER IV

## HEURISTICS FOR INITIAL ESTIMATES

As previously mentioned in Chapter 3.1, the performance of an EM algorithm will heavily depend on its initialization because of its nature as a steepest-descent algorithm. This chapter tackles this problem for the specific domain of pitch estimation in polyphonic music, and also features other heuristics that are necessary for the current system to work well on real data. Chapters 4.1 through 4.3 discuss heuristics for single-pitch cases, and 4.4 leaves everything untouched except for the core heuristic, which is modified to support polyphony.

### 4.1 Limiting the Frequency Range ("Cut-Off")

The optimization described in this section may seem insignificant, but it can have a tremendous effect on the number of iterations needed for the (unconstrained) EM algorithm to converge. The idea is that we can disregard the highest frequency portion of the magnitude spectrum of a frame. In many cases, the significant peaks are located in the lower regions; sometimes even up to $60-70 \%$ of the spectrum does not provide much discrimination information. Hence, reducing the data becomes appealing, as illustrated by the magnitude spectrum of a real piece of music (sampled at 16 kHz ) shown in Figure 13.

The cutoff frequency should be determined for each frame independently. It was experimentally determined to be well chosen by

$$
\begin{equation*}
\omega_{\text {cutoff }}=\arg \max _{\omega_{k}}\left(\left\{\omega_{k} \mid \phi\left(\omega_{k}\right) \geq 0.025 \cdot \max \left(\phi\left(\omega_{k}\right)\right)\right\}\right) \tag{33}
\end{equation*}
$$

in other words, the largest frequency for which the spectral density value is larger than or equal to 0.025 of the global maximum. In Figure 13, only the spectrum inside the red box will be passed on to any further processing. The rest is discarded since it does not carry useful information for our purposes.

Figures 14 and 15 show spectrogram results for the (unconstrained) EM algorithm but with and without the spectral cutoff heuristic (and without further heuristics). This simple


Figure 13: Reducibility of the spectrum


Figure 14: Spectrogram result without spectral cutoff


Figure 15: Spectrogram result with spectral cutoff
heuristic yields a $47 \%$ reduction of EM algorithm iterations for the given spectrogram. The saving will generally depend on musical timbres and pitches contained in the data.

### 4.2 Leveraging the Spacing of the Harmonics

As discussed in the introduction, we know that the locations of the peaks belonging to a particular pitch should adhere closely to frequencies that are integer multiples of the fundamental (1). Assuming we are given a recording with reasonable SNR, we can also assume that the peaks we seek generally have significantly higher amplitudes than any noise-related peaks. Supposing we can find a set of such peaks that follow these principles with some tolerance, we can use this set as a good initial estimate for the EM algorithm.

### 4.2.1 Peak-Picking Function

A code module that finds all local maxima of a curve had to be added (low complexity). This module finds all the local maxima, no matter where they are located, and no matter how


Figure 16: Summary of cutoff ranges
flat the peaks might be. Further filtering of these peaks had to be applied; the discussion follows.

### 4.2.2 Peak Filtering Heuristic

Examination of a lot of data has shown that there may be a low-frequency peak with a high amplitude that does not fit into the harmonic structure. Although it is not entirely clear where such peaks come from, we suspect it may be a recording issue. It is necessary to exclude them ( $f_{k} \leq 45 \mathrm{~Hz}$ ). Furthermore, we exclude extremely low amplitude peaks from consideration as components $\left(\right.$ those $\left.\leq 0.05 \cdot \max \left(\phi\left(\omega_{k}\right)\right)\right)$. This results in the following peak exclusion window (after peak-picking):

This means that the set of peaks will be initialized according to:

$$
\begin{align*}
& \Lambda:=\max _{\text {local }}\left(\phi\left(\omega_{k}\right)\right) ;  \tag{34}\\
& \Lambda:=\Lambda \backslash\left\{\lambda_{k} \mid \lambda_{k}<45 H z \vee \phi\left(\lambda_{k}\right)<0.05 \cdot \max (\Lambda)\right\} ;
\end{align*}
$$

### 4.2.3 Core Heuristic: Basic Idea

Here, the core heuristic will be explained for the single pitch case, and in Chapter 4.4 extended to support polyphony.

Suppose the peak-picking function returns a set of frequency locations for the peaks,
denoted by $\Lambda=\left\{\lambda_{k}\right\}$, where $k=1 \ldots K$, and $K$ is the total number of peaks found, which is not predictable because of noise. The basic idea of the heuristic described in this section can be expressed in pseudocode as (let the initial guesses we are seeking be denoted by $\hat{\mu}_{m}^{0}$ ):

$$
\begin{align*}
& m=1: \hat{\mu}_{1}^{0}:=\lambda_{1} ; \\
& \text { for }\left(m:=2 ; m:=m+1 ; m<\left\lfloor\frac{\lambda_{K}}{\hat{\mu}_{1}^{0}}\right\rfloor\right) \\
& \text { if } \exists \lambda_{k} \in \Lambda:\left|m \hat{\mu}_{1}^{0}-\lambda_{k}\right| \leq \delta_{f} \\
& \text { then } \\
& \qquad \hat{\mu}_{m}^{0}:=\arg \min _{\lambda_{k}}\left|m \hat{\mu}_{1}^{0}-\lambda_{k}\right|  \tag{35}\\
& \text { else } \\
& \qquad \hat{\mu}_{m}^{0}:=\infty \\
& \text { endif } \\
& \text { endfor }
\end{align*}
$$

In other words, we assume the peak at the lowest frequency to be the fundamental. Then we proceed to look for integer multiples of it (within some tolerance) across the entire frequency range. The tolerance with which the peaks are allowed to deviate from their ideal locations $(+$ or -$)$ is defined as $\delta_{f}$. This value has to be determined experimentally. We also have to consider the pitch range we want to investigate. The difficulty is that for low pitches, we would want $\delta_{f}$ to be low because otherwise we might consider a tolerance corridor that extends over adjacent peaks. For higher pitches, peaks are far apart, and then it would make sense to increase $\delta_{f}$, to allow for a little more variation. From experience, we currently recommend using

$$
\begin{equation*}
\delta_{f}=\min \left(50 H z, \hat{\mu}_{1}^{0} / 4\right) \tag{36}
\end{equation*}
$$

which prevents the tolerance from increasing unboundedly. This yields reasonable values for a pitch range from the second to the ninth octave (approx. $65-15800 \mathrm{~Hz}$ ). The lowest pitches are hard to estimate, because the fluctuation of the peak locations over the frequency range of the many peaks is quite significant. Chapter 5 discusses this problem further.

### 4.2.4 Interpolation for "Missing Peaks"

Suppose we have a low-amplitude harmonic in the mid-frequency range that is so weak it would be excluded from the set of peaks $\Lambda$ after the peak-filtering heuristic discussed in 4.2.2. Instead of only recognizing the first consecutive set of found harmonics, counting from lowest frequency on, and discarding the rest, it makes sense to fill any "holes" or "missing peaks" of "not found" harmonics up to the highest valid found harmonic.

For example, if we have a set of found harmonics
$50100150200 \operatorname{Inf} 300350 \operatorname{Inf} \operatorname{Inf}$,
where the Inf entries represent harmonics that are not found in the set of filtered peaks, we interpolate the Inf entry between 200 and 300, but the two others are deemed uncertain and/or imprecise. However, this is not done by filling in the mean of the two adjacent values, instead, the heuristic looks at the set of peaks that was returned before peak-filtering occurred. If there is nothing found within that frequency range, only the entries of 50 to 200 are considered to belong to this pitch (this more significance for the polyphonic case treated later). This modification is not given here analytically because it is straightforward. It can be followed in the code.

### 4.3 Dynamic Determination of Number of Mixtures

If the peak-picking heuristic works well, the problem of the user having to enter the correct number of mixtures is also solved. It is clear that the above formulation (Chapter 4.2.3) can terminate at any step, and hopefully, we will receive an $M$, equaling the number of peaks found, that will be equal to the true number of components for that pitch.

### 4.4 Extension to Polyphonic Cases

### 4.4.1 Maximum Number of Polyphonic Tones

In the implementation of the GUI, a field was added that lets the user specify the maximum number of simultaneous tones played in the data (the user/developer should ideally know this for testing purposes). This value serves as upper termination condition for the polyphonic-capable heuristic developed next.

### 4.4.2 Iterative Scheme

Suppose we are given $\Lambda$, the set of peaks, as outlined earlier in this chapter. For notational simplicity, we drop the superscript zero indicating the initial estimate, since this should be understood by now. We define $\hat{\mu}_{m}^{p}$ as the initial estimate of the $m^{\text {th }}$ mixture mean of the $p^{\text {th }}$ pitch, and $P$ as the user-defined maximum of simultaneous pitches available in the signal. The extension to the polyphonic case is realized with this iterative scheme:

```
\(p:=1 ;\)
while \((p<P \&|\Lambda|=K>0)\)
    \(m=1: \hat{\mu}_{1}^{p}:=\lambda_{1} ;\)
    for ( \(\left.m:=2 ; m:=m+1 ; m<\left\lfloor\frac{\lambda_{K}}{\hat{\mu}_{1}^{\Gamma}}\right\rfloor\right)\)
        if \(\exists \lambda_{k} \in \Lambda:\left|m \hat{\mu}_{1}^{p}-\lambda_{k}\right| \leq \delta_{f}\)
```

            then
                \(\hat{\mu}_{m}^{p}:=\arg \min _{\lambda_{k}}\left|m \hat{\mu}_{1}^{p}-\lambda_{k}\right| ;\)
                \(\Lambda:=\Lambda \backslash \hat{\mu}_{m}^{p} ;\)
                else
                \(\hat{\mu}_{m}^{p}:=\infty\)
            endif
        endfor
        \(p:=p+1 ;\)
    endwhile
It is thus possible for the algorithm to terminate before $p=P$ has been reached, and this will be the case if all significant peaks have been used up in prior set(s) of peaks. This is useful because $P$ is entered as a global parameter for a piece of music, in which the number of simultaneous tones can vary, and the above iterative scheme merely works on a single frame of typically 64 ms .

All non-initialized means for a certain value of $p$ and its successors can be subsequently regarded as a pitch (or pitches) which is (or are) not available at this time frame. Consequently, no EM algorithm will be executed for those means.

### 4.4.3 Constrained EM Algorithm

In this subsection, the EM algorithm working on GMM will be modified in such a manner as to only allow component means that adhere to a harmonic structure. The motivation for this is that we expect both the initial estimates by the heuristic and the final estimates of the EM algorithm to be in that range. Constraining the EM algorithm also saves us precious computation time, since the constraint should cut down on the number of iterations necessary for convergence.

Let the initial estimates, the outputs of the heuristic, be denoted by $\hat{\mu}_{m}^{p}(0)$ and $\hat{\sigma}_{m}^{p}(0)$. These are the estimates of the means and variances for each pitch $p=1 . . P$ and each component $m=1 . . M(p)$ as previously defined (see Chapter 3.7 on polyphonic GMM), at $t=0$ (init). While there is only a loosely enforced harmonic spacing condition in these initial estimates, the subsequent EM algorithm iterations will place strict requirements on the component means.

- Auxiliary EM Initialization: $\hat{\mu}_{m}^{p}(0)=m \cdot \hat{\mu}_{1}^{p}(0)$, or any other pitch computation scheme as discussed in Section 5.1 (here, we use the "fundamental only" scheme)
- Expectation-step:

$$
\begin{equation*}
\Gamma\left(p, m, \omega_{k}\right)^{t}=\frac{w_{m}^{p}(t) \cdot N\left(\omega_{i} ; \hat{\sigma}_{m}^{2 p}(t), \hat{\mu}_{m}^{p}(t)\right)}{\sum_{k=1}^{M} w_{k}^{p}(t) \cdot N\left(\omega_{i} ; \hat{\sigma}_{k}^{2 p}(t), \hat{\mu}_{k}^{p}(t)\right)} . \tag{38}
\end{equation*}
$$

- Maximization-step:

$$
\begin{align*}
& \hat{w}_{m}^{p}(t+1)=\sum_{k=1}^{\Omega} \phi\left(\omega_{k}\right) \cdot \Gamma\left(m, \omega_{k}\right)^{t}, \\
& \hat{\mu}_{m}^{p}(t+1)=\frac{1}{\hat{w}_{m}^{p}(t+1)} \cdot \sum_{k=1}^{\Omega} \omega_{k} \cdot \phi\left(\omega_{k}\right) \cdot \Gamma\left(p, m, \omega_{k}\right)^{t},  \tag{39}\\
& \hat{\sigma}_{m}^{2 p}(t+1)=\frac{1}{\hat{w}_{m}^{p}(t+1)} \cdot \sum_{k=1}^{\Omega}\left(\omega_{k}-\hat{\mu}_{m}^{p}(t)\right)^{2} \cdot \phi\left(\omega_{k}\right) \cdot \Gamma\left(p, m, \omega_{k}\right)^{t} .
\end{align*}
$$

- Enforcing strict harmonic spacing for next iteration:

$$
\begin{align*}
& \Delta_{f}(t+1)=\frac{1}{M(p)} \cdot \sum_{m=1}^{M(p)} \frac{1}{m}\left(\hat{\mu}_{m}^{p}(t+1)-\hat{\mu}_{m}^{p}(t)\right),  \tag{40}\\
& \hat{\mu}_{m}^{p}(t+1)=m \cdot\left(\hat{\mu}_{1}^{p}(t)+\Delta_{f}(t+1)\right) .
\end{align*}
$$

This deserves some explanation. In the previous two steps, the same unconstrained EM formulation given before has been used. The parameters can flow freely. After
that, we look at how the means changed in comparison to the previous means. The formula for $\Delta_{f}$ is the mean of these changes. It will be positive if the sum of all components moving up in frequency is larger than those moving downward, thus following the overall dominant trend. This will be accounted for in a correction to the component means, which are in this step strictly located according to the harmonics condition.

- Modified termination condition. Iterate until

$$
\begin{align*}
& \left|\Delta_{f}\right| \leq 0.01 \text { or } \\
& \left|\Delta_{f}(t)+\Delta_{f}(t+1)\right|=0 \text { or }  \tag{41}\\
& \hat{\mu}_{m}^{p}(t+1) \geq \hat{\mu}_{m}^{p}(0) \cdot 2^{1 / 24} \text { or } \\
& \hat{\mu}_{m}^{p}(t+1) \leq \hat{\mu}_{m}^{p}(0) \cdot 2^{-1 / 24} .
\end{align*}
$$

The first line is defined in analogy with the previously used termination condition. The second line is necessary because, in some cases, subsequent $\Delta_{f}$ have the same magnitude but alternate in sign, and this is regardless of the first condition. Conditions three and four are simple sanity-checks for the moving of the first mixtures we do not want them to move more than half a semitone (see the introduction for details).

### 4.5 Chapter Review

This chapter discussed how initial estimates for polyphonic music can be found using straightforward heuristics, and how the EM algorithm can be constrained and terminated using knowledge of the application. At this stage, the EM algorithm is executed serially for each harmonic structure reflecting one pitch, assuming at each time that it were the only pitch present. Some problems with this assumption will be discussed in Chapter 5.2.

## CHAPTER V

## PRELIMINARY RESULTS AND ANALYSIS

### 5.1 Implementation Remarks: Single-Pitch Prototype

Figure 17 shows a screenshot of a MATLAB-based visualizer that plots the spectrogram and computes pitch estimates dependent on some user input parameters. The input parameters include:

- Frame length (defaulting $64 m s$ )
- Number of components (typically 15 to 20 ), constant over all frames
- Method selector: how to determine pitch value from the sets of mixture mean estimates

The last bullet has not been discussed yet; and several different approaches may come to mind. We implemented the following:

- Fundamental only: $p=\hat{\mu}_{1}^{T}$ ( $T$ is not a matrix transpose, it denotes the final value of the estimation process)
- Second harmonic minus fundamental: $p=\hat{\mu}_{2}^{T}-\hat{\mu}_{1}^{T}($ if $M>1)$
- Mean: $p=\frac{1}{M} \sum_{m=1}^{M} \frac{\hat{\mu}_{m}^{T}}{m}($ if $M>1)$
- "Sum and Divide": $p=\frac{1}{M(M+1) / 2} \sum_{m=1}^{M} \mu_{m}^{T}($ if $M>1)$

The idea for the latter is that if, in the ideal case, $\hat{\mu}_{m}^{T}=m \hat{\mu}_{1}^{T}$ for all $m$, then

$$
\begin{equation*}
\sum_{m=1}^{M} \hat{\mu}_{m}^{T}=\sum_{m=1}^{M} m \hat{\mu}_{1}^{T}=\hat{\mu}_{1}^{T} \sum_{m=1}^{M} m=\hat{\mu}_{1}^{T} \frac{M(M+1)}{2} \tag{42}
\end{equation*}
$$

where $\hat{\mu}_{1}^{T}$ is the pitch coinciding with the fundamental that we would like to discover. Other parameters include:

- Frame overlap set to $50 \%$


Figure 17: GUI screenshot of single-pitch prototype

- Hamming-windowing of the time frames
- Sample rate usually 16 kHz (but works with all of course)

The EM algorithm is initialized by the following procedure:

- Spread mixture means ( $\mu$ 's) equally across the entire frequency range
- Set all variances equal to a fraction of the global variance
- Initialize all mixture weights with $1 / M$ (note the weights sum to unity)

This GUI system running in MATLAB was used in all the later parts of the pitch estimation project for visual verification purposes. During the course of the research, it underwent some changes that will be discussed later.

The graph in Figure 17 should be interpreted as follows. The spectrogram shows the development of the signal's spectral magnitude over time (the time axis is scrollable for
better resolution of long files). The harmonics are apparent in the graph as thin red horizontal lines, with red indicating larger amplitude. These are the peaks of each frame. Blue indicates low amplitude. Due to the $1 / f$ falloff typical of acoustic instruments, higher frequency regions tend towards blue.

The white lines show where the EM estimated the harmonics to be. Ideally, we would want them to lay exactly on top of the red lines of the data. Finally, the only black line in the graph shows the actual pitch estimate. It should ideally be close (but not necessarily identical) to the fundamental, which is the lowest of the red lines.

### 5.1.1 Results for Single Pitch Cases

We tested with labeled data [19], so we know what the true numerical pitch values of the recorded notes should be. We discovered that for many of the samples, the pitch was not identical to the scientific pitch notation [33]. We verified the operation of the EM algorithm by looking at how well the means matched the peaks of the given data. After all, if we have an A4 recording and the first peak is at 449 Hz and the second at 893 Hz , instead of 440 Hz and 880 Hz , for instance, it does not make sense to try to adapt an implementation to the latter. Our observations do not mean that the basic truths (labeling) are wrong ${ }^{1}$, just that the perceived pitch is a bit more involved, as laid out earlier in the introduction. Hence, we currently have no sophisticated numerical statistics summarizing our tests.

### 5.1.2 Potential Difficulties in Adopting GMM For Pitch Estimation

- The Gaussian function will inevitably return numerical values so small that they are considered zero. This happens in frames whose peaks are widely spread across the whole spectral range. The normalization of the E-step would divide by zero, which clearly has to be prevented. This might seem unlikely to happen, but we found it occurred quite frequently.
- Peaks are sometimes extremely narrow, so variances that become numerically close to zero are possible. Additional precautions have to be taken when implementing to

[^0]prevent this from causing trouble. In particular, the Gaussian function, which takes in the frequency range, as well as the parameter estimates of mean and variance, has to be modified also to work with zero variance inputs if we do not prohibit this by defaulting to a minimum nonzero variance. In such cases, we returned a vector that evaluates to zero everywhere but at the specified mean, at which some arbitrary large value is chosen (we chose 1000).

- Some noise peaks may confuse the EM algorithm. Smoothing may resolve this issue and make the peaks more similar to Gaussians. We tried a little bit of running-average smoothing with a window of two and experimented with other values as well. If we smooth using this technique, the window should not be chosen large, because this could merge larger peaks into one. This is generally not desired, unless the heuristics explicitly account for such cases.


### 5.2 Discussion of the Final Implementation

The main limitations of the system (see Figure 18 for a screenshot) with all the described heuristics (GUI screenshot in Figure ) for polyphonic operation is that it is dependent on a user-specified maximum number of tones $P$. If this maximum is correct, we enjoy relatively robust number of tone detection. However, for certain cases, the algorithm will find excess pitches not present in the data if $P$ is chosen too high.

Figure 19 shows pitch estimation results superimposed over the spectrogram of a highpitched triad (C diminished in octave 6, ideally having pitches of $1047,1319,1480 H z$ ). The results are quite good. We can see that at about $t=2 s$, the C 6 pitch fades out, one of the pitch lines discontinues (blue), and the two remaining lines take over the remaining two pitches. Also, the precision of the pitch estimates seems to visually follow the spectrogram lines of the fundamentals. Note that the rapid decline of the last remaining pitch at the end is due to offset artifacts.

The verbose MATLAB console printout for the experiment shown in Figure 19 looks like:
runEM: frame 22 of 73


Figure 18: GUI screenshot of the final implementation


Figure 19: Heuristic working well in a high-pitch example

```
myEM (tone 1): init(heur.) means: 1062.5 2125
myEM (tone 1): init(pitch) means: 1062.5 2125
myEM (tone 1): after EM: 1063.8381 2127.6762
myEM (tone 2): init(heur.) means: 1343.75 2687.5
myEM (tone 2): init(pitch) means: 1343.75 2687.5
myEM (tone 2): after EM: 1341.3972 2682.7945
myEM (tone 3): init(heur.) means: 1500 3000
myEM (tone 3): init(pitch) means: 1500 3000
myEM (tone 3): after EM: 1500.8629 3001.7259
```

In other octaves, more problems occur. In Figure 20, an undesired effect occurs at $t=1.8 s$, where the estimated number of pitches does not change, but in reality, a tone fades out at this point. The blue (top) line is not desired here. There are several reasons for this. The fundamentals of the tones E4 and G4 lie close together, and in this case appear


Figure 20: Excess pitches in medium-pitch example
to be wider than usual. Furthermore, there is some apparent imprecision in the recordings (pitches should be more distant from each other, $\mathrm{E} 4=329.6 \mathrm{~Hz}, \mathrm{G} 4=392 \mathrm{~Hz}$ ), and the moving average process also contributes to the merging of the two fundamentals into one. Thus, the peak-picking heuristic will come up with only one value here. The following heuristic picks one fundamental and then looks for integer multiples of it, however, without assuming that certain peaks can be "shared." It then follows that there will be a higher spurious peak detected within the range of the second harmonics, which of course separate out better. In reality this is not an actual new fundamental of another pitch, but just the harmonic of one of the pitches that unfortunately merged together as just described.

Figure 21 shows the behavior of the algorithm when the spectral magnitude is not smoothed prior to executing the heuristics and the EM algorithm. It becomes clear that the E4 and G4 fundamentals now separate better, considering all the places where the blue and white pitch lines lie close together, as they should.

Another effect that potentially occurs when two harmonics of two different pitches lie


Figure 21: Excess pitches in a medium-pitch example - without spectral smoothing
close together is that the EM algorithm, in its current implemented form of being serially executed for each pitch, will tend to try combining the two harmonics into a larger one, thus shifting the initial estimate towards the middle of the two adjacent peaks. Precautions have been taken, however, to limit this behavior (namely small component variance and a maximum deviation from the initial pitch estimate as one of the termination conditions). This effect may be leveled out by the opposite tendencies of higher harmonics. In our tests, we generally did not find this issue to cause much concern.

The third issue we mention is that despite allowing for some variation from the ideal harmonic spacing, the highest harmonics of a complex tone (i.e. one with a high number of harmonics) may escape the heuristic and thus lead to spurious high pitch estimates that do not exist in the data. This is demonstrated in Figure 22.


Figure 22: Excess pitches in medium-pitch example - escaping the heuristic

### 5.2.1 Summary of Issues

Clearly, this approach is highly dependent on the quality of the heuristics. The EM algorithm is a great mathematical tool for making precise estimates of the needed parameters, but its nature as a steepest descent algorithm demands that great effort go into finding good initialization. Our heuristic techniques are not claimed to be highly sophisticated, but they should encourage further investigations into this approach. In particular, the assumption that one peak always reflects only a single harmonic is too simplistic and quickly leads to problems when the degree of polyphony increases.

Some concepts for that may be useful in this application can be found in [34].
The main issues found mey be summarized as follows:

- We need to set a small constant variance for the Gaussian mixtures, so that the EM algorithm will not gravitate too much toward an adjacent peak that belongs to a different pitch.
- Limited peak resolution, the smearing of two peaks close to each other but belonging to different pitches, is aggravated by spectral smoothing.
- The number of pitches is not estimated very reliably in all octaves.
- This approach is good if peaks have some minimum frequency distance between each other.
- If there are many harmonic belonging to a pitch, the highest harmonics seem prone to escape the heuristic, and spurious high pitches are detected.


## CHAPTER VI

## CONCLUSIONS AND DIRECTIONS FOR FUTURE RESEARCH

### 6.1 Conclusions

The final system we developed allows for good estimates of pitches in the middle to high octaves. It uses a low number of EM steps, especially compared to our earlier unconstrained version without the heuristics. The total number of iterations (for a piece of music of length $2-3 s)$ has been reduced from the order of 10,000 to the order of 100 , while also usually coming up with better estimates. The implementation requires entering the maximum number of pitches $(P)$ played simultaneously, which is a much more practical parameter than the number of components found in an early version. Many instruments, such as winds and bowed strings (excluding special cases), are only capable of producing a single pitch at a time, so for some ensembles in which the instrumentation should be known, this could already be a usable application. For other instruments such as the piano and the harp, $P$ might not be obvious at all.

### 6.2 Directions for Future Work

### 6.2.1 Account for Harmonic Overlapping of Pitches

The possibility of one peak being "shared" by two different harmonic structures should be incorporated into our heuristics. Due to the harmonic composition principles of western music, the "sharing" of harmonics is quite likely to happen.

### 6.2.2 Follow Klapuri Closely

The paper by Klapuri [17] proposes that after a decision about the membership of a set of mixtures to a specific pitch has been made, the spectrum should be recomputed. This computation simply takes the previous spectrum and subtracts the set of estimated mixtures
belonging to the current pitch. The next steps would then involve re-picking the peaks, running the heuristic, and subtracting again, until the maximum number of pitches is reached.

Klapuri's proposal makes sense, because neighboring peaks have a mutual influence on their weights: the closer they are together, the higher they will appear, when in reality, they are lower. This is not important for pitch estimation itself because we are only interested in the mixture means, but suppose we were interested in identifying what instruments played the pitches; then, the weights are the crucial parameters (and to some minor extent, maybe also the variances). Some problems discussed in the previous chapter could also be resolved.

### 6.2.3 Confidence Score for Missing Harmonic Interpolation

The decision of what to do with "missing harmonics" has to consider two notions:

- The harmonics found after the gap could also be the first harmonics of other pitches,
- The pitch's harmonics may actually have ended at the missing harmonic (such timbres exist; we only need to mix the "right" instruments)

To better account for this, some kind of "confidence score" could allow us to decide whether or not to interpolate the missing harmonic.

### 6.2.4 Tune Some of the EM Numerical Parameters such as Termination

More testing with the termination condition described in Section 4.4.3 and possibly other termination conditions, should be done.

### 6.2.5 Set Up an Instrument Identifier on the Separated Sources

If we receive a good separation of pitches, we can use our work to identify different instruments. Some further literature review would need to be done prior to this.

## APPENDIX A

## PITCH TABLE

| Table of note frequencies - Key: Frequency in hertz (semitones from middle C |  |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Octave Note $\downarrow$ | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| C | $\begin{aligned} & \hline 16.35 \\ & (-48) \\ & \hline \end{aligned}$ | $\begin{aligned} & \hline 32.70 \\ & (-36) \\ & \hline \end{aligned}$ | $\begin{aligned} & \hline 65.41 \\ & (-24) \\ & \hline \end{aligned}$ | $\begin{aligned} & 130.8 \\ & (-12) \end{aligned}$ | $\begin{aligned} & \hline 261.6 \\ & (0) \\ & \hline \end{aligned}$ | $\begin{aligned} & 523.3 \\ & (+12) \end{aligned}$ | $\begin{aligned} & \hline 1047 \\ & (+24) \end{aligned}$ | $\begin{aligned} & \hline 2093 \\ & (+36) \end{aligned}$ | $\begin{aligned} & \hline 4186 \\ & (+48) \end{aligned}$ | $\begin{aligned} & 8372 \\ & (+60) \end{aligned}$ |
| C\# | $\begin{aligned} & 17.32 \\ & (-47) \\ & \hline \end{aligned}$ | $\begin{aligned} & 34.65 \\ & (-35) \\ & \hline \end{aligned}$ | $\begin{aligned} & 69.30 \\ & (-23) \\ & \hline \end{aligned}$ | $\begin{aligned} & 138.6 \\ & (-11) \\ & \hline \end{aligned}$ | $\begin{gathered} 277.2 \\ (+1) \\ \hline \end{gathered}$ | $\begin{array}{\|l} 554.4 \\ (+13) \\ \hline \end{array}$ | $\begin{aligned} & 1109 \\ & (+25) \\ & \hline \end{aligned}$ | $\begin{aligned} & 2217 \\ & (+37) \\ & \hline \end{aligned}$ | $\begin{aligned} & 4435 \\ & (+49) \\ & \hline \end{aligned}$ | $\begin{aligned} & 8870 \\ & (+61) \\ & \hline \end{aligned}$ |
| D | $\begin{aligned} & 18.35 \\ & (-46) \\ & \hline \end{aligned}$ | $\begin{aligned} & \hline 36.71 \\ & (-34) \\ & \hline \end{aligned}$ | $\begin{aligned} & 73.42 \\ & (-22) \end{aligned}$ | $\begin{aligned} & \hline 146.8 \\ & (-10) \\ & \hline \end{aligned}$ | $\begin{gathered} 293.7 \\ (+2) \end{gathered}$ | $\begin{aligned} & 587.3 \\ & (+14) \end{aligned}$ | $\begin{aligned} & \hline 1175 \\ & (+26) \end{aligned}$ | $\begin{aligned} & \hline 2349 \\ & (+38) \end{aligned}$ | $\begin{aligned} & \hline 4699 \\ & (+50) \end{aligned}$ | $\begin{aligned} & 9397 \\ & (+62) \end{aligned}$ |
| D\# | $\begin{aligned} & 19.45 \\ & (-45) \\ & \hline \end{aligned}$ | $\begin{aligned} & 38.89 \\ & (-33) \\ & \hline \end{aligned}$ | $\begin{aligned} & 77.78 \\ & (-21) \\ & \hline \end{aligned}$ | $\begin{aligned} & 155.6 \\ & (-9) \\ & \hline \end{aligned}$ | $\begin{aligned} & 311.1 \\ & (+3) \\ & \hline \end{aligned}$ | $\begin{array}{\|l} \hline 622.3 \\ (+15) \\ \hline \end{array}$ | $\begin{aligned} & 1245 \\ & (+27) \\ & \hline \end{aligned}$ | $\begin{aligned} & 2489 \\ & (+39) \\ & \hline \end{aligned}$ | $\begin{aligned} & 4978 \\ & (+51) \\ & \hline \end{aligned}$ | $\begin{aligned} & 9956 \\ & (+63) \\ & \hline \end{aligned}$ |
| E | $\begin{aligned} & \hline 20.60 \\ & (-44) \end{aligned}$ | $\begin{aligned} & \hline 41.20 \\ & (-32) \\ & \hline \end{aligned}$ | $\begin{aligned} & 82.41 \\ & (-20) \end{aligned}$ | $\begin{aligned} & 164.8 \\ & (-8) \end{aligned}$ | $\begin{aligned} & 329.6 \\ & (+4) \end{aligned}$ | $\begin{aligned} & 659.3 \\ & (+16) \end{aligned}$ | $\begin{aligned} & \hline 1319 \\ & (+28) \end{aligned}$ | $\begin{aligned} & 2637 \\ & (+40) \end{aligned}$ | $\begin{aligned} & 5274 \\ & (+52) \end{aligned}$ | $\begin{aligned} & 10548 \\ & (+64) \end{aligned}$ |
| F | $\begin{aligned} & 21.83 \\ & (-43) \\ & \hline \end{aligned}$ | $\begin{aligned} & 43.65 \\ & (-31) \end{aligned}$ | $\begin{aligned} & 87.31 \\ & (-19) \\ & \hline \end{aligned}$ | $\begin{aligned} & 174.6 \\ & (-7) \\ & \hline \end{aligned}$ | $\begin{aligned} & 349.2 \\ & (+5) \\ & \hline \end{aligned}$ | $\begin{aligned} & 698.5 \\ & (+17) \end{aligned}$ | $\begin{aligned} & \hline 1397 \\ & (+29) \\ & \hline \end{aligned}$ | $\begin{aligned} & 2794 \\ & (+41) \\ & \hline \end{aligned}$ | $\begin{aligned} & 5588 \\ & (+53) \\ & \hline \end{aligned}$ | $\begin{aligned} & 11175 \\ & (+65) \\ & \hline \end{aligned}$ |
| F\# | $\begin{aligned} & 23.12 \\ & (-42) \end{aligned}$ | $\begin{aligned} & 46.25 \\ & (-30) \end{aligned}$ | $\begin{aligned} & 92.50 \\ & (-18) \end{aligned}$ | $\begin{aligned} & 185.0 \\ & (-6) \end{aligned}$ | $\begin{aligned} & 370.0 \\ & (+6) \end{aligned}$ | $\begin{aligned} & 740.0 \\ & (+18) \end{aligned}$ | $\begin{aligned} & \hline 1480 \\ & (+30) \end{aligned}$ | $\begin{aligned} & 2960 \\ & (+42) \end{aligned}$ | $\begin{aligned} & 5920 \\ & (+54) \end{aligned}$ | $\begin{aligned} & 11840 \\ & (+66) \end{aligned}$ |
| G | $\begin{aligned} & 24.50 \\ & (-41) \\ & \hline \end{aligned}$ | $\begin{aligned} & 49.00 \\ & (-29) \end{aligned}$ | $\begin{aligned} & 98.00 \\ & (-17) \end{aligned}$ | $\begin{aligned} & 196.0 \\ & (-5) \end{aligned}$ | $\begin{aligned} & 392.0 \\ & (+7) \end{aligned}$ | $\begin{aligned} & 784.0 \\ & (+19) \\ & \hline \end{aligned}$ | $\begin{aligned} & 1568 \\ & (+31) \end{aligned}$ | $\begin{aligned} & 3136 \\ & (+43) \end{aligned}$ | $\begin{aligned} & 6272 \\ & (+55) \\ & \hline \end{aligned}$ | $\begin{aligned} & 12544 \\ & (+67) \end{aligned}$ |
| G\# | $\begin{aligned} & 25.96 \\ & (-40) \end{aligned}$ | $\begin{aligned} & 51.91 \\ & (-28) \end{aligned}$ | $\begin{aligned} & 103.80 \\ & (-16) \end{aligned}$ | $\begin{aligned} & 207.7 \\ & (-4) \end{aligned}$ | $\begin{aligned} & 415.3 \\ & (+8) \end{aligned}$ | $\begin{aligned} & 830.6 \\ & (+20) \end{aligned}$ | $\begin{aligned} & 1661 \\ & (+32) \end{aligned}$ | $\begin{aligned} & 3322 \\ & (+44) \end{aligned}$ | $\begin{aligned} & 6645 \\ & (+56) \end{aligned}$ | $\begin{aligned} & 13290 \\ & (+68) \end{aligned}$ |
| A | $\begin{aligned} & 27.50 \\ & (-39) \end{aligned}$ | $\begin{aligned} & 55.00 \\ & (-27) \end{aligned}$ | $\begin{aligned} & 110.00 \\ & (-15) \end{aligned}$ | $\begin{aligned} & 220.0 \\ & (-3) \\ & \hline \end{aligned}$ | $\begin{aligned} & 440.0 \\ & (+9) \end{aligned}$ | $\begin{aligned} & 880.0 \\ & (+21) \end{aligned}$ | $\begin{aligned} & 1760 \\ & (+33) \\ & \hline \end{aligned}$ | $\begin{aligned} & 3520 \\ & (+45) \end{aligned}$ | $\begin{aligned} & 7040 \\ & (+57) \end{aligned}$ | $\begin{aligned} & 14080 \\ & (+69) \end{aligned}$ |
| A\# | $\begin{aligned} & 29.14 \\ & (-38) \end{aligned}$ | $\begin{aligned} & 58.27 \\ & (-26) \end{aligned}$ | $\begin{aligned} & 116.50 \\ & (-14) \end{aligned}$ | $\begin{aligned} & 233.1 \\ & (-2) \end{aligned}$ | $\begin{aligned} & 466.2 \\ & (+10) \end{aligned}$ | $\begin{aligned} & 932.3 \\ & (+22) \end{aligned}$ | $\begin{aligned} & 1865 \\ & (+34) \end{aligned}$ | $\begin{aligned} & 3729 \\ & (+46) \end{aligned}$ | $\begin{aligned} & 7459 \\ & (+58) \end{aligned}$ | $\begin{aligned} & 14917 \\ & (+70) \end{aligned}$ |
| B | $\begin{aligned} & 30.87 \\ & (-37) \end{aligned}$ | $\begin{aligned} & 61.74 \\ & (-25) \\ & \hline \end{aligned}$ | $\begin{aligned} & 123.50 \\ & (-13) \end{aligned}$ | $\begin{aligned} & 246.9 \\ & (-1) \end{aligned}$ | $\begin{aligned} & 493.9 \\ & (-1) \end{aligned}$ | $\begin{aligned} & 987.8 \\ & (+23) \end{aligned}$ | $\begin{aligned} & 1976 \\ & (+35) \end{aligned}$ | $\begin{aligned} & 3951 \\ & (+47) \end{aligned}$ | $\begin{aligned} & 7902 \\ & (+59) \end{aligned}$ | $\begin{aligned} & 15804 \\ & (+71) \end{aligned}$ |

## APPENDIX B

## CODE LISTINGS

## B. 1 GUI Code

```
function varargout= pitchVisGUI(varargin)
\% varargout= pitchVisGUI(varargin)
\% GUI for pitch estimation and replay of sound files
\% Author: Dominik Loeffler. (c) 2005/2006
gui_Singleton= 1 ;
gui_State \(=\) struct ('gui_Name', mfilename, ...
    'gui_Singleton', gui_Singleton, ...
    'gui_OpeningFcn', @pitchVisGUI_OpeningFcn, ...
    'gui_OutputFcn', @pitchVisGUI_OutputFcn, ...
    'gui_LayoutFcn', [] , ...
    'gui_Callback', []);
    if nargin \& ischar (varargin \(\{1\}\) )
    gui_State.gui_Callback \(=\operatorname{str} 2 f u n c(v a r a r g i n\{1\}) ;\)
    end
    if nargout
    [varargout \(\{1\) : nargout \(\}]=\) gui_mainfcn(gui_State, \(\operatorname{varargin\{ :\} );~}\)
    else
    gui_mainfcn(gui_State, varargin \(\{:\})\);
end
21 \% -- Executes just before pitchVisGUI is made visible.
```

22 function pitchVisGUI_OpeningFcn(hObject, eventdata, handles, varargin)

23 handles.pitchPlotHandle= [];
24 handles.sliderHandle= [];
25 handles.EMOut= [];
26 handles.fileLocCurrentPlot $={ }^{\prime}{ }^{\prime}$;
27 \% Update handles structure
28 handles.output $=$ hObject;
$29 \operatorname{set}\left(\mathbf{g c f},{ }^{\prime}\right.$ DefaultAxesColorOrder, , $\left[\begin{array}{llllllllllllll}0 & 0 & 0 ; & 1 & 1 & 1 ; & 0 & 0 & 1 ; & 0 & 1 & 0 ; & 1 & 0\end{array}\right.$ 0]) ;
guidata(hObject, handles);
\% UIWAIT makes pitchVisGUI wait for user response (see UIRESUME)
32 \% uiwait(handles.figure1); handles)

