Hearing
Technologies

Why linear processing is insufficient to compensate for hearing loss

Mimi Hearing Technologies

Research Study

By Mimi Hearing Technologies GmbH, Berlin, Germany
June 10th, 2024

1. Hearing Loss is non-linear

Many sound personalization solutions on the market like JBL's Personi-Fi, Audiodo, Sonarworks and of course Mimi offer hearing loss compensation. However, in most solutions, hearing impairment is treated in a very simplistic manner. The audiogram serves as a tool to identify frequency regions with raised pure tone thresholds (where a higher signal level is needed for the user to be able to hear a tone). The required gain in each frequency band is then applied statically to the signal to compensate for the hearing loss resulting in constant signal coloration. Early hearing aids operated in similar ways, and while this approach does benefit the hearing-impaired listener, it does not sufficiently account for the complexity of auditory processing and the nature of hearing loss.

The concept of loudness is especially relevant in media consumption. It describes the subjective intensity of a sound as perceived by a listener within the mix. In hearing impaired listeners, loudness perception can change drastically, particularly in frequency regions with significantly elevated thresholds. In solutions with linear processing positive or negative gain is applied to the signal regardless of its temporal and dynamic behavior. This is identical to how an EQ functions and in the following these EQ-like, linear solutions are referred to as "gain-only". However, simply amplifying the sound level in frequency regions where listeners experience hearing difficulties does not fully address the problems resulting from the changes in loudness perception. This is due to the phenomenon of abnormal loudness growth, where the discrepancy in loudness perception varies across signal levels for a given frequency. While quieter signal components can be lost for hearing impaired listeners, louder signal components can be perceived as loud compared to how a healthy listener would perceive them (Figure 1).
Due to the linear nature of their processing solutions like Audiodo's or Sonarworks' can only lift all signal components equally as shown in the black line (Figure 2, left). This leads to several problems:

- Loud components can quickly become too loud and even harmful to the listener (see [4])
- Due to the limited available headroom
    - loud signal components either prevent a thorough restoration of quiet signal components
    - Or clipping occurs much more often (see [3])
- Low level signal components can be masked by intense signal components that are amplified more than necessary

In the context of music, which is often highly dynamic with a range of low and high intensity components, applying non-linear, compressive dynamics processing is essential to account for the perceptual effects of hearing loss. To fully compensate for the abnormalities in loudness perception in hearing-impaired listeners, signal components with various levels of intensity must be amplified differently to achieve the desired target function (Figures 1 and 2). This approach ensures that the level of both quiet and loud elements of the music is adjusted appropriately, providing a more accurate and satisfying listening experience for people with hearing impairment.

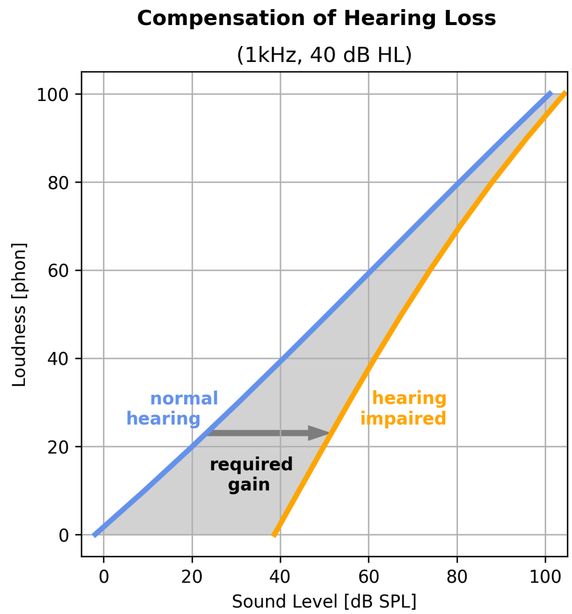

Figure 1: Normal hearing and hearing impaired loudness perception as a function of sound level at 1kHz

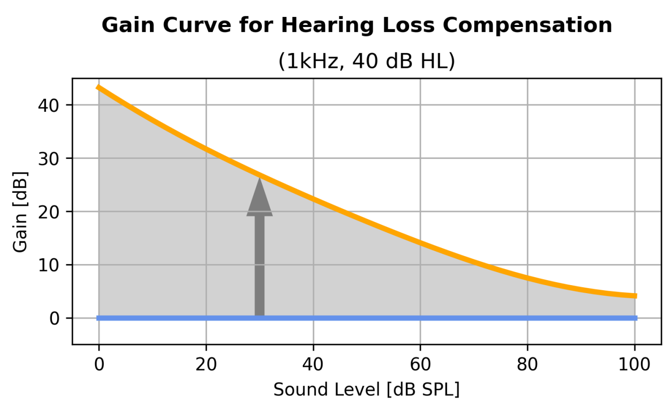

Figure 2: Gain curve for Hearing Loss Compensation at 1kHz with 40dB HL hearing loss as a function of sound level

2. Non-linear processing can compensate more hearing loss than EQ solutions

The amount of gain that can be added to a signal without running into clipping and distortions is limited if all signal intensity levels are treated equally. In [1] it was explained that more gain is needed for quiet signal components than for louder components.

