Audio Engineering Society

Convention Paper

Presented at the 128th Convention
2010 May 22–25 London, UK

A Computer-Aided Audio Effect Setup Procedure for Untrained Users

Sebastian Heise[1], Michael Hlatky[1] and Jörn Loviscach[2]

[1] Hochschule Bremen (University of Applied Sciences), 28199 Bremen, Germany
Sebastian@h3e.eu, Hlatky@acm.org

[2] Fachhochschule Bielefeld (University of Applied Sciences), 33602 Bielefeld, Germany
jl@j3l7h.de

ABSTRACT

The number of parameters of modern audio effects easily ranges in the dozens. Expert knowledge is required to understand which parameter change results in a desired effect. Yet, such sound processors are also making their way into consumer products, where they tend to overburden most users. Hence, we propose a procedure to achieve a desired effect without technical expertise based on a black-box genetic optimization strategy: Users are only confronted with a series of comparisons of two differently processed sound examples. Learning from the users' choices, our software optimizes the parameter settings. We conducted a study on hearing-impaired persons without expert knowledge who used the system to adjust a third-octave equalizer and a multiband compressor to improve the intelligibility of a TV set.

1. INTRODUCTION

Highly skilled specialists in the field of signal processing have long since taken over an important role in the audio production process. Since computers have become ubiquitous in recording and mixing music, an ever-growing number of parameters to be set are exposed to the audio engineer. The operator is required to understand the effect of each parameter of a signal processor---be it in hardware or in software. Otherwise, the audio processing task can easily become tedious guesswork.

Professionals in the studio usually have learned to cope with complex technology. However, more and more similarly complex signal processors are finding their way into multimedia consumer products. Here, they are often (believed to be) demanded by the "prosumer", a user which has adequate expert knowledge to cope with the variety of setup options. The greater part of the user group, the "consumers," normally struggle with the number of control options, and often leaves the device's settings unchanged. In the context of audio, however, sound reproduction quality and intelligibility could often benefit from an appropriate setting of a device's built-in signal processors.

Since the Western nations' population is aging, the number of hearing-impaired persons is ever-growing. This group as well uses standard multimedia consumer devices, and often struggle with being able to grasp the content of the played audio. Especially this group of users could greatly benefit from the built-in signal processors, as they provide options to compress the audio output, and boost certain frequency bands of the audio spectrum, thereby improving for instance speech intelligibility. However, a high percentage of this user group often does not manage to fully exploit a device's potential for sophisticated sound quality enhancement. Thus, an easier way to adjust the settings of these devices would be a great improvement.

But even in a wider scope, non-expert users of multimedia technology today are confronted with highly complex topics in the field of audio processing. Complex equalizers and compressors in MP3 players, reverberation settings in home stereos, even standard car radios stride away from the once so simple, one-button-per-function user interaction. Existing approaches to avoid overburdening users may provide presets that can easily be toggled, but their canned settings are far away from a personalized sound setting.

2. PERSONALIZED SOUND SETTINGS

Already very early radio and television receivers provided means of adjusting the perceived audio output not only in terms of the playback volume, but also in tone, often using two simple electronic filters that allowed adjusting the amount of "bass" and "treble". Some consumer audio devices have carried over these well-understood, simple adjustment controls into the digital age. Others, which expose the available power of their underlying technology to the user, offer a vast number of adjustable audio settings: Multiple filters, artificial reverberation, delays and compression, each with a large number of parameters. The manufacturers of such devices, however, often leave it up to the user to gain an understanding of the actual influence of each parameter on the perceived sound. Adjusting for instance a fully parametric equalizer is not a challenging task for a trained audio professional. Already an "obvious" fact may present a huge hurdle to many untrained users: At 0 dB gain setting, the frequency and bandwidth parameters of a filter don't have any effect.

Setting up individual filters, or other effect parameters is therefore often restrained, consumer devices hence provide a number of designed presets, which are usually named using vague or even meaningless descriptions such as *Concert Hall*, *Stadium*, *Jazz*, or *Rock*. Usually none of the presets allows for an adequate enhancement of the sound reproduction quality or the speech intelligibility of the played audio signal. Also today's ubiquitous mobile phones usually offer no way at all of adjusting a caller's audio signal, apart from simple volume settings.