36 \% Get default command line output from handles structure
37 varargout $\{1\}=$ handles. output;
str2double(get(hObject, 'String')) returns contents of
frameLengthField as a double ;

44 guidata(hObject, handles);
\% ——— Executes during object creation, after setting all properties
function frameLengthField_CreateFcn(hObject, eventdata, handles)
\% Hint: edit controls usually have a white background on Windows.
\% See ISPC and COMPUTER.
if ispc \&\& isequal (get (hObject, 'BackgroundColor'), get ( 0, ,
defaultUicontrolBackgroundColor'))
set (hObject, 'BackgroundColor', 'white ');
end
function noTonesField_Callback(hObject, eventdata, handles)
\% Hints: get(hObject, 'String') returns contents of noTonesField as
text
\% str2double(get(hObject, 'String')) returns contents of
noTonesField as a double
set (hObject, 'UserData', str2double (get (hObject, 'String')));
guidata(hObject, handles);
\% ——— Executes during object creation, after setting all properties
function noTonesField_CreateFcn(hObject, eventdata, handles)
\% Hint: edit controls usually have a white background on Windows.
\% See ISPC and COMPUTER.
if ispc \& isequal (get (hObject, 'BackgroundColor'), get ( $0,{ }^{\prime}$
defaultUicontrolBackgroundColor ') )
set (hObject, 'BackgroundColor', ' white');
end
\% ——— Executes on button press in browseButton.
function browseButton_Callback(hObject, eventdata, handles)

-analyzed') ;
if isstr(fileName) \% user made a valid selection
$\mathrm{s}=$ [filePath fileName]; \% the current selection
if ${ }^{\text {strcmp }}(\boldsymbol{g e t}($ hObject, $\quad$ UserData'), s) \% inequality with
previous selection
handles.meansMatrix $=[] ;$
handles.timeAxis= [];
end
set (handles.runButton, 'Enable', 'On') ;
set (handles.playButton, 'Enable', 'On');
set (handles.FileNameDisplayText, 'String', fileName);
else \% this happens if dialog gets cancelled
$\mathrm{s}=$ ' none ';
set (handles.runButton, 'Enable', 'Off');
set(handles.playButton, 'Enable', 'Off');
set (handles.FileNameDisplayText, 'String', s);
end
set (hObject, 'UserData', s);
guidata(hObject, handles);
\% ——— Executes on button press in saveFigureButton.
function saveFigureButton_Callback(hObject, eventdata, handles)

include」extension') ;
if ischar (fileName)
plotCp= findobj(gcf, 'Type', 'axes'); \% 1) legend 2) axes
$\mathrm{p}=\boldsymbol{\operatorname { g e t }}(\mathrm{plot} \mathrm{Cp}(2), \quad$ Position' $)$;
shift $=\left[\begin{array}{llll}-.00 & -.1 & 0.02 & 0.09\end{array}\right] ; \%$ if you enlarge the copy in the
new window,
$\mathrm{p}=\mathrm{p}+\mathrm{shift} ; ~ \% ~ m a t l a b$ will erase everything (bug?)
$\boldsymbol{\operatorname { s e t }}(\mathrm{plot} \mathrm{Cp}(2), \quad$ Position', p$)$;
newFig= figure('Visible', 'off');
newAxes $=$ copyobj (plotCp, newFig) ;
$\boldsymbol{s e t}(\operatorname{plotCp}(2), \quad$, Position', $\mathrm{p}-\mathrm{shift})$;
saveas (newFig, [pathName fileName]) ; \% only for whole figures!
delete (newFig) ;
end
function runButton_Callback(hObject, eventdata, handles)
fileLocation $=\operatorname{get}($ handles.browseButton, 'UserData') ; \%get: [path filename]
tonesNo $=$ get (handles.noTonesField, 'UserData');
windowLength $=\boldsymbol{g e t}($ handles.frameLengthField, 'UserData');
tic;
clc
handles.EMOut= runEM(fileLocation, windowLength, tonesNo);
set(handles.statisticsText, 'String', ['exec」time: $\lrcorner^{\prime}$ num2str(toc) ' $\mathrm{s} ; \quad$ '...
'total EM_iterations: 」' num2str(handles.EMOut.iterEM)]);
\% display specgram. w/ time slider if necessary.
imagesc ([handles.EMOut.timeAxis (1) handles.EMOut.timeAxis (end) ],... [handles.EMOut.freqAxis (1) handles.EMOut.freqAxis (round (end $/ 2$ )) ] ,.
$\log ($ handles.EMOut. spectralMatrix $(1: \operatorname{round}($ end $/ 2),:))) ;$
set (gca, 'YDir', 'normal') ; \% b/c image-function behaves abnormally

119

], 'Interpreter', 'none');
121 xlabel('Time」in」seconds');


shiftPlot $=\left[\begin{array}{llll}0 & 0.02 & 0 & -0.02\end{array}\right] ;$
$\operatorname{pos}=\operatorname{get}($ gca, $\quad$ position' $) ;$
if ${ }^{\sim}$ isempty (handles.sliderHandle) \% clear old stuff
delete(handles.sliderHandle);
handles.sliderHandle= [];
set (gca, 'position', [ pos(1) pos(2) pos(3) pos(4)]-shiftPlot);
end
if handles.EMOut.timeAxis (end) $>2$;
$\mathrm{d} \mathrm{x}=2 ; \quad \% d x$ is the width of the axis 'window'
\% Set appropriate axis limits and settings
set (gcf, 'doublebuffer', 'on');
\% This avoids flickering when updating the axis
set (gca, 'xlim', [windowLength dx], 'ylim', [0 handles.EMOut.
freqAxis (round (end/2))]);
$\operatorname{pos}=\operatorname{get}($ gca, $\quad$ position' $) ;$
set (gca, ' position', $[\operatorname{pos}(1) \operatorname{pos}(2) \operatorname{pos}(3) \operatorname{pos}(4)]+\operatorname{shiftPlot);~}$
\% Generate constants for use in uicontrol initialization
scrollBarPos= $[\operatorname{pos}(1) \operatorname{pos}(2)-0.1 \operatorname{pos}(3) 0.05] ;$
\% This will create a slider which is just underneath the axis
\% but still leaves room for the axis labels above the slider
xmax $=$ handles.EMOut.timeAxis (end);
$S=\left[' \operatorname{set}\left(\right.\right.$ gca $, ',{ }^{\prime} \operatorname{xlim}{ }^{\prime}$, , get (gcbo,,$'$ value $\left.{ }^{\prime}{ }^{\prime}\right)+['$ num2str $($
windowLength ) ' ', num2str(dx) ']) '];
\% Setting up callback string to modify XLim of axis (gca)
\% based on the position of the slider (gcbo)
\% Creating Uicontrol
handles.sliderHandle= uicontrol('style', 'slider',$\ldots$
'units', 'normalized', ' position', scrollBarPos , ...
'callback', S , 'min' , 0 , ' $\max { }^{\prime}$, xmax-dx);
end
handles.fileLocCurrentPlot= fileLocation; \% update fileLocation to
current
hold on; \% add means traces and pitch to specgram
\% add pitch estimation plot to specgram
handles. pitchPlotHandle= plot(handles.EMOut.timeAxis, handles.EMOut . pitches, 'LineWidth', 2);
hold off;
for $t=1$ :tonesNo

end
legHandle= legend (handles. pitchPlotHandle, legendStrings);
set (legHandle, ${ }^{\prime}$ Color', $\left.\left[\begin{array}{lll}.7 & .9 & .7\end{array}\right]\right)$;
set (handles.axes1, 'Visible', 'On');
set (handles.saveFigureButton, 'Enable', 'On');
guidata(hObject, handles);
\% - Executes on button press in playButton.

169 function playButton_Callback(hObject, eventdata, handles)
170 \% hobject handle to playButton (see GCBO)
171 \% eventdata reserved - to be defined in a future version of MATLAB
172 \% handles structure with handles and user data (see GUIDATA)
173 fileLocation= get(handles.browseButton, 'UserData');
174 [wave, Fs, bits]= wavread(fileLocation);
175 wavplay (wave, Fs);

## B. 2 STFT Analysis

```
function out= runEM(fileLocation, windowLength, tonesNo)
% out= runEM(fileLocation, windowLength, tonesNo)
% STFT analyis of file (50% overlap), invokes EM for pitch
    estimation
% - fileLocation: path and filename, concatenated
% - windowLength: unit shall be seconds.
% common definition: windowLength==64msec
% - tonesNo: number of tones assumed to be hidden in file
%-out.spectralMatrix: for spectrogram plotting of
% the input file, unfiltered
% - out.freq: frequency points for which DFT was computed
% - out.timeAxis: vector of frame analysis start points
% (in sec), for plotting the out.pitches and specgram
% - out.pitches: columns correspond to frames, rows to
% (maximally tonesNo) pitches. Inf entry/entries indicate
% that for this frame: no of actual pitches < tonesNo
% - out.iterEM: total number of EM iterations for the file
%
% Author: Dominik Loeffler Date: Mar 15, 2006
% Version: 7
if nargin }~=3 || \mp@subsup{\mathrm{ nargout }}{~}{~}=
```



```
                    ');
```

end
if windowLength $<=2 *$ eps


54
end
［path，fileName，ext，versn］＝fileparts（fileLocation）；
try

```
    [wave, Fs, bits]= wavread(fileLocation);
```