Mimi's dynamic range multiband compressor attenuates loud signal components based on their intensity level, creating the foundation for restoring the intended loudness of subtle details. Following this, static gain is applied to ensure that both loud and quiet elements of the audio are appropriately balanced and audible to the listener.

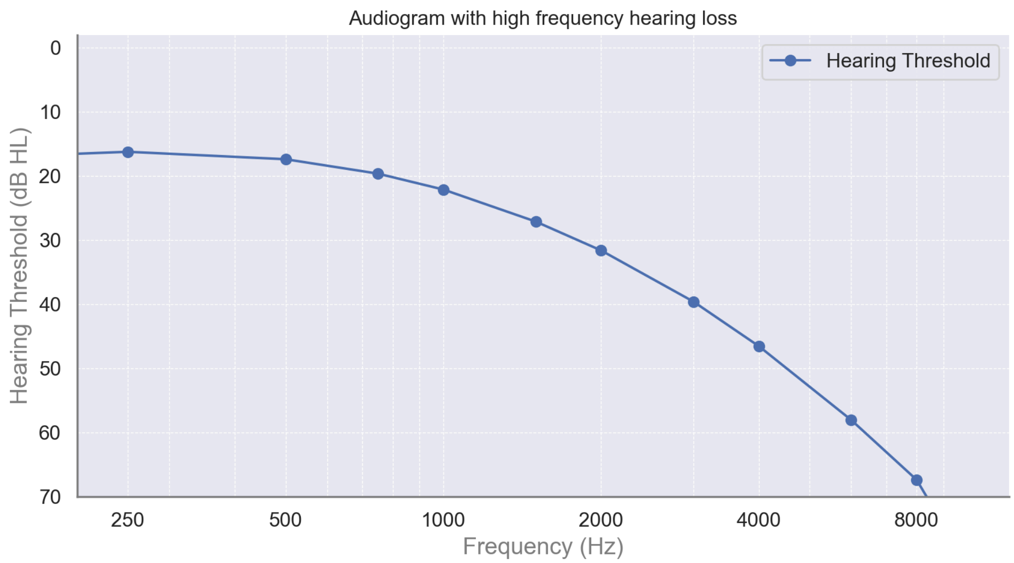

Figure 3: Exemplary high frequency hearing loss

To demonstrate and quantify the difference between Mimi's solution and a linear gain-only solution, the following exemplary analysis was done for one hearing loss profile (Figure 3). Also see Figure 4 for a visual representation of the preparation stage.

Preparation Stage:

- A hearing loss profile for high frequency hearing loss was defined (Figure 3)
- Mimi's loudness restoration algorithm was used to calculate a fitting
- The fitting profile's equivalent continuous sound level (LEQ) was analyzed with a 3rd octave filterbank (FIgure 5) and mimicked with a gain-only stage, using programme simulation noise (BS EN 50332)