People faced with even slight deafness, which is in general a gradual process starting already at early ages, could benefit tremendously from a simple means of personally adapting the perceived sound of their electronic audio devices in regard to their hearing. Standard hearing aids use multiple high-gain filters and multiband compression to counteract on even extreme hearing losses. The settings of those devices are, however, usually never exposed to their users. The traditional adaptation procedure for a hearing aid comprises that an acoustician adjusts filters and compressors according to a previously inquired personal audiogram of the patient. If the resulting sound is not convenient to the patient, often a laborious procedure of repeatedly tweaking the settings based on the patient's problem description sprouts.

Various strategies for a personalized adaptation of hearing aids have been proposed, most of them based on Lybarger's half-gain rule [1][2], a linear fitting rationale. In his studies, Lybarger provided patients who suffered from sensorineural hearing loss, i.e. deafness due to abnormalities in the hair cells in the cochlea, with hearing aids that were adjusted in order to mirror the individual patient's abnormality in the threshold of hearing. If a patient would have been attested with a hearing loss of 40 dB at a frequency range, the hearing aid would apply a 40 dB gain at this part of the audio spectrum. The patients were asked to adjust the overall volume of the hearing aid until it would sound comfortable to them. The study found that in average patients would reduce the overall gain of their hearing aids by half of the maximum applied gain. At this level, patients reported to be able to catch both soft and loud sounds while still being able to tolerate the output.

McCandless and Lyregaard proposed the Prescription of Gain and Output (POGO) method for adjusting hearing aids [3]. They enhanced the half-gain rule in that they in addition reduce the gain at 250 Hz by 10 dB and 5 dB at

500 Hz, but boost the high-frequency spectrum in order to further improve speech intelligibility. Even less gain is given to the spectrum below 250 Hz in order to reduce the risk of an upwards spread of masking. The method further recommends setting the Maximum Power Output (MPO) of the hearing aid device in order not to excess a patient's Uncomfortable Listening Level (UCL). Due to its design, POGO is limited in its appliance to patients with sensorineural hearing loss.

An enhancement of POGO is the National Acoustic Laboratories of Austria (NAL) linear fitting methods [4], which can be applied to conductive hearing losses, i.e. deafness due to the inability of sound waves propagating through the hearing channel, and sensorineural hearing loss. The NAL methods amplify the spectrum of speech so that all frequencies are perceived equally loud at a comfortable listening level. The advanced NAL-RP fitting method prescribes less high-frequency gain than POGO for steeply sloping hearing losses, and can generally be described by that frequencies at which a severe-profound hearing loss occurred are amplified less than frequencies with better hearing.

Non-linear fitting methods such as NAL-NL1 [4][5] take also into account that the threshold of pain—the UCL—does not shift upwards with the Hearing Threshold (HTL) for a hearing-impaired patient, it rather stays the same as for a healthy person. Therefore, in order to not harm a patient by applying high gain filters, the dynamic range of the auditory environment should be reduced in order to fit between the HTL and the UCL. Kates [6] conducted studies to find appropriate temporal compression settings for hearing aids. A fast temporal response of the compressor allows for rendering all portions of the audio signal, whereas slower temporal response times preserve the overall speech envelope modulation for the listener. Using a non-linear fitting method such as NAL-NL1, the intelligibility of soft sounds as well as rather loud sounds can be enhanced, in contrary to linear fitting methods that can be adapted to only one excitation level.

Apart from NAL-NL1, a fitting method named DSL (Desired Sensation Level) [7][8] developed at the University of Western Ontario specifically for pediatric patients are the most popular fitting rationales used by acousticians today. With the DSL fitting method, the holistic audio spectrum, not only the limited speech spectrum is equalized to be perceived at equal loudness. The goal of this method is to allow especially hearing-impaired children to learn about their whole auditory environment, not only mere speech.

All personalized fitting procedures mentioned above have in common that adjustments to the sound of a hearing aid can only be carried out from an acoustician. If a patient is not pleased with the currently perceived sound, he first has to report to an acoustician about the problem. The acoustician then needs to translate the patient's problem description into an appropriate adjustment of the filters and the compressors, which to no surprise is prone to errors.

3. PERSONALIZING SOUND SETTINGS USING OPTIMIZATION STRATEGIES

It is important to state that the patients' acceptance rates for wearing hearing aids are greatly linked to the initial experience with the device. If a hearing aid does not improve the intelligibility of the perceived sound or even diminishes it after an initial setup procedure, it is most likely that a patient will not carry on wearing it. In contrast to visual aids, hearing aids usually do not improve a patient's hearing with the same initial benefit. This may partially be related to the imprecision of the audiograms determined by inquiry of the patients. Audiograms are compiled by testing a patient's hearing threshold at various frequencies. Especially elderly patients often report that they have problems exactly identifying their hearing threshold for the individual frequencies.