catch
errordlg ('Could„notょopen」file.') ;
error ('pitchEstEM: „could」not」open」file.');
end
\％windowSize equals to 1024 samples at 16 kHz and 0.064 s windowSize $=$ round $($ Fs $*$ windowLength $) ;$
if $\bmod ($ windowSize, 2$)==1$
windowSize $=$ windowSize +1 ；\％ensure $i t$＇s always even
end
$\operatorname{disp}([$＇Analyzing＿file＇fileName］）；
fftSize＝ $\max ($ windowSize，256）；
\％freqs for which DFT will be computed out．freqAxis $=0:$ Fs／windowSize：Fs／2－Fs／windowSize；
\％advancing window in windowSize／2 steps always（ $->50 \%$ overlap）
indices＝windowSize：windowSize／2：length（wave）；
\％standard hamming window；size＝＝windowSize $x 1$
hamm $=$ hamming（windowSize）；
out．spectralMatrix $=\boldsymbol{z e r o s}($ windowSize $/ 2$ ，length（indices）$) ;$
$\operatorname{specSmoothTemp}=\operatorname{zeros}($ windowSize $/ 2,1) ;$
$\mathrm{nAvg}=2 ; \%$ smoothing parameter for running average
\％for each frame，save all the obtained means for each mixture out．pitches＝zeros $($ tonesNo，length（indices $))$ ；
spec $=\operatorname{zeros}($ windowSize， 1$)$ ；

```
    iterFrame= 0;
    out.iterEM= 0;
    for k= indices
    iterFrame= iterFrame +1;
    disp(['runEM:_frame_' num2str(iterFrame) ',\triangleleftof `'...
        num2str(length(indices))]);
    spec= hamm.* wave(k-windowSize +1:k);
    if sum(realpow (spec,2))/windowSize <= 7*10^-8 % silence
        detector
        spec= abs(fft(spec, fftSize));
        out.spectralMatrix (:, iterFrame)= spec(1: windowSize / 2);
        out.pitches(:, iterFrame)= Inf*ones(tonesNo, 1);
        disp('runEM:\_silence_detected; EM_skipped_for\_this frame');
        continue;
    end
    spec= abs(fft(spec, fftSize));
    out.spectralMatrix (:, iterFrame ) = spec (1: windowSize / 2);
    specSmoothTemp= filter(ones(1,nAvg)./nAvg, 1, ...
        out.spectralMatrix(:, iterFrame));
    %specSmoothTemp= out.spectralMatrix (:, iterFrame);
    % find hi-freq cutoff frequency
    ind= find (specSmoothTemp>=0.1*max (specSmoothTemp));
    cutoffInd= ind(end)+round(windowSize/2/100);% add 1% of total
        length
    [pitches, priorProbs, variances, iterations]= ...
        myEM(out.freqAxis (1:cutoffInd) , ...
            specSmoothTemp(1:cutoffInd)', tonesNo);
    out.iterEM= out.iterEM + iterations;
    out.pitches(:, iterFrame)= pitches;
```

82 end
83

84 out.timeAxis= indices./Fs;

86 end \% function runEM

## B. 3 Core Heuristic

3 \% Core heuristic for EM algorithm initialization
\% - freq: row vector (only 1-D data!) frequency pts for which spectrum def.
\% - spec: fft magnitude value, abs(fft)
$\%$ - tonesNo: number of polyphonic tones we assume to be hidden in source
\% - harmonicsFreq: freq points at which heuristic supposes harmonics of $a$
\%
pitch. cell array of (tonesNo,?) dimensions, 2nd dim expanded
\% dynamically for each tone
\% - hInd: indices in freq corresponding to harmonics in harmonicsFreq
$11 \%$
12 \% Author: Dominik Loeffler Date: Mar 30, 2006
13 \% Version: 7 (+ weak harmonics interpolation)
14 if ~isequal(size(freq), size(spec))
15

end
17 Nfreq= length (freq);
18 if tonesNo $>$ Nfreq
19
20
 'estimating $\quad$ ' num2str(tonesNo) ' „tones']);
end
[peakVal, peakInd] $=\operatorname{lmax}(\operatorname{spec}, 0, f a l s e) ;$
\% filter out very small peaks which are unlikely real harmonics;
\% cut off some low-frequency peaks that cause trouble when we assume them
\% to be the fundamental below. mostly occurs at very low pitches.
cutoffPeakAmpl=0.05;
$\operatorname{sig} \operatorname{Indices}=$ find $($ peakVal $>=$ cutoffPeakAmpl $* \max ($ peakVal) \& $\ldots$
freq $($ peakInd $)>45)$;
$\operatorname{sig} P e a k=$ peakVal (sigIndices) $;$
sigPeakFreq= freq (peakInd (sigIndices));
harmonicsFreq= num $2 \operatorname{cell}(\mathbf{I n f} *$ ones $($ tonesNo, 1$))$;
hInd= num2cell(zeros (tonesNo, 1));
$\operatorname{sig}$ PeakFreqTemp $=\operatorname{sig}$ PeakFreq;
for $t=1$ : tonesNo
\% need to have_at least_ as many significant peaks as tones
if isempty (sigPeakFreqTemp) \% no more peaks available $\rightarrow$ quit
 - for $\quad$ '...

 return;
end
currentFundamental $=\operatorname{sig}$ PeakFreqTemp (1) ;
harmonicsFreq $\{t\}=\operatorname{Inf} *$ ones $(1, \quad$ round $(\operatorname{sigPeakFreqTemp}($ end $) /$ currentFundamental)) ;
$\operatorname{hInd}\{\mathrm{t}\}=\operatorname{zeros}(\boldsymbol{\operatorname { s i z e }}($ harmonicsFreq$\{\mathrm{t}\}))$;
\% tolerance in Hz hat harmonics can deviate at the max from \% their ideal frequency locations
tol $=\min (50$, currentFundamental $/ 4) ;$
harmonicsFreq $\{\mathrm{t}\}(1)=$ currentFundamental ;
\% remove so in next iteration choose different fundamental
$\operatorname{sig} P e a k F r e q T e m p(1)=[] ;$
$\operatorname{hInd}\{\mathrm{t}\}(1)=\operatorname{find}(\operatorname{freq}=\operatorname{currentFundamental}) ;$
for $k=2: \operatorname{length}($ harmonicsFreq $\{t\})$
binIndices $=\boldsymbol{f i n d}(\mathbf{a b s}(\operatorname{sig}$ PeakFreqTemp $-\mathrm{k} *$ currentFundamental $)$ $<=$ tol) ;
if ${ }^{\sim}$ isempty (binIndices)
indClosest $=$ findapprox ( $\operatorname{sigPeakFreqTemp(binIndices)~,~k*~}$ currentFundamental);
harmonicsFreq $\{\mathrm{t}\}(\mathrm{k})=\operatorname{sig}$ PeakFreqTemp (binIndices $($ ind Closest) ) ;
\%harmonicsFreq $\{t\}(k)=$ sigPeakFreqTemp (binIndices (round) ( length(binIndices)/2)));
$\operatorname{hInd}\{\mathrm{t}\}(\mathrm{k})=\operatorname{find}(\operatorname{freq}=$ harmonicsFreq$\{\mathrm{t}\}(\mathrm{k}))$;
\% remove the peak just found
$\operatorname{sig} P e a k F r e q T e m p(b i n I n d i c e s(i n d C l o s e s t))=[] ;$
end \% else: leave the preinitialized Inf entries there
end
$\% * * * V 7$ : optimize this here by interpolating missing entries
$* * *$
\% remove any trailing consecutive Inf entries
indNotInf $=$ find $\left(\right.$ harmonicsFreq $\left.\{t\}^{\sim}=\mathbf{I n f}\right) ;$
$\operatorname{harmonicsFreq}\{\mathrm{t}\}=\operatorname{harmonicsFreq}\{\mathrm{t}\}(1: \operatorname{ind} \operatorname{NotInf}($ end $))$;
$\operatorname{hInd}\{\mathrm{t}\}=\mathrm{hInd}\{\mathrm{t}\}(1: \operatorname{ind} \operatorname{Not} \operatorname{Inf}($ end $)) ;$

```
    % now interpolate remaining Inf entries, if any
    infInd= find(harmonicsFreq{t}==Inf);
    if ~isempty(infInd) % GB some other condition?
        for k= infInd
        binIndices= find(abs(freq(peakInd)-k*currentFundamental)
            <= tol);
        if ~isempty(binIndices) % attempt to interpolate with
                low-amp peak
                indClosest= findapprox(freq(peakInd(binIndices)), k*
                    currentFundamental);
        harmonicsFreq{t}(k)= freq(peakInd(binIndices(
                indClosest)));
        hInd{t}(k)= find(freq=harmonicsFreq{t}(k));
                %disp('pitchTrace: interpolated missing harmonic');
        else % interpolation failed - stop here
            harmonicsFreq{t}= harmonicsFreq{t}(1:k-1);
            hInd {t}=hInd{t}(1:k-1);
            break;
        end
        end
    end
end % for
    % ___ only if verbose is selected
    if ~}\mathrm{ verbose
        return;
    % graphical output (optional)
```

    end
    figure (2);
cla;
hold on;
plot(freq, spec); \% spectrum
\%stem(sigPeakFreq, sigPeak, 'gd', 'LineWidth', 1.6); \%significant peaks

```
plotProp= {'ro', 'k*', 'r^',, 'y'};
```

legendString $\{1\}=$ ' spectrum ';
for $t=1$ :tonesNo
stem(harmonicsFreq $\{\mathrm{t}\}$, ...
$(1.1-\mathrm{t} * 0.2) * \max (\operatorname{peakVal}) * \operatorname{ones}(1, \operatorname{length}(\operatorname{harmonicsFreq}\{\mathrm{t}\}))$,
plotProp $\{\bmod (\mathrm{t}, 4)\}) ; \%$ pitch trace
legendString $\{\mathrm{t}+1\}=\left[\right.$ 'tone ${ }^{\prime}$ ' num2str $\left.(\mathrm{t})\right]$;
$\% \operatorname{stem}\left(\operatorname{freq}(h \operatorname{Ind}\{t\}), \operatorname{spec}(h \operatorname{Ind}\{t\}), \quad{ }^{\prime} g^{\wedge}\right)$;
end
hold off;
legend (legendString);

') ;
xlabel ('f/Hz') ;
ylabel ( $\left.{ }^{\prime} \mathrm{FFT}(\mathrm{f}) / \max (\mathrm{FFT}(\mathrm{f}))^{\prime}\right)$;
disp(['pitchTrace:„sign.』peaks: 」' num2str(freq(peakInd))]);
for $t=1$ : tonesNo

num2str (harmonicsFreq\{t $\}$ )]);
\%disp (['pitchTrace (tone ' num2str ( $t$ ) '): we found '...
\% num2str(100*length(find(harmonicsFreq\{t $\left.\left.\}^{\sim}=\operatorname{Inf}\right)\right) \ldots$

127 end \% function pitchTrace

## B. 4 Expectation Maximization Algorithm

function [pitches, priorProbs, variances, iterations]= myEM(freq, spec, tonesNo)