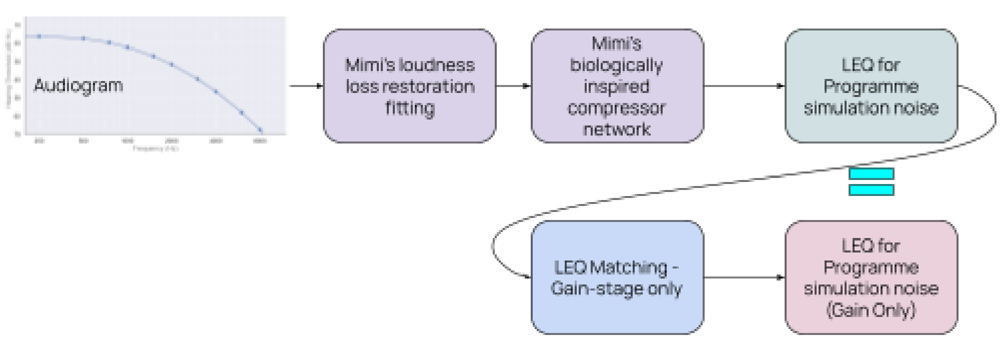

Figure 4: Preparation stage

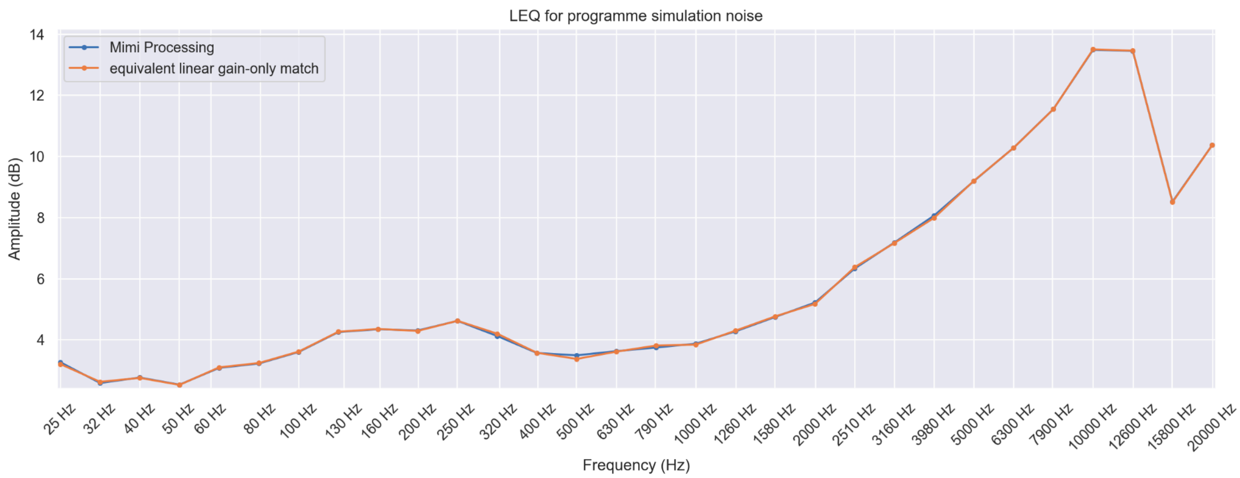

Figure 5: LEQ boost profile for Mimi processing, linear gain-only processing for programme simulation noise

Analysis Stage (Figure 6)

- A loudness loss model[1] which characterizes the impact of hearing loss (defined in terms of absolute threshold) on loudness perception was used to calculate loudness levels in phon (L_N). This was done across 17 frequency bands (½ octave filterbank ranging from 90Hz -20kHz) while the signal was split into 100ms time frames (0% overlap).
- The loudness model was used to model loudness estimates of three different configurations that could then be compared:
  - The reference loudness derived from a model of normal hearing
  - Hearing impaired loudness derived from a model of hearing impairment
  - Pre-processed hearing impaired loudness derived from the same model of hearing impairment, but with the audio input augmented.
- Restoration was quantified using granular frame-wise estimations of loudness distances, in phon, between a normal-hearing person, a hearing-impaired person, and a hearing-impaired person with pre-processed audio. Specifically, the restoration metric was based on the percentage of these loudness distances between a normal-hearing listener and a hearing-impaired listener that could be reduced or eliminated through pre-processing.
  - A corpus of 50 songs from varying genres were used (rock, pop, jazz, classical music, electronic music)

[1] Based on human perceptual data published in https://pubmed.ncbi.nlm.nih.gov/25920842/

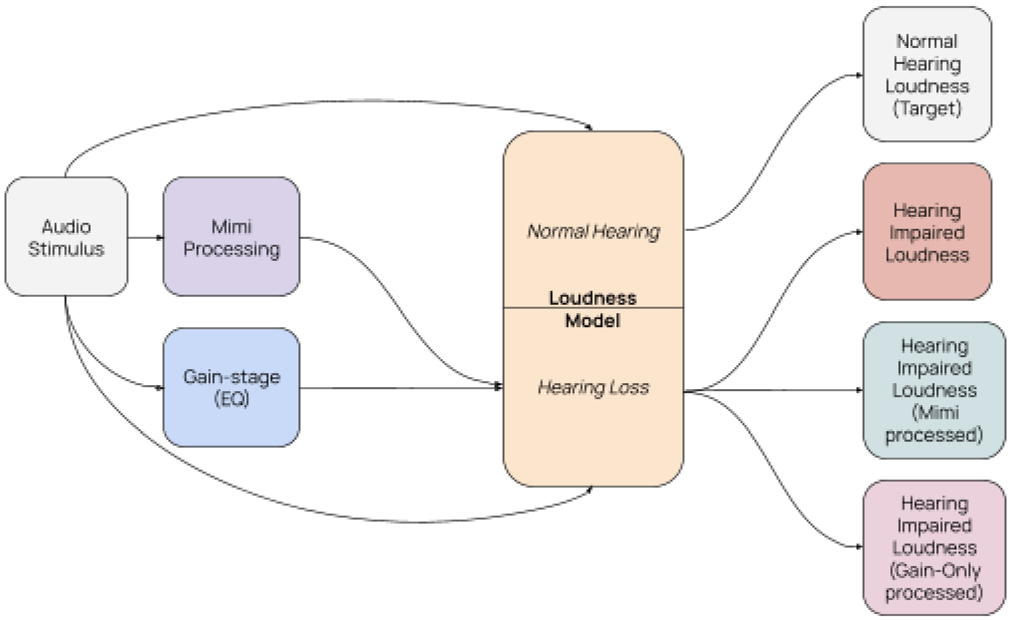

Figure 6: Analysