Towards libre piano tuning software based on psychoacoustic features

University of Music and Performing Arts Graz

Toningenieur-Projekt für das Masterstudium Elektrotechnik-Toningenieur

Author: André Menrath (andre.menrath@posteo.de)

Advisor: Univ.Prof. Alois Sontacchi

29.06.2021

Contents

1 Introduction to piano tuning

Piano tuning has been a respected craftsmanship for as long as modern pianos have existed. Apart from the purely technical difficulties that arise when turning the tuning pin and the enormous piano construction background knowledge that is necessary in order not to cause any damage the piano, the real mastery is to consistently tune all 88 keys of the piano evenly well, especially in terms of their relations, i.e. perceptual quality of the resulting intervals. This refers, among other things, to the fact that nowadays all 12 notes of an octave are treated as equally important in a standard tuning. Assuming that the fundamental frequency of the octave of a piano note would be exactly twice as high, one could calculate the semitone steps with the frequency ratio $\sqrt[12]{2}$. Such a simplified approach, though, is not sufficient due to the acoustic properties of piano strings.

1.1 Partials and (octave-)stretching

The sound of a piano tone always consists as a combination of the fundamental frequency, and its harmonics. For a good piano tuning, it is less important that the fundamentals of two keys are in a certain frequency ratio, but rather that the nearby partial frequencies of different notes are as closely together as possible. This is because when two sinusoidal components with very similar frequency collide, disturbing beatings are perceived. Beatings can be described as two different psychoacoustic sensations: fluctuation strength and roughness, as discussed in detail in section 4.

Due to the stiffness and thickness of the piano strings the harmonics are slightly higher than the multiples of the fundamental. Equation 1 shows a model for calculating the frequency of the k-th harmonic via the basic frequency f_0 and the inharmonicity factor B (Fletcher, Blackham, and Stratton 1962).

$$f_k = kf_0\sqrt{1 + Bk^2} \tag{1}$$

The inharmonicity factor B changes over the pitch ranges of the piano. It has been noted (Rigaud, David, and Daudet 2013), that pianos are designed so that this change is as continuous as possible, although there are typical inconsistencies in the bass break where the strings are connected to the soundboard via a different bridge, and they begin to be wrapped. Taking the above criterion into account, considering an octave, for example, this leads to a frequency ratio greater than two for both the fundamental frequencies. This is called stretching. However, there is no optimal stretching, even for a single specific piano. It is a compromise, since this condition can never be implemented equally well for all partial tones involved in an interval, not even in the most consonant interval: the octave.

Finding an appropriate stretching for each individual piano and furthermore for each individual occasion is the essence problem in piano tuning. It is obvious that one of the most decisive components in the finding of this compromise is the loudness with which two partials collide. Fluctuation strength and roughness, both depend on the perceived loudness of the partial pure tones involved. The situation becomes even more complex because the weighting of the loudness of the partials of a tone also changes while it is fading out.

1.2 Aural tuning

1.2.1 Strechting

Finding the optimal stretching (especially for an octave, the most elementary and undoubtedly most consonant interval in music) works as follows for the piano tuner working by ear: perception is the determining factor, supported by experience.

1.2.2 Temperament

Modern music is usually based on equal temperament, which means that all intervals of all tones are equally important. Achieving an equal temperament aurally is rather complex due to the fact that the piano tuner can only compare two notes at a time. Therefore, the aural tuner is limited to certain procedures in which only certain interval combinations are compared.

1.2.3 Tuning procedure

The most common practice in aural tuning is to tune all keys in a single octave first. The octave from F3 to F4 is usually used for this, because the strings are longest and thinnest there, and this range therefore has the least inharmonicity in the piano. For this procedure, there are also several methods, which intervals will be compared with each other to achieve the most uniform result. The remaining keys are tuned primarily by listening to important intervals (octaves, duo-decimas, double-octaves etc.) available from the already tuned keys. Witch intervals are considered depends on which part of the piano is being tuned (Rasch and Heetvelt 1985). Like for the bass mainly octaves and multiple octaves are important, for the middle region fifths and tenths are regularly used as well and for the treble end quarts are taken into account frequently. This positioning of which intervals are more important in which range could also likely be found in average piano music.

1.3 The advantages of piano tuning algorithms

For determining the appropriate stretch for an octave, for a particular case: a particular piano, a particular octave, and beyond that for a particular musical occasion, the experienced aural tuner is likely superior to the algorithms. However, the individual perception of a tuner may well differ from the generally average perception and symptoms of fatigue in the course of a tuning may have a negative impact on the tuning quality.

Another disadvantage of aural tuning is that, as explained above, that a person can only approximate a equally tempered tuning, while an algorithm if the conditions for a *good sounding interval* would be known could solve this problem more accurately and consistently.

Finally, tuning with the help of software is possibly more cost-effective, as it usually takes less time.

2 Outline of existing non-proprietary procedures for piano tuning

2.1 Parametric model

In 2013 François Riguad and Bertrang David (Rigaud, David, and Daudet 2013) presented a parametric model and estimation techniques for the tuning of the piano. It is based on the fact that the inharmonicity coefficient B varies more or less continuously over the pitch range of the piano. B is parametrized for the whole piano range. The more notes are recorded in advance before the tuning is calculated, the better the approximation for B is likely to become.