To improve the hearing aid fitting process for patients who lack knowledge about the effect of the individual parameters on the perceived sound, the scientific community has proposed various systems that enable a user-guided setup procedure, leveraging neural networks and genetic optimization algorithms, or presenting the user simplistic and easy-to-grasp user interfaces. In almost every case, patients rated settings generated on their own as sounding better than a personalized setting produced by an acoustician.

Wessel et al. [9] proposed a system that allows a user to find appropriate compressor settings for music stimuli. In their software, a user selects settings by pointing at a position on a two-dimensional map. The two-dimensional map representing similar sounding settings of the compressor in close-by areas is generated in a previous step using a neural network trained with user-generated dissimilarity ratings of sound snippets.

Kuk and Pape [10] utilized a downhill-simplex optimization method for hearing aid fitting. On average, their setup procedure takes 2.5 hours. The user test suggested that patients would choose similar settings each time they underwent the optimization process. That indicates a certain degree of reliability of the method. Takagi [11] proposed a system which uses an evolutionary optimization algorithm for hearing aid fitting. His system presented the patient a 4x4 matrix of different adjustments on a computer screen. Each setting could be rated by the user as “good” or “bad”. The evolutionary algorithm updated the population formed by the settings each time a user had voted on all presented settings.

Durant [12][13] conducted several experiments to examine the performance of genetic algorithms in user-guided hearing aid fitting tasks. He found that the resulting setting chosen by the patients did not depend on the initial settings.

The more similarthe initial and the resulting settings were, the faster the patients carried out the setup procedure. In pursuit of a simple user interface for his tests, he created a proprietary hand-held device that supported a paired comparison of individual settings.

Zakis et al. [14] conducted tests with a hearing aid that provided a user-trainable compression algorithm. Users trained the system over the course of several weeks in their daily environment using only a single rotary knob. The user could adjust three parameters of a three-band compressor: the overall volume, a parameter called bump, which controlled the relation of the mid-frequencies band’s gain to the high and the low frequencies bands’ gains, and a parameter called slope, which controlled the relation between the gain of the low band and the gain of the high band. All other parameters influencing the compressors were derived from the user’s adjustments. The method proved reliable. Users were confused, however, concerning which of the three parameters the knob currently controlled.

Baskent et al. [15]**Fehler! Verweisquelle konnte nicht gefunden werden.** examined in how far genetic algorithms can be applied in audio control operations. In a user test, they asked subjects to improve the intelligibility of distorted speech signals using a vocoder effect. The setting of the vocoder was controlled through a genetic optimization algorithm. By rating different settings, users were able to produce solutions with enhanced intelligibility.

Siemens AG introduced a line of user-trainable hearing aids [16], which utilize a proprietary training method called “SoundLearning”. The system provides the user with two knobs, one of which controls the overall gain and the other controls the gain for high frequencies. The system adjusts compression and filter settings by analyzing in which situations a user attenuated the high-frequencies or the overall volume. The adaptation process of the hearing aid takes several weeks, however.

In a comparative study, Eiler et al. [17] argued that for a successful implementation of optimization strategies in the adjustment of hearing aids, the setup time has to be reduced drastically. They computed that the average setup time in the optimization strategies proposed by the scientific community is about 18 minutes, insufficient duration that seems to be too long for an application in a busy clinical context. In order for patients to be able to carry out the optimization process on their own, they state that the design and the usability of the interfaces employed in the optimization process would have to be improved.

In an earlier work [18], one of us showed that using a simplified user interface, the setup time for a parametric five-band equalizer can be reduced. In the interface, users can simply draw a desired transfer function. An evolutionary optimization strategy adjusts an automatically chosen number of biquad filters so that the overall transfer function resembles the hand-drawn user input as closely as possible.

To adjust the settings of a ten-band equalizer Mecklenburg and Loviscach [19] used words that describe the perceived timbre such as “boomy”, “boxy”, “warm” and “nasal”, They employed a two-dimensional self-organized map to lay out the different equalizer settings in accordance to their related words. With this system users can adjust the equalizer by pointing at a position on the map.

In another earlier work [20], we compared the performance of different optimization strategies in the context of automatically setting the parameters of generic audio effects. We found that the time needed for an automatic adjustment is not significantly different among a generic evolutionary optimization algorithm, a Particle Swarm Optimization (PSO) [21], and the downhill simplex optimization method known as Nelder-Mead Algorithm [22], no measureable differences occur in terms of the length of time needed for the adjustment. We tested the algorithms in automatically adjusting a synthetic reverberation plug-in in order to resemble recorded impulse responses of existing rooms, and on automatically adjusting standard synthesizer plug-ins in order to resemble a recording of a natural sound as close as possible [23]. Our software rated the settings of the plug-in's parameters by comparing the output and the reference sounds using psychoacoustic measures. We tested the settings found by the optimization process against expert-created settings and found in double-blind user tests that the machine-crafted settings were perceived as more similar.

4. SYSTEM DESIGN

Attempting to develop a system that allows for a user-defined, quick adjustment of complex audio processors without exposing the actual parameters to the user, we employ a Nelder-Mead optimization strategy in combination with a simplistic graphical user interface that enables a paired comparison of sound examples representing two different sets of settings of the audio processors. The system basically consists of three different parts: audio processor, optimization engine and user interface.

4.1. Audio Effect Processor

We have implemented the audio effect processor as a VST plug-in host. In this way, we are able to choose from a large variety of different generic audio effects that are freely available and well-known to many audio professionals. In addition, these effects tend to provide a large number of parameters that are not commonly understood by untrained users.

4.2. Optimization Engine

We decided to employ a Nelder-Mead downhill simplex optimization strategy in order to achieve a quick process while at the same time confronting the user with as few as possible "random" settings. Generic evolutionary optimization algorithms and also less-controlled approaches such as the PSO algorithm guarantee to extensively probe the available parameter space for the problem solution. To this end they make heavy use of randomness in selecting new parameter settings. These approaches usually lead to good results and tend to not get stuck in a local minimum, which looks beneficial for the problem at hand. These methods are, however, frustrating and also time-consuming due to the randomness: Many of the solutions suggested by the algorithm may not be suitable at all, but must nonetheless be listened to and rated by the user. The highly controlled Nelder-Mead optimization strategy allows for a quick solution of an optimization problem; it is, however, often prone to get stuck in a local minimum. The possibility for this can, however, be reduced by initializing the algorithm with a good estimation of the expected result. In our design, we opted for randomized initial settings, as the estimation would for instance have to be derived from a personal audiogram of each listener, which requires an examination with special equipment.

In our system, each input of the user is treated as an evaluation of the fitness function. For each iteration of the optimization process, the *worst* setting from the current group of settings has to be identified.

4.3. User Interface

The interface for the optimization procedure has been implemented in three versions: as a native .NET application, as AJAX Web page for mobile devices such as the Microsoft Zune and as native application for the Apple iPod Touch. All user interfaces provide the same functionality and use the same elements. The user is presented with two buttons in order to perform the paired comparison between the two presented examples. We decided to reflect the individual settings of the two examples in a procedural design of the buttons; buttons that represent similar settings have a similar appearance.

Several methods exist to perform an acceptance gesture—for instance a double-click,a wipe-gesture on a touch-enabled device, or also dragging the icon representing the chosen setting onto a drop-zone. Nonetheless, we elected to provide the user with a third button that allows for a distinct accepting of the currently selected setting. In our user interface design, the third button appears only after the user has listened at least once to each of both sound settings.

The procedural button images representing the sound settings are inspired by a work of Kolhoff et al. [49], who presented content-based icons for audio files. In our implementation, the parameters of each effect setting are mapped to the colors and shapes of two hypertrochoid curves, resulting in flower-like forms (see Fig. 1).

The background of the currently selected button is highlighted to provide constant feedback. On selecting one of the buttons, a smooth animation of the button's size further indicates the selection. After accepting one of the settings as better-sounding by pressing the third button, all buttons disappear and the two main buttons containing two newly generated procedural images fade in with a smooth scaling effect. The procedural design of the buttons possesses an additional benefit: Even though two sounds in a successive paired comparison may not be perceived as sounding different, the button's appearance will always change at least slightly as long as there is a difference in the parameters. Thus, the user is assured that the system continues the optimization, although he or she may not have perceived a change and immediate success yet.

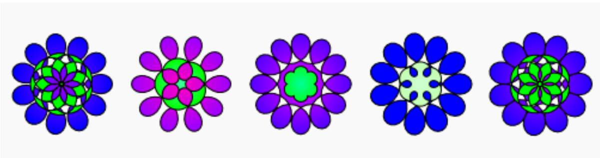

Fig. 1: The flower-shaped procedural button images are computed from the individual audio effect parameters. Similar effect settings result in a similar appearance of the images.

Fig 2: The user interface containing two flower-shaped, procedural buttons, running as native application on the Apple iPod Touch, and as an AJAX Web page on the Microsoft Zune.

5. USER TEST

In order to evaluate the performance of our system, we conducted a user test, in which two VST plug-ins were used to resemble the digital signal processing chain commonly found in modern hearing aids: equalization and multi-band compression.

5.1. Test Subjects

In order to test users with no possible expert knowledge, we decided to select a test panel of seniors. Some of the tested subjects further suffered from severe hearing losses. Altogether we tested twelve participants (six female, six male). The mean age of the test subjects was 71, the maximum age 80, the minimum age 34. All test subjects are seasoned members of the listening panel of the Hörzentrum Oldenburg[1], where they participate regularly in testing new hearing aid technology. Even though they did not receive training concerning signal processing, they are well trained in assessing sound reproduction quality. The audiograms of all test subjects were available to us for evaluation.

5.2. Software Setup

The test participants were asked to adjust the parameters of two chained VST plug-ins: a three-band compressor by "slim slow slider"[2] that provides 26 parameters in total, and an equalizer by Karma FX[3] that provides 31 filter bands at constant frequencies and bandwidth at a maximum gain of 24 dB per band.

As finding appropriate settings for each parameter of the two processors via the optimization would have been too time-consuming, we opted for control only a limited, yet effective set of parameters via the optimization process. We control six parameters of the compressor, namely the ratio and the threshold setting for each of the three bands. The crossover frequencies of the bands were set in advance to 670 Hz and 4700 Hz. The attack time for each band was set to 10.0 ms, the release time for each band to 200.0 ms. All gain settings for the bands were set to neutral. The compression knee for each band was set to 75% hardness.

We controlled the 31-band equalizer with only six parameters. This was not accomplished by dividing the 31 bands into six blocks, as this would introduce jumps into the overall transfer function of the equalizer between the parts of the spectrum associated with each block. Instead, we controlled the equalizer using the first six power bands of the Fourier-transformed gain settings of all filter bands. That is: The first parameter controls the gain of all filter bands; the second parameter controls the gradual ratio in gain between the high and the low filter bands; the third parameter controls the

[1] http://www.hoerzentrum-oldenburg.de
[2] http://www.geocities.jp/webmaster_of_sss/vst/index.html
[3] http://karmafx.net

gradual ratio in gain between the middle against the high and the low frequency bands; and so on. Using this technique, even complex settings of the equalizer can be described by only six factors.

5.3. Hardware Setup

The optimization process was fully controlled by the user and solely through the application running on an Apple iPod Touch. All signal processing was carried out on a laptop computer that was linked wirelessly to the iPod Touch. All communication between the laptop computer and the iPod was carried out via simple HTTP requests. The user listened to the test stimuli via Sennheiser HDA-200 headphones, which allow for extreme amplifications particularly for applications in audiology, and are equalized so that they provide a near-flat frequency response.

5.4. Test Stimuli

In order to simulate a everyday listening behavior, two real-life test stimuli were chosen: a 15-minute audio recording of the most-watched German TV news program “Tagesschau”[4], and a 15-minute excerpt of a piano recording of Chopin’s sonata No. 3, Allegro Maestoso, performed by Leone Magiera in 1993.

5.5. Test Procedure

The first part of the test consisted in an oral description of the user interface and instructions on the test procedure; in addition, each subject could familiarize him- or herself with the iPod touch control in a not-recorded test run. Due to the age of the test subjects, we set no strict time frame of the instruction phase, as we did not want to introduce additional pressure into the test situation. The instructions were, if necessary, repeated until each user operated the user interface properly. All users achieved this after five to eight minutes.

In the second part of the user test, the subjects adjusted the sound of the processing setup starting from randomly initialized settings while listening to the recording of the Tagesschau news presenters. After that, the users were verbally questioned about the perceived quality and intelligibility of the speech signal. In the third part of the user test, the subjects again adjusted the processor settings, this time listening to the 15-minute recording of the classical piano piece.

Afterwards, the subjects were asked about the perceived quality of the music playback. Subjects were further asked to rate the overall sound quality when compared to unaided listening, and to rate the quality of the resulting sound when compared to their own, expert-adapted hearing aids. Subjects were further questioned about their ability to make a choice between the two presented sound setting examples during the course of the optimization, and about their general opinion on the usability of the user interface. The answers were allowed to range in seven steps from “very bad” to “very good”; they were collected orally by the test instructor.