2 \% [pitches, priorProbs, variances, iterations]= myEM(freq, spec, tonesNo)

3 \% Expectation Maximization Algorithm, constrained, working on GMM
$\%$ - freq: row vector (only $1-D$ data!) frequency $p t s$
5 \% for which spectrum defined
6 \% - spec: fft magnitude value, abs(fft)
$7 \%$ - tonesNo: number of tones assumed to be hidden in file
$8 \%-$ pitches: pitch estimates for each tone
$9 \%$ - priorProbs: best estimate of prior prob for each mixture $(P(X \mid$ mix_i))

10 \% - variances: analogous...
11 \% - iterations: iterations needed for EM alg to converge
$12 \%$
13 \% Author: Dominik Loeffler Date: Mar 15, 2006
14 \% Version: 7
15
16 Nfreq= length(freq);
17 if ~isequal(size(freq), size(spec))
 precisely');

19
end
20 sumSpec= sum(spec);
21 if sumSpec/Nfreq < eps \% should be caught by silence detector on caller side

silence! ) ') ;
end
spec $=$ spec./sumSpec ;
\% use this heuristic to find good initial estimates
\% meansOld: a cell array of tonesNo elements, each element with
variable
\% components as determined by the heuristic "pitchTrace"
$[$ meansOld, meansInd $]=$ pitchTrace (freq, spec, tonesNo, false);
\% 2nd dim will be expanded in the following loop
means $=\operatorname{num} 2 \operatorname{cell}(\mathbf{I n f} *$ ones $($ tonesNo , 1) $) ;$
variances $=$ means ; \% stays constant in this version
priorProbsOld= means; priorProbs= means;
mixtures $=\operatorname{zeros}(1$, tonesNo $) ; \%$ no of peaks for each tone
pitches $=\mathbf{I n f} *$ ones (tonesNo, 1);
iterations $=1$;
for $t=1$ :tonesNo
\% heuristic couldn't come up with anything sensible
if meansOld $\{\mathrm{t}\}(1)=\mathbf{I n f}$
$\operatorname{disp}\left(\left[{ }^{\prime} \mathrm{myEM}_{\lrcorner}\left(\right.\right.\right.$tone $_{\sqcup}{ }^{\prime}$ num2str$(\mathrm{t})$. .

return;
end
mixtures $(t)=\operatorname{length}(\operatorname{meansOld}\{t\}) ;$
variances $\{\mathrm{t}\}=10 * \operatorname{ones}(1$, mixtures $(\mathrm{t}))$;
priorProbsOld $\{\mathrm{t}\}=\operatorname{spec}(\operatorname{meansInd}\{\mathrm{t}\}) ; \%$ then normalize
priorProbsOld $\{\mathrm{t}\}=\operatorname{priorProbsOld}\{\mathrm{t}\} . / \operatorname{sum}(\operatorname{priorProbsOld}\{\mathrm{t}\}) ;$
$\operatorname{means}\{\mathrm{t}\}=\operatorname{zeros}(1, \quad \operatorname{mixtures}(\mathrm{t}))$;
priorProbs $\{\mathrm{t}\}=\operatorname{means}\{\mathrm{t}\}$;

${ }^{\prime}$ ) : „init (heur.) „means: 」, num2str(meansOld\{t\})]);
\%disp (['myEM: initial priorProbs: , num2str(priorProbsOld\{t\})])
;
\%disp(['myEM: initial variance: ' num2str(variances $\{t\})]$ );
$\mathrm{y}=\operatorname{zeros}($ mixtures (t), Nfreq);
\%meansOld $\{t\}=\ldots$
$\% \quad[1: \operatorname{mixtures}(t)] . *\left(\operatorname{sum}\left((1: \operatorname{mixtures}(t)) .^{\wedge}-1 . * \operatorname{meansOld}\{t\}\right) \ldots\right.$
\% /mixtures ( $t$ ) ; \% resetting with good initial pitch
estimate
firstPeak= meansOld $\{\mathrm{t}\}(1)$;
$\operatorname{meansOld}\{\mathrm{t}\}=[1: \operatorname{mixtures}(\mathrm{t})] . * \operatorname{meansOld}\{\mathrm{t}\}(1) ;$

$(\operatorname{meansOld}\{t\})]) ;$
deltaF $=0 ;$ deltaFOld $=0$;
while (1) \% loop for the EM of _one_ pitch
$\operatorname{sum} Y=\operatorname{zeros}(1$, Nfreq $) ; \% E-s t e p$
for $k=1:$ mixtures $(t)$
$\mathrm{y}(\mathrm{k},: \mathrm{a})=\operatorname{priorProbsOld}\{\mathrm{t}\}(\mathrm{k}) * \ldots$
$\operatorname{myGauss}($ freq,$\quad[\operatorname{sqrt}($ variances $\{t\}(k)), \operatorname{meansOld}\{t\}(k$
)], $\quad$ true) $;$
$\operatorname{sum} Y=\operatorname{sum} Y+y(k,:) ;$
end
sumYZeroInd= find $($ sumY $==0)$;
if ${ }^{\text {isempty ( }}$ isumYeroInd)
$\operatorname{sumY}(\operatorname{sumYZeroInd})=$ eps; \%disp ( [ 'myEM (tone ' num2str ( $t$ ) ... \% '): zero sum of likelihoods occurred, fixing it now']) ;
end
for $k=1$ : mixtures $(t)$

$$
\mathrm{y}(\mathrm{k},:)=\mathrm{y}(\mathrm{k},:) . / \mathrm{sumY} ;
$$

specMultY= spec. $* \mathrm{y}(\mathrm{k},:) ;$
sumSpecMult $Y=\operatorname{sum}($ specMult $Y) ;$
if sumSpecMult $Y==0$
\% seems to happen when unsmoothed or few distro values
means $\{\mathrm{t}\}=\operatorname{zeros}(1, \operatorname{mixtures}(\mathrm{t}))$;
priorProbs $\{\mathrm{t}\}=\operatorname{zeros}(1, \operatorname{mixtures}(\mathrm{t}))$;
variances $\{t\}=\operatorname{zeros}(1, \operatorname{mixtures}(t))$;
 $==0) \quad$ ') ;
return;
end
priorProbs $\{\mathrm{t}\}(\mathrm{k})=$ sumSpecMultY; \% M-step
means $\{\mathrm{t}\}(\mathrm{k})=\operatorname{sum}($ freq $\cdot *$ specMultY) $/$ sumSpecMultY;
end
\%deltaF= $\operatorname{sum}($ means $\{t\}-$ means Old $\{t\}) / \operatorname{sum}($ abs $($ means $\{t\}-$ means $\operatorname{Old}\{t\})$ );
\% if pos: increases in frequency higher than decreases deltaF $=\operatorname{sum}\left((1: \operatorname{mixtures}(\mathrm{t})) .^{\wedge}-1 \cdot *(\operatorname{means}\{\mathrm{t}\}-\operatorname{meansOld}\{\mathrm{t}\})\right) /$ mixtures (t) ;

121 end \% function myEM

## B． 5 Gaussian Function

```
function \(g=\) myGauss(input, params, dirtyFix)
\% g= gauss(input, [sigma mean], dirtyFix)
\% like gaussmf, but gaussmf is not good enough:
\% need cases with sigma==0 and the scaling factor
\% 1/(sqrt(2*pi)*sigma) for most general uses.
\% note that sigma denotes standard deviation==sqrt(variance)
\% - input: \(1-D\) vector with real values only
\%-params: \(1 \times 2\) vector:
\% params(1)=standard deviation, params(2)=mean
\% - dirtyFix: logical, indicates whether or not a case of
\% variance=0 should be fixed by returning 1000 for the input
\% value closest (not necessarily equal!) to the mean specified
\% in "params". This guarantees that the output vector be nonzero.
\%
\% Author: Dominik Loeffler Date: Oct 20, 2005 Last update: Nov 06
\% Version: 5 (+ real values optimization w/ realpow,
\(\% \quad+\) dirtyFix
if \(\boldsymbol{n a r g i n}^{\sim}=3\)
    error ('myGauss:」wrong_number_of」input」arguments');
end
if ( \({ }^{\text {isequal }}\left(\operatorname{size}(\right.\) params \(\left.\left.),\left[\begin{array}{ll}1 & 2\end{array}\right]\right) \| \operatorname{params}(1)<0\right)\)
    error ('myGauss:」wrong」input」parameters」"params"') ;
end
if \({ }^{\text {isvector (input) }}\)
    error ('myGauss: input」is not」a」vector');
end
```

    if isempty (input)
    
end
if ~islogical(dirtyFix)

end
$\operatorname{sigma}=\operatorname{params}(1) ;$
mean $=$ params (2);
if $\operatorname{sigma}==0 \% \rightarrow$ gauss goes to infinity in continuous case
$\mathrm{g}=\operatorname{zeros}(\operatorname{size}($ input $)) ;$
if dirtyFix
\% choose the next two lines if your "input" vector can
\% have some arbitrary values, i.e. its elements are
\% not unique
\%closest=find(abs(input-mean) $==\min (a b s($ input - mean $)))$;
$\%$ (closest(round(length(closest)/2)))=1000;
$[\mathrm{m}$ mIndex $]=\min (\mathbf{a b s}($ input - mean $)) ;$
$g($ mIndex $)=1000 ;$
else
equal $=$ find $($ input $=$ mean $)$;
if isempty (equal)

'vector」due」to」zero』variance']) ;
end
$g($ equal $)=1000 ;$
end
else
$\mathrm{g}=(1 /(\mathbf{s q r t}(2 * \mathbf{p} \mathbf{i}) * \operatorname{sigma})) * \ldots$

```
                exp(-realpow (input-mean, 2) / (2*realpow (sigma, 2)) );
```

    end
    59

60 end \% function myGauss

## B. 6 Local Maxima

```
    function [lmval, indd]=lmax(xx, filt, strict)
    %LMAX [lmval, indd]=lmax(xx,filt). Find local maxima
    % in vector XX, where LMVAL is the output vector with
    % maxima values, INDD are the corresponding indexes,
    % FILT is the number of passes of the small running
    % average filter in order to get rid of small peaks.
    % Default value FILT=0 (no filtering). FILT in the
    % range from 1 to 3 is usually sufficient to remove
    % most of a small peaks
```

10
\% - - - - - - - - - - - 06/03/97
15
filt $=0$;
21 else
$\mathrm{x} 1=\mathrm{x}(1) ;$
$\mathrm{x} 2=\mathrm{x}\left(\mathrm{len}_{-} \mathrm{x}\right) ;$
for $\mathrm{j}=1: \mathrm{filt}$
$\mathrm{c}=\boldsymbol{\operatorname { c o n v }}(\mathrm{fltr}, \mathrm{x}) ;$
$\mathrm{x}=\mathrm{c}\left(2: \operatorname{len}_{-} \mathrm{x}+1\right)$;
$\mathrm{x}(1)=\mathrm{x} 1$;

$$
x\left(\operatorname{len}_{-} x\right)=x 2
$$

end
end
$\operatorname{lmval}=[] ; \quad$ indd $=[] ;$
$\mathrm{i}=2 ; \quad \%$ start at second data point in time series
while $\mathrm{i}<$ len_x -1
if $\mathrm{x}(\mathrm{i})>\mathrm{x}(\mathrm{i}-1)$

$$
\text { if } \mathrm{x}(\mathrm{i})>\mathrm{x}(\mathrm{i}+1) \% \text { definite } \max
$$

lmval $=[$ lmval $x(i)] ;$
indd $=[$ indd i$] ;$ elseif $(x(i)=x(i+1)) \& \&(x(i)=x(i+2)) \%{ }^{\prime}$ long', flat spot if ~strict
lmval $=[\operatorname{lmval} x(i)] ;$
indd $=[$ indd i$]$;
end
$\mathrm{i}=\mathrm{i}+2 ; \%$ skip 2 points elseif $\mathrm{x}(\mathrm{i})=\mathrm{x}(\mathrm{i}+1) \%$ 'short, flat spot $\mathbf{i f} \sim$ strict lmval $=[\operatorname{lmval} x(i)] ;$ indd $=[$ indd i$] ;$ end $\mathrm{i}=\mathrm{i}+1 ; \%$ skip one point
end
end
$\mathrm{i}=\mathrm{i}+1$;
end
if filt $>0 \& \&{ }^{\sim}$ isempty (indd)
if $(\operatorname{indd}(1)<=3) \|(\operatorname{indd}(\operatorname{length}(\operatorname{indd}))+2>\operatorname{length}(x x))$
rng $=1 ;$ \% check if index too close to the edge
else
$\mathrm{rng}=2 ;$
end
for $\mathrm{i}=1$ length(indd) \% Find the real maximum value
$[\operatorname{val}(\mathrm{ii}) \operatorname{ind}(\mathrm{ii})]=\max (x x(\operatorname{indd}(\mathrm{ii})-r n g: \operatorname{indd}(\mathrm{ii})+r n g))$;
iind (ii $)=\operatorname{indd}(\mathrm{i} i)+\mathrm{iind}(\mathrm{i} i)-\mathrm{rng}-1$;
end
indd= iind;
lmval= val;
end

## B. 7 Findapprox

1 function $\mathrm{i}=$ findapprox (vector, scalar)
2 \% i= findapprox(vector, scalar)
$3 \%$ Find closest element in vector, akin to find.
4 \% - scalar: value we're trying to match against vector
$5 \%$ - vector: of same datatype as scalar
6 \% if there is no element in vector equal to scalar, the index
7 \% of the element being closest to it will be returned (thus
8 \% the name findAPPROX). If more that could be more than one index,
9 \% the first (leftmost) is taken only.
$10 \%$
11 \% USE WITH CAUTION. it doesn't check for input errors.
12 \% Author: Dominik Loeffler Last update: Nov 26, 2005
end \% function findapprox

## REFERENCES

[1] Akaike, H., "A new look at the statistical model identification," IEEE Transactions on Automatic Control, vol. 19, no. 6, pp. 716-723, 1974.
[2] AssociatedPress, "Online and wireless music sales tripled in 2005." WWW, January 2006. Accessed: March 06, 2006.
[3] Brown, J. C., "Calculation of a constant q spectral transform," The Journal of the Acoustical Society of America, vol. 89, no. 1, pp. 425-434, 1991.
[4] Casey, M. A. and Westner, A., "Separation of mixed audio sources by independent subspace analysis," in Proceedings of ICMC, pp. 154-161, 2000.
[5] Duda, R. O., Hart, P. E., and Stork, D. G., Pattern Classification. Hoboken, NJ: John Wiley \& Sons, second ed., 2001.
[6] Fechner, G. T., Elemente der Psychophysik. Bristol, UK: Thoemmes Press, second ed., 1889. First ed. 1860.
[7] Feder, M. and Weinstein, E., "Parameter estimation of superimposed signals using the em algorithm," IEEE Transactions on Acoustics, Speech, and Signal Processing, vol. 36, no. 4, pp. 477-489, 1988.
[8] Fitzgerald, D., Automatic Drum Transcription and Source Separation. PhD thesis, Dublin Institute of Technology, 2004.
[9] Fitzgerald, D., "Generalised prior subspace analysis for polyphonic pitch transcription," in Proceedings of DAFX, 2005.
[10] Fitzgerald, D., Lawlor, R., and Coyle, E., "Prior subspace analysis for drum transcription," in Proceedings of AES $114^{\text {th }}$ Convention, 2003.
[11] Greenberg, S. and Kingsbury, B. E. D., "The modulation spectrogram: in pursuit of an invariant representation of speech," in Proceedings of ICASSP, 1997.
[12] Howard, D. M. and Angus, J., Acoustics and Psychoacoustics. Burlington, MA: Focal Press/Elsevier, 1996.
[13] Hyvarinen, A., Karhunen, J., and Oja, E., Independent Component Analysis. Hoboken, NJ: John Wiley \& Sons, 2001. GT library call number: QA278 .H98 2001.
[14] ISO, "Acoustics - standard tuning frequency (standard musical pitch): Iso 16." WWW, 1975.
[15] Kameoka, H., Nishimoto, T., and Sagayama, S., "Accurate f0 detection algorithm for concurrent sounds based on em algorithm and information criterion," in Proceedings of Special Workshop in Maui (SWIM), 2004.
[16] Kameoka, H., Nishimoto, T., and Sagayama, S., "Separation of harmonic structures based on tied gaussian mixture model and information criterion for concurrent sounds," in Proceedings of ICASSP, 2004.
[17] Klapuri, A., Virtanen, T., and Holm, J.-M., "Robust multipitch estimation for the analysis and manipulation of polyphonic musical signals," in Proceedings of DAFX, 2000.
[18] Laird, N. M., Dempster, A. P., and Rubin, D. B., "Maximum likelihood from incomplete data via the em algorithm," Ann. Royal Statistical Society, pp. 1-38, 1977.
[19] University of Iowa, Electronic Music Studios, "Musical instrument samples." WWW. Accessed: March 2006.
[20] Meredith, D. and Wiggins, G., "Comparing pitch spelling algorithms," in Proceedings of ISMIR, pp. 280-287, 2005.
[21] Mitchell, T. M., Machine Learning. New York, NY: McGraw-Hill, 1997.
[22] Ohgushi, K. and Ano, Y., "The relationship between musical pitch and temporal responses of the auditory nerve fibers," in Integrated Human Brain Science: Theory, Method, Application (Music) (NaKADA, T., ed.), pp. 357-364, Elsevier Science B. V., 2000.
[23] Pantev, C., Eulitz, C., Verkindt, C., Hampson, S., Schruierer, G., and ElBert, T., "Specific tonotopic organizations of different areas of the human auditory cortex revealed by simultaneous magnetic and electric recordings," Electroencephalography and clinical Neurophysiology, vol. 94, pp. 26-40, 1995.
[24] Sagayama, S., Takahashi, K., Kameoka, H., and Nishimoto, T., "Specmurt anasylis: A piano-roll-visualization of polyphonic music signal by deconvolution of log-frequency spectrum," in Proceedings of ISCA. SAPA, 2004.
[25] Saito, S., Kameoka, H., Nishimoto, T., and Sagayama, S., "Specmurt anasylis of multi-pitch music signals with adaptive estimation of common harmonic structure," in Proceedings of ISMIR, pp. 84-91, 2005.
[26] Schmidt-Jones, C., "Connexions project (cnx): Harmonic series." WWW. Accessed: March 2006.
[27] Stevens, S. S., "On the psychophysical law," Psychological Review, vol. 64, no. 3, pp. 153-181, 1957.
[28] Tervaniemi, M., "Automatic processing of musical information evidenced by eeg and meg recordings," in Integrated Human Brain Science: Theory, Method, Application (Music) (Nakada, T., ed.), pp. 325-335, Elsevier Science B. V., 2000.
[29] Tervaniemi, M., Kujala, A., Alho, K., Virtanen, J., Ilmoniemi, R., and NÄÄtÄnen, R., "Functional specialization of the human auditory cortex in processing phonetic and musical sounds: A magnetoencephalographic study," NeuroImage, no. 9, pp. 330-336, 1999.
[30] Titinen, H., Alho, K., Huotilainen, M., Ilmoniemi, R., Simola, J., and NÄÄTÄNEN, R., "Tonotopic auditory cortex and the magneto encephalographic (meg) equivalent of the mismatch negativity," Psychophysiology, vol. 30, pp. 537-540, 1993.
[31] Virtanen, T. and Klapuri, A., "Separation of harmonic sounds using linear models for the overtone series," in Proceedings of ICASSP, vol. 2, pp. 1757-1760, 2002.
[32] von Helmholtz, H. L. F., On the Sensations of Tone as a Physiological Basis for the Theory of Music. New York, NY: Dover Publications, 1954. First ed. (German) 1863. GT library call number: ML3820.H42 1957.
[33] Wikipedia, "Scientific pitch notation." WWW.
[34] Xue, W. and Sandler, M., "A partial searching algorithm and its application for polyphonic music transcription," in Proceedings of ISMIR, pp. 690-695, 2005.


[^0]:    ${ }^{1}$ Unfortunately, the authors of [19] did not return feedback concerning the question of how tuning was arranged and which pitch standard was used.