2.1.1 Tuning algorithm

After the parameters for the parametrized model of the inharmonicity coefficient B of the whole piano range have been computed, at first all the octave intervals relative to the reference note A4 are tuned. From these computed pitches the tuning is interpolated for the whole piano.

2.1.2 Determine the right stretch for the octave

When tuning octaves, different partials collide, e.g. 2:1 (second partial of the higher octave note collides with the first partial of the upper one) or 4:2, etc. Which collision is most important to be eliminated is determined via the arbitrarily determined factor $\rho(m)$, which starts at a maximum of 5 in the bass and gradually decreases to 1 in the treble. If the partial pairs 2:1 and 4:2 would be most important this would lead to $\rho(m)$ between 1 and 2.

The fundamental frequency of an octave $F_0(m+12)$ can be computed like shown in equation 2. m is the key index and B is the inharmonicity factor.

$$F_0(m+12) = 2F_0(m)\sqrt{\frac{1+B(m)\cdot 4\rho(m)^2}{1+B(m+12)\cdot \rho(m)^2}} \tag{2}$$

However, tuning an octave via this continuous parametrization is missing one important fact: it is not taking the intensities of the partials into account directly, though via the estimation curves for ρ this effect is indirectly taken into account.

2.2 Entropy Piano Tuner

In the *Entropy Piano Tuner* (Hinrichsen 2012) the compromise of which partials of different piano tones overlap more and which less is based on trying to achieve a minimum entropy H of the joined power spectrum S, based on a logarithmic frequency discretization with the frequency bins m.

$$H = -\sum_m S(m)\log_2(S(m)) \tag{3}$$

$$S(m) = \sum_{keys} S_{key}(m) \tag{4}$$

2.2.1 Pre-recording

In order for the *Entropy Piano Tuner* to work, all notes of the piano must first be recorded separately. From this, the power density spectrum is calculated for each tone and the frequency bins are reorganized logarithmically, also known as constant Q transformation (CQT). The resolution of the spectrum is 1200 bins per octave. The lowest bin m_0 corresponds to the frequency $f_0 = 20.6017Hz$ and the highest $10542Hz$. Consequently, there is a total of 10800 CQT frequency components. The interval between two frequency bins corresponds to one cent and is thus smaller than the differential perceptibility threshold of pitch. The tuning of a single string can now be simulated approximately by shifting all values of the CQT.

2.2.2 Average-basic-tuning

Next, a kind of basic tuning is set up, based on empirical values or averages for the stretching of pianos tuned by ear in conjunction with the determined inharmonicity coefficients of each tone.

2.2.3 (Psychoacoustic) Preparation of the power density spectrum

Before these data are used for the core algorithm a preprocessing of the power density spectrum, also according to psychoacoustic aspects, takes place. The following processes are performed separately for each note.

Normalization First, the values for the power density spectrum are weighted so that the sum of all frequency components per string results in the value 1. This is to ensure that each note of the piano is equally important. However, equal weighting does not necessarily follow from equal importance, since different areas of the piano may emit different amounts of sound energy in average.

De-Noising the Spectrum All spectral components, which are not in the neighborhood of a presumed partial tone are cut off. To do this, the existing spectrum is multiplied by an envelope that has the value range $[0, 1]$. Based on the recorded fundamental frequency and the calculated inharmonicity, it is estimated where partials are located. In order to prevent possible errors, which follow from an inaccurate calculation of the inharmonicity, this envelope has peaks which become broader towards more audible frequencies. Figure 1 illustrates this graphically with an example.

To account for human hearing perception, the spectra are then weighted with the A-weighting curve (Figure 2). In addition, all frequency components that fall below the estimated hearing threshold are immediately set to 0.

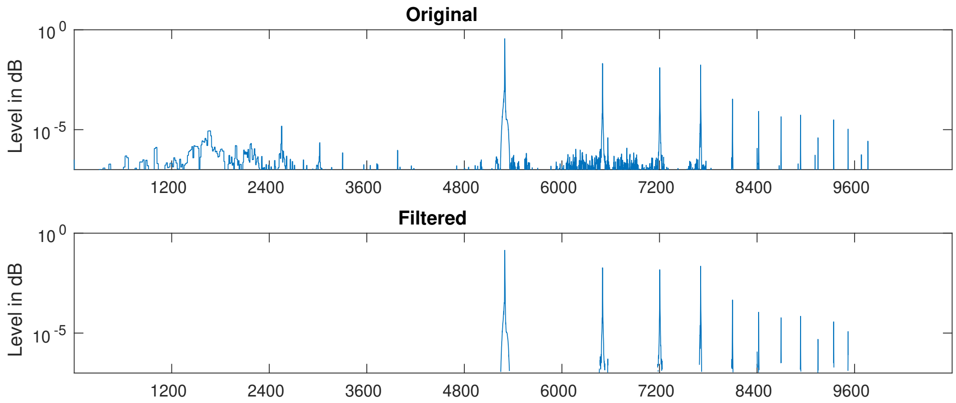

Figure 1: The logarithmic spectrum of an A4 before and after de-noising as described in 2.2.3. The x-axis is showing the frequency in hertz.

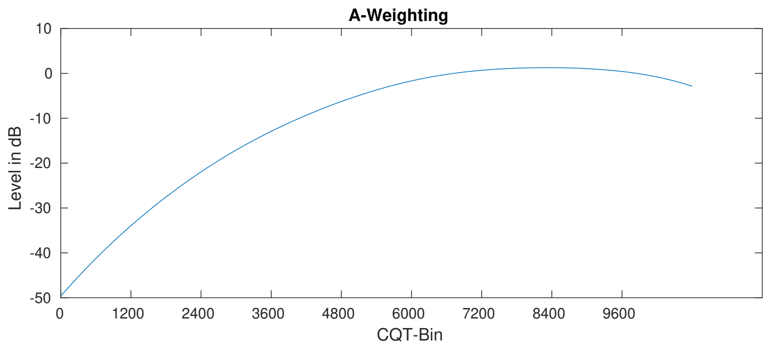

Figure 2: A-Weighting depending on the CQT-frequency-bin (see 2.2.1).

Enhancing the high notes According to the documentation in the source code (Haye Hinrichsen 2018), in order to counteract the weak reproduction of high frequencies above $5000Hz$ from built-in microphones in electronic devices such as notebooks or smartphones, this the high notes are enhanced. Figure 3 shows the effect of this function using two tones as examples. A calculation of this function takes place here only for the tones A4 and higher as an exception.

Mollify For the core tuning algorithm to work correctly, it is necessary that especially the strongest partial tones of different strings related to each other overlap spectrally at the beginning, i.e. after estimating a basic tuning via inharmonicity. Otherwise, one runs the risk that a tone will not settle to its pitch, but to the pitch of one of its tone neighbors. Since pianos differ most from each other in the bass registers due to their size, the estimate obtained is also the coarsest in this register. Moreover, the deviations in the low register are often immense. It is not uncommon for the fundamental frequency of the lowest note A0 of a tuned piano to actually be a semitone lower than it would have to be if it were tuned ignoring the stretching.

To overcome this issue, the spectral peaks are stretched and softened towards low frequencies whereas they are sharpened towards high frequencies. The exact

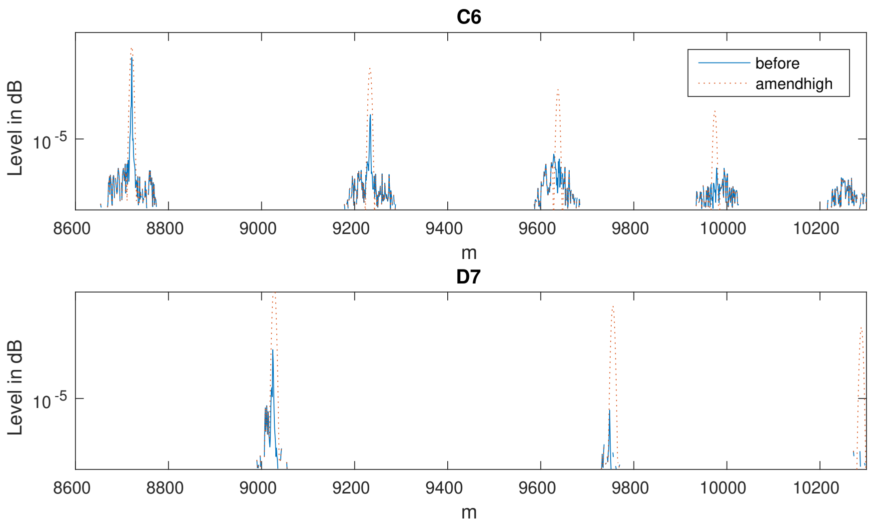

Figure 3: Enhancement of high notes, shown for a C6 and a D7, where m is the index of the CQT-frequency-bin (see 2.2.1).

calculation process of the resulting power spectrum $S_{mollified}$ is given in equation 5 and the range of influence, respectively the slope of the filter Δm depending on CQT-Bin is shown graphically in figure 4.

$$S_{mollified}(m) = \frac{\sum_{m_s=m-\Delta m}^{m+3\Delta m} S(m_s) \cdot e^{-\frac{(m_s-m)^2}{\Delta m^2}}}{\sum_{m_s=m-\Delta m}^{m+3\Delta m} e^{-\frac{(m_s-m)^2}{\Delta m^2}}} \tag{5}$$

$$\Delta m = round\left(1200 \cdot log_2\left(\frac{f+\Delta f}{f_0}\right)\right) - m \tag{6}$$

$$\Delta f = \frac{55}{f} + \frac{f}{2000} \tag{7}$$

$$f = f_0 \cdot 2^{\frac{m}{1200}} \tag{8}$$

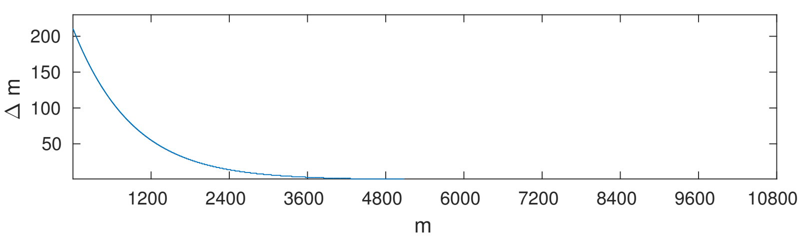

Figure 4: Δm in depending on the CQT-frequency-bin as described in section 2.2.1

What this function will do for two different sine tones is shown in Figure 5.

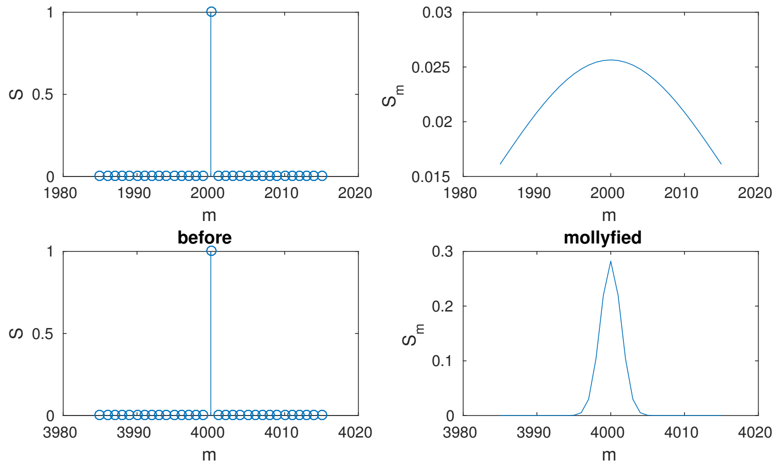

Figure 5: Exemplary demonstration of the algorithm described in section 2.2.3, which broadens spectral peaks towards low frequencies.

2.2.4 Minimizing the entropy

The core of the tuning algorithm is the following loop: Using a modified *Monte Carlo* random algorithm, the power spectrum of one or more strings is randomly retuned via shifting the frequency bins, the power spectra of all strings are summed up and the entropy is calculated. If it is lower than the entropy before, the change is kept, otherwise it is discarded.

The main problem here is that the entropy minimum the algorithm is converging to is only a local one. Due to the preceding basic average tuning however, it is supposed that too bad local minima can be avoided.

2.3 Towards Automatic Piano Tuning

In (J. Tuovinen, Hu, and Välimäki 2019) a novel tuning algorithm has been presented within a fully automatic piano tuning system.

2.3.1 Tuning process

In principle, it imitates the tuning process of the piano, as it is the case when tuning by ear. It tunes the reference A4 first, followed by the A3 and the reference tempered octave from F3 to F4. Afterwards, the high and low registers are tuned via comparing

1. octave
2. octave and double-octave
3. otave, duo-dezima, double-octave, double-octave + quint (and for the treble even double-octave + third and triple-octave)

2.3.2 Determine the quality of the intervals

To take psychoacoustic perception into account, at first A-weighting is always used prior to any other computations. Two different parameters are varied to determine the stretching of the tuning: Which set of intervals were compared and which overlapping partials were included. In their conclusion, they note that the best results, compared to tuning by an experienced aural tuner, were obtained when only the first matching partials of each interval were considered, but different sets of intervals were chosen for the different pitch regions of the piano.

Trying to take all colliding partials into account using a proposed weighting function did not lead to better results. An attempt was made to determine which partials of two tones should overlap more and which should overlap less by first discarding closely spaced partials of two different notes in the calculation if one psychoacoustically masks the other. Then, the ratio of the maximum loudness of each remaining closely spaced pair of partial is used as a weighting factor. The reason that this model could not provide better results most likely is, as the author states as well, that it oversimplifies the perception of beatings. On the one hand, according to Sethares (Sethares 2005), the minimum volume of overlapping partials seems to be the more appropriate measure. Furthermore, it does not include basic characteristics regarding the perception of fluctuation strength and roughness, such as the fact that fluctuation strength typically has a peak at 4 Hz.

3 Goals of this work

At the beginning the motivation of this work was to improve the basic algorithm of the *Entropy Piano Tuner*. It was investigated if by considering that the inharmonicity of a string changes within varying its tension during tuning, the tuning results could be optimized. But this effect is probably not worth considering in a normal tuning process because it is much too insignificant (Joonas Tuovinen 2019).

After I got in touch with several piano technicians and began to better understand the demands they placed on piano tuning software, several reasons spoke against basic characteristics of the *Entropy Piano Tuner*. Recording all piano notes in advance is too time-consuming. The nondeterministic tunings connected with the random local minima of the algorithm are considered problematic, as these results could prioritize different tonalities and lead to less equal temperament. Furthermore, a piano tuner can ideally influence the parameters of a piano tuning algorithm in a way that is tangible for by the methods from aural practice. For example, the widely used proprietary software *Verituner* (Veritune 2011–2021) allows the tuner to set the weights for the beat rates of coincident partials for specific intervals depending on the pitch-region. In addition to these very specific settings, various basic profiles are provided, for example for different degrees of stretching in general. All of these aspects can not be implemented without changing the fundamental design principles of the base algorithm in the *Entropy Piano Tuner*.

Entropy minimization has already been successfully re-implemented using the particle swarm optimization algorithm (Szwajcowski and Pilch 2020), allowing reproducible results to be obtained. In addition, minor improvements were made to the psychoacoustic preprocessing of the spectrum. Nevertheless, this method remains difficult to be parameterized intuitively and the quality of the computed tunings is not yet guaranteed to be sufficient. In addition, this approach is quite computationally demanding.

At the beginning of this work an additional improvement of the tuning was aimed at by optimizing the psychoacoustic preprocessing of the power spectrum. Attempts to tune a single octave by simply shifting the CQT-frequency bins of the preperated spectrum were rather sobering. The stretch of the octaves that were calculated in this way was about twice as large as usual and undoubtedly outside a consonant range. As a result and taking the entire problem into account it seemed reasonable to focus on the essential problem, to find a better understanding of how the stretching of piano tuning is related to sensory consonance, in other words to grasp how the trade-off between the many coinciding partials in two-tone piano intervals may work and maybe is related to already defined psychoacoustic sensations.

It is an unsolvable problem that no matter what fine-tuning is used, all partials of two different piano notes can never be perfectly overlapped. This is also true for consonant intervals, such as the octave. The perceptual sensations leading to sensory dissonance caused by closely aligned partials are called fluctuation strength and roughness (Terhardt 1968) (Zwicker and Fastl 2013). The term fluctuation strength is used to describe beatings from 0 to 20 Hz. Beyond this, the term roughness is used, although the transition between the two is rather fluid.

4 Models

4.1 Roughness

In principle, research by N. Giordano has already shown that the stretching of piano tuning is related to sensory dissonance (Giordano 2015). Giordano's research has been limited to roughness perception alone and has left out fluctuation strength. His work is based on Sethares' roughness model (Sethares 2005) as shown in equation 9. This model describes the perceived roughness of two sinusoids with frequencies f_1 and f_2.

$$R \sim e^{-b_1 s|f_2 - f_1|} - e^{b_2 s|f_2 - f_1|} \tag{9}$$

$$s = \frac{s_0}{s_1 min(f_1, f_2) + s_2} \tag{10}$$

$b_1 = 3.5$, $b_2 = 5.75$, $s_0 = 0.24$, $s_1 = 0.021$, and $s_2 = 19$.

4.1.1 Roughness of two complex tones

For tuning a piano two-note interval we are not interested in the intrinsic roughness of a single tone but rather in the by colliding partials from separate tones. When knowing which interval is compared, it is also known which partials will collide. Using only this partial-pairs for the computation can speed up the algorithm a lot.

4.1.2 Weighting

For computing an overall roughness based on the individual pairs of partials, those will be weighted by their minimal loudness and summed up. The relative perceived loudness l of each partial is approximated via $l \approx 2^{log(S)/10}$ where S is its A-weighted amplitude (Sethares 2005). However, this summation is crucial point since it is not well researched. Sethares suggests that it should be at least quite valid as long the summands are not within the same critical band. For the upper part of the piano the notes have very few partials and this condition is most likely full-filled, but for the bass it isn't. Figure 6 shows how much the count, density and intensity of partials typically varies from the bass to the treble range of the piano.

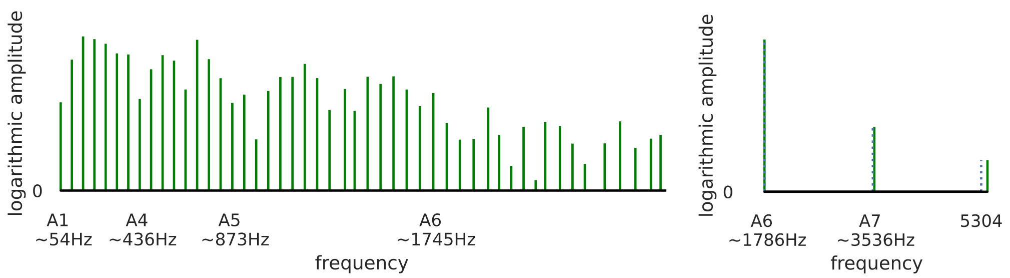

(a) A1 with $f_0 \approx 54Hz$ (b) A6 with $f_0 \approx 1786Hz$

Figure 6: Power spectrum for A1 and A6, representing bass and treble. The dotted lines indicate the frequencies of the harmonics without inharmonicity.

4.1.3 Alternative weighting model

Furthermore, in this project the weighting during the roughness-summation by the minimal loudness of each pair, has been compared to another model proposed by

Vassilakis (Vassilakis and Kendall 2010). Their proposed model is similar regarding the frequency terms but proposes different terms for the amplitude weighting as shown in equation 11.

$$weighting = (A_{min} \cdot A_{max})^{0.1} \left(\frac{2A_{min}}{A_{min} + A_{max}} \right)^{3.11} \tag{11}$$

4.1.4 Secondary beats

Roughness or fluctuation strength occurs not only when two sinusoids sound simultaneously with a frequency differing by a few hertz, but also when one frequency is close to a multiple of the other. Listening tests confirmed that this effect is present and furthermore showed that depending on how many octaves are in between, the strength of the effect decreases (Cong 2016). For this project Congs proposed model is used where the perceived roughness of secondary beatings gets decreases by n, where n is the order, i.e. how many octaves there are between the frequencies.

4.2 Fluctuation strength

For fluctuation strength, in contrast to roughness, hardly any freely available implementations exist. The models that do exist use the time domain and are far more computationally demanding than Sethares' roughness model for sinusoids 9.

This may be one possible explanation why previous studies on the tuning of pianos did not pay much attention to fluctuation strength. However, piano tuning is often about minimizing annoying beatings in the range of just a few Hz. Essentially, the perception of fluctuation strength has a peak around $4Hz$. This fact could be of crucial importance when tuning inharmonic piano intervals.

The peak of fluctuation strength can be roughly modeled, based upon the model of Zwicker and Fastls, as shown in equation 12 (Zwicker and Fastl 2013).

$$F \propto \frac{\Delta L}{\frac{f_{mod}}{4Hz} + \frac{4Hz}{f_{mod}}} \tag{12}$$

The auditory experiments of this model are based on amplitude, as well as on frequency modulated signals, and do not necessarily correspond to the auditory perception of two colliding partials. Nonetheless, there are certain similarities. For the following experiments, the modulation frequency f_{mod} was replaced with the frequency difference of the colliding partials. The modulation depth ΔL, which would be the proper weighting factor, was arbitrary approximated using the minimal loudness, like in the roughness model above.

5 Experiments

5.1 Recordings

For the analytical part of this work samples from all A-notes from a 160 cm long modern grand piano were taken, after it has been tuned by a professional. The samples only use one string per note and were done via a Zoom H1 handy-recorder at 48 kHz with open lid and a distance about 20 to 30 cm to the position, where the hammers hit the string. The velocity the string were hit with was between *forte* and *fortissimo*.

5.2 Preprocessing

For the treble region, up to 50 ms from the attack noise, were cut off, and for the bass up to 100 ms. Then the partials frequencies and their corresponding amplitudes have been extracted using 0.1 to 1 second long windows from the treble to the bass. They were found via peak detection near the estimated frequencies and fine-tuned via quadratic interpolation. The amplitudes have been A-weighted prior to any of the following analyses.

For simulating the tuning, it is assumed that the pitch of the note can easily be shifted applying equation 13 to all partials of a tone, where c is the amount we want to change the pitch in cents and f_n is the partials frequency. This approximation should be good enough, since the inharmonicity and amplitude of the partials do not change too much for such small pitch changes.

$$f_{n,shifted} = f_n \cdot 2^{\frac{c}{1200}} \tag{13}$$

5.3 Analyzing octaves in different pitch ranges

Figure 7 shows the roughness and fluctuation strength progressions for all A-octaves of the piano using the models introduced in section 4. Including the secondary beats sometimes leads to less stretching and sometimes to a wider minimum.

For the lowest octave (Figure 7a), all the models suggest a wider stretching compared to the measured tuning. Tuning the very bass notes in small grands and in small uprights is very problematic due to the relatively short and thick strings and personal preferences there diverge widely. In addition to this the tuning curve shown in figure 8 shows quite an unusual curve for this region. The proposed stretching for the last octave based on the models therefore seems reasonable.

For the octaves A4-A5 and A5-A6 (Figures 7d and 7e) the minima resulting from Vassilakis' model deviate significantly from the existing tuning. Such a deviation in this pitch range leads to a definitely non-fitting tuning.

For the highest octave the minima of all models suggest about 12 cents less stretching (Figure 7g), which would lead to an absolute stretching of only a few cents. As in the bass range, the perception of intervals in the treble range of the piano is very special and a much wider range is usually considered to be acceptable. It is interesting to notice that the used fluctuation model has a local minimum at -12 cents, and in fact at the measured tuning the fluctuation strength is quite similar.

5.4 Stretching

Figure 8 shows the deviation in cents from the fundamental frequency for each A of the equal temperament as tuned by the professional tuner. The other values show the deviation with one after the other tuned octaves (A3 tuned only using the A4, A2 tuned only using the resulting A3 etc.) via the roughness model using the minimal loudness weighing.

6 Discussion

6.1 Tuning octaves

Tuning octaves with the roughness model using the minimum loudness as weighting seems quite promising. Whether the inclusion of secondary beats could be an improvement cannot be conclusively judged, but the wider minimum suggests a strong relation to practical experience, namely that there is a small range that sounds equally good overall. Then characteristics other than fluctuation strength and roughness may come to the fore. Among piano tuners and pianists there is then often talk about *brighter* or *darker* octaves. Furthermore, the use of the proposed weightings is even more crucial in this case, and probably perceived audibility and masking effects should be considered for more appropriate results.

6.2 Whole-scale tuning

This process of tuning a piano has much in common with the intonation process which is common to many other instruments, especially string instruments. Studies on the relation between intonation preference and roughness can be found e.g. in (Villegas et al. 2010). The only difference is that when tuning a piano, it is done beforehand by the piano tuner. However, even if there would exist a precise and stable tool for determining the sensory dissonance of piano intervals, the problem how all the possible intervals in a piano are balanced, still would have to be solved. Undoubtedly, the octave should be one of the least dissonant sounding intervals.

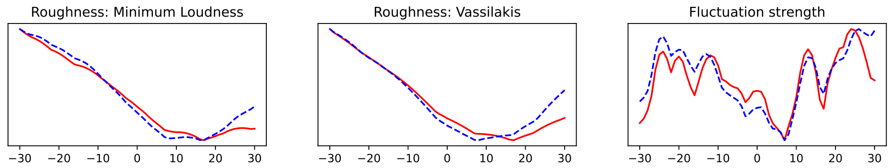

(a) A0 & A1

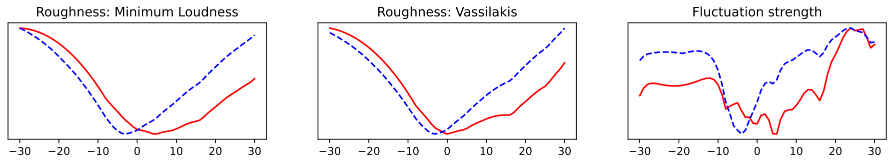

(b) A1 & A2

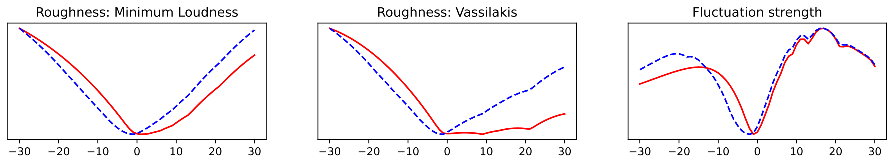

(c) A2 & A3

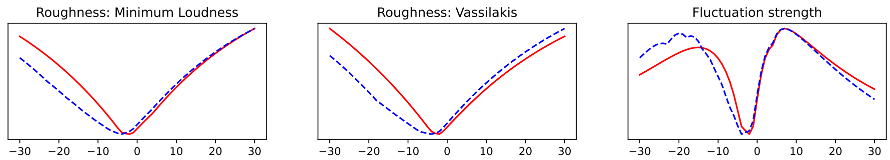

(d) A3 & A4

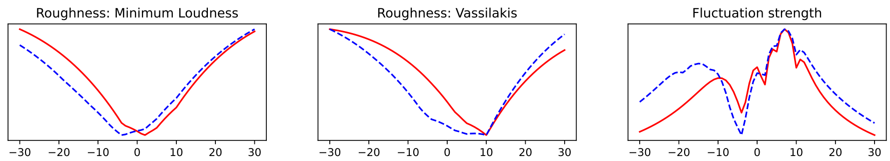

(e) A4 & A5

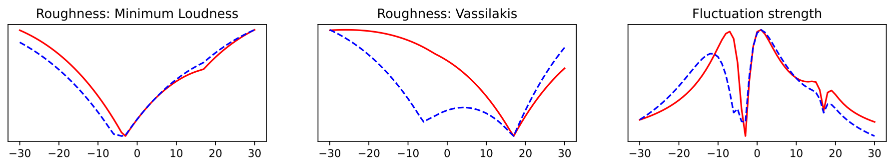

(f) A5 & A6

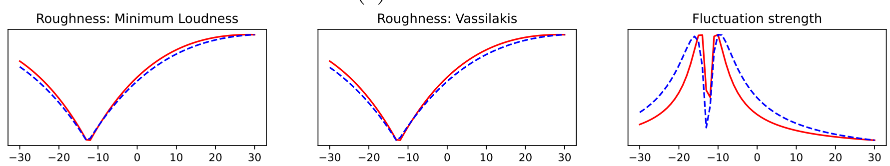

(g) A6 & A7

Figure 7: Roughness and fluctuation strength of several piano octaves for different octave stretchings. The x-axis shows the deviation in cents from the reference professional tuning. Positive values mean greater stretching, vice versa. The y-axis is a normalized linear scale. The dashed lines include the model of secondary beats described in 4.1.4.

But even the octave has to make a compromise with other important consonant intervals, like stated in section 1.2.3.

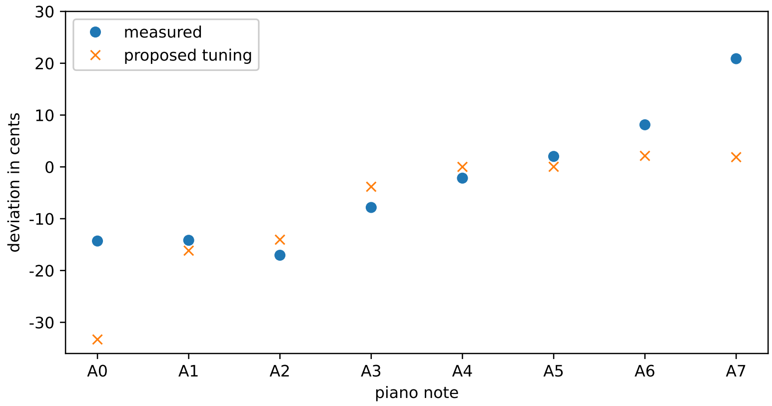

Figure 8: Comparing measured tuning with very simple *one octave after the other* tuning starting with A4 using roughness with minimal loudness weighting.