6. RESULTS

Immediately after the first test run, all participants reported that they “liked” the system’s user interface and feel. In an informally conducted blind test after the actual user test, eleven out of the twelve test subjects were able to pick their crafted setting against a linear setup.

6.1. User Interface Performance

During both recorded optimization sessions, all users showed an almost constant toggle and decision rate. Mostly, users toggled 25 to 30 times per minute between the two presented sound settings via the procedural buttons and made six to eight decisions per minute concerning which of the examples sounded better than the other. In less than 2% of all actions, the users tried to advance by accepting one sound although they did not yet evaluate the other example setting. In most cases, participants reported that they could not perceive any noticeable difference between the single settings after seven to ten minutes into the optimization session, which suggests that the optimization process already converged after this time.

6.2. Subjective Evaluation

Fig. 3. shows a box-and-whisker diagram of the evaluation of the user questionnaire. Both the perceived playback quality for speech and music are rated

[4] http://www.tagesschau.de

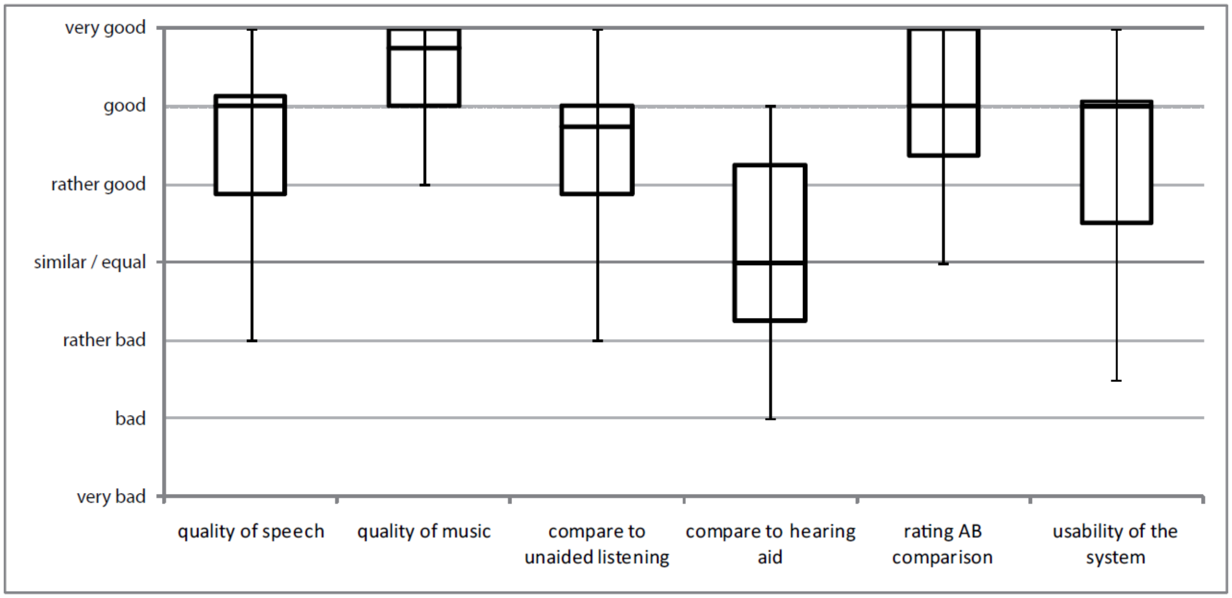

Fig. 3: Box-and-whisker plot showing the results of the user questionnaire.

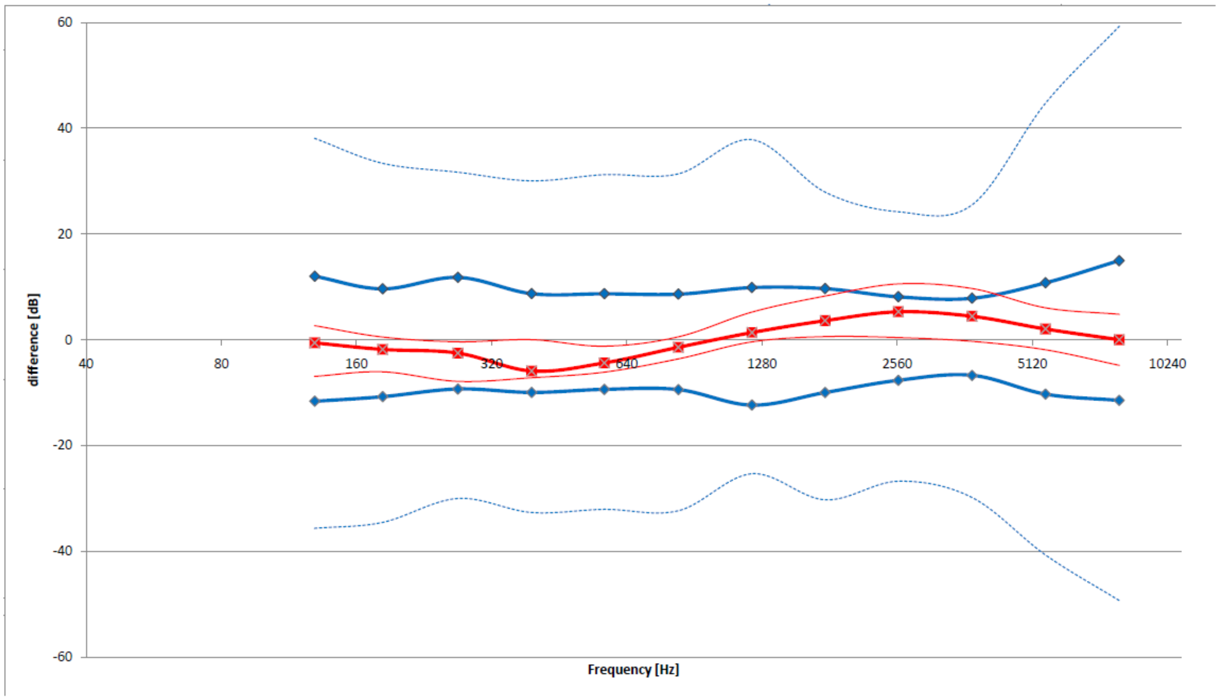

Fig. 4: Deviation of the resulting frequency responses before and after optimization. The upmost and lowest dotted lines show the maximum difference between the randomized initial settings in the frequency spectrum. The lines marked with diamonds denote the 25% and 75% percentiles of the average deviation between the randomly initialized settings. The inner thin lines show the 25% and 75% percentile of the average difference between the optimized settings. The line marked with squares shows the median of the frequency response of all optimized settings.

on average between very good and good. In addition, the average perceived quality after the optimization in comparison to a playback with zero settings is rated good on average. Generally, the test subjects rated the perceived sound quality as similar to their own hearing aids. On average, the test subjects rated their ability to distinguish between the two presented sounds as good, with a leaning towards very good, and the overall usability of the system as good, with a leaning to rather good.

6.3. Objective Evaluation

In general, the subjects did not reach sound settings that compensated their hearing loss as indicated by their personal audiogram. Some participants drove the optimization to a nearly perfect mirror image of the frequency response curve in their audiogram. Others, however, even decreased the parts of the frequency spectrum in which they suffer from hearing loss. The correlation coefficient between the individual subject's optimization settings from the speech and the music run is 0.4. These observations suggest that the system does not provide a means of repeatedly setting a desired frequency and compression response starting out the optimization from randomly initialized settings. Fig. 4, however, shows that the differences in the individual random start settings for the two sessions are significantly higher than the differences in the settings after the user test. This indicates that the result of the optimization process does not entirely occur by accident.

7. CONCLUSION

We have created a system that allows users without expert knowledge to control sound effect parameters, such as those of equalizers and compressors, for. In a series of paired comparisons, users decide which setting they prefer. Software running in the background continuously optimizes the overall best setting dependent on the choices made. An experiment with elderly, hearing-impaired users confirmed the general applicability of such a system.

The resulting sound settings, however, revealed significant contradictions to established hearing aid fitting rationales for hearing-impaired patients. The test subjects nevertheless reported a rather high liking for the resulting sound quality and intelligibility and rated the system very usable. This indicates that the system may also be a valuable enhancement for standard multimedia gear, allowing non-expert users to personalize the sound settings.

8. REFERENCES

[1] Staab, J. W.: "Hearing Aid Selection: An Overview" in Sandlin, R, E. (ed).: "Hearing Aid Amplification", Singular Publishing Group, pp. 55-134, San Diego, US (2000).

[2] Lybarrger, F. S., and Lybarger, H. E.: "A Historical Overview" in Sandlin, R, E. (ed).: "Hearing Aid Amplification", Singular Publishing Group, pp. 1-33, San Diego, US (2000).