Another task yet to be examined further is, how the information of the bass-break and the possible rather discontinuous change in inharmonicity and amplitude envelope of the partials can be included. Usually professional tuning software records at least all A-notes, thus recommends some additional notes in the bass break. After calculating a base tuning, the actual tuning is usually performed stepwise from the lowest to the highest note and allows for recalculating the tuning as during the process information of the strings already played can be gained and used. If one considers the previous paragraph, this also makes sense, because when in the lower half of the piano almost exclusively (multi-) octaves are decisive, the pitches of the low A's can already be calculated with sufficient certainty before the actual tuning process begins.

6.3 Future Work

Extracting the exact partial frequencies and proper amplitudes in a robust and efficient way is not trivial, even if the key played is known. The approach used in this work is kept fairly straightforward and though may lead to more or less inaccurate results and may not work at all for other piano sample sources. This is of utmost importance, because the models rely solely on these features.

In the future, it seems most advisable to pay targeted attention to certain regions of the piano, such as the treble or the bass break, and to include a wide variety of samples (concert grand, small grand, small and high upright piano) for this purpose. Research has shown that for example for tuning high piano notes, exactly matching the first (and likely by far most powerful) colliding partials could be the decisive measure (Sneha Shah 2020) and the other collisions of partials could be ignored.

It is still a long way for a free and libre (non-proprietary), robust and practical software for tuning pianos.

References

Cong, Ning (2016). "A Multi-Dimensional Analytical Model for Musical Harmony Perception." PhD thesis. University of York.

Fletcher, Harvey, E Donnell Blackham, and Richard Stratton (1962). "Quality of piano tones". In: *The Journal of the Acoustical Society of America* 34.6, pp. 749–761.

Giordano, N. (Oct. 2015). "Explaining the Railsback stretch in terms of the inharmonicity of piano tones and sensory dissonance". In: *The Journal of the Acoustical Society of America* 138, pp. 2359–2366. DOI: 10.1121/1.4931439.

Haye Hinrichsen, Christoph Wick (2018). *Entropy-Piano-Tuner*. https://gitlab.com/tp3/Entropy-Piano-Tuner/.

Hinrichsen, Haye (Mar. 2012). "Entropy-based Tuning of Musical Instruments". In: *Revista Brasileira de Ensino de Fuedsica* 34. DOI: 10.1590/S1806-11172012000200004.

Rasch, Rudolf A. and Vincent Heetvelt (Dec. 1985). "String Inharmonicity and Piano Tuning". In: *Music Perception* 3.2, pp. 171–189. ISSN: 0730-7829. DOI: 10.2307/40285331. eprint: https://online.ucpress.edu/mp/article-pdf/3/2/171/191145/40285331.pdf. URL: https://doi.org/10.2307/40285331.

Rigaud, François, Bertrand David, and Laurent Daudet (May 2013). "A parametric model and estimation techniques for the inharmonicity and tuning of the piano". In: *The Journal of the Acoustical Society of America* 133, pp. 3107–3118. DOI: 10.1121/1.4799806.

Sethares, William (Jan. 2005). *Tuning, Timbre, Spectrum, Scale*. ISBN: 1852337974.

Sneha Shah, Vesa Välimäki (Mar. 2020). "Automatic Tuning of High Piano Tones". In: *Applied Sciences* 10, p. 1983. DOI: 10.3390/app10061983.

Szwajcowski, Adam and Adam Pilch (Sept. 2020). "Optimization of piano tuning by means of spectral entropy minimization". In: *Applied Acoustics* 166, p. 107359. DOI: 10.1016/j.apacoust.2020.107359.

Terhardt, Ernst (1968). "Acoustic roughness and fluctuation strength". In: *Acta Acustica united with Acustica* 20.4, pp. 215–224.

Tuovinen, J., Jamin Hu, and V. Välimäki (2019). "Toward Automatic Tuning of the Piano". In.

Tuovinen, Joonas (2019). "Signal Processing in a Semi-Automatic Piano Tuning System". English. Master's thesis. Aalto University. School of Electrical Engineering, pp. 51+7. URL: http://urn.fi/URN:NBN:fi:aalto-201903172337.

Vassilakis, Pantelis and Roger Kendall (Feb. 2010). "Psychoacoustic and cognitive aspects of auditory roughness: Definitions, models, and applications". In: vol. 7527, p. 75270. DOI: 10.1117/12.845457.

Veritune, Inc. (2011–2021). *Dave Carpenter, RPT*. https://www.veritune.com/.

Villegas, Julián et al. (May 2010). "Influence of Roughness on Preference of Musical Intonation". In.

Zwicker, Eberhard and Hugo Fastl (2013). *Psychoacoustics - Facts and Models*. Berlin Heidelberg: Springer Science and Business Media. ISBN: 978-3-662-09562-1.