[3] McCandless, A. G., and Lyregaard, P. E.: "Prescription of Gain/Output (POGO) for hearing aids.", Hearing Instruments, Vol. 34, No. 1, pp. 16-21 (1983).

[4] Venema, H. T.: "The NAL-NL1 Fitting Method", Audiology Online (2001), http://www.audiologyonline.com/articles/pf_article_detail.asp?article_id=253 (last access: 12.03.2010).

[5] Dillon, H.: "NAL-NL1: A New Procedure for Fitting Non-Linear Hearing Aids", The Hearing Journal, Vol. 52 No. 4 pp. 10-16 (1999).

[6] Kates, M. J.: "Digital Hearing Aids", Plural Publishing, San Diego, CA (2008).

[7] Beck, L. D., Moodie, S., and Speidel, D.: "Pediatric Hearing Aid Fittings and DSL v5.0", The Hearing Journal, Vol. 60, No. 6, pp. 54-58 (2007).

[8] Scollie, S., Seewald, R., Cornelisse, L., Moodie, S., Bagatto, M., Laurnagaray, D., Beaulac, S., and Pumford, J.: "The Desired Sensation Level Multistage Input/Output Algorithm", Trends in Amplification, Vol. 9, No. 4, pp. 159-197 (2005).

[9] Wessel, D., Fitz, K., Batternberg, E., Schmeder, A. and Edwards, B.: "Optimizing Hearing Aids for Music Listening", 19th International Congress On Acoustics (2007).

[10] Kuk, K. F., and Pape, C. M. N.: "The Reliability of a Modified Simplex Procedure in Hearing Aid Frequency-Response Selection", Journal of Speech and

Hearing Research, Vol. 35, No. 12, pp. 418-429 (1992).

[11] Takagi, H., and Ohsaki, M.: "IEC-based Hearing Aid Fitting", IEEE SMC '99 Conference Proceedings, Vol. 3, pp. 657-662 (1999).

[12] Durant, A. E.: "Hearing Aid Fitting with Genetic Algorithms", PhD Thesis, The University of Michigan (2002).

[13] Durant, A. E., Wakefield, H. G., Van Tasell, J. D., and Rickert, E. M.: "Efficient Perceptual Tuning of Hearing Aids With Genetic Algorithms", IEEE Transaction on speech and audio processing, Vol. 12, No. 2, pp. 144-155 (2004).

[14] Zakis, A. J., Dillon, H., and McDermott, J. H.: "The Design and Evaluation of a Hearing Aid with Trainable Amplification Parameters", Ear and Hearing, Vol. 28, No. 6, pp. 812-830 (2007).

[15] Baskent, D., Eiler, C., and Edwards, B.: "Using Genetic Algorithms with Subjective Input from Human Subjects: Implications for Fitting Hearing Aids and Cochlear Implants", Ear and Hearing, Vol. 28, No. 3, pp. 370-380 (2007).

[16] Chalupper, J., Junius, D., and Powers, T.: "Algorithm lets Users Train Aid to Optimize Compression, Frequency Shape, and Gain", The Hearing Journal, Vol. 62, No. 8 pp. 26-30 (2009).

[17] Eiler, C., Baskent, D., Recker, K., and Edwards, B.: "Genetic Algorithms: Are they the Future of Hearing Aid Fittings?", The Hearing Journal, Vol. 61 No. 12, pp16-19, (2008).

[18] Loviscach, J.: "Graphical Control of a Parametric Equalizer", AES Paper 7437, AES 124th Convention (2008).

[19] Mecklenburg, S., and Loviscach, J.: "SubjEQt: Controlling an Equalizer through Subjective Terms", ACM SIGCHI 2006, pp. 1109-1114, (2006).

[20] Heise, S., Hlatky, M., and Loviscach, J.: "Automatic Adjustment of Off-the-Shelf Reverberation Effects", AES Paper 7758, 126th AES Convention (2009).

[21] Kennedy, J., and Eberhart, R.: "Particle Swarm Optimization", Presented at the IEEE Conference on Neural Networks, pp 1942–1948, (1995).

[22] Nelder, A. J., and Mead, R.: "A Simplex Method for Function Minimization", Computer Journal, Vol. 7, No. 4 pp. 308-313 (1965).

[23] Heise, S., Hlatky, M., and Loviscach, J,: "Automated Cloning of Recorded Sounds by Software Synthesizers", AES Paper 7858, 127th AES Convention,(2009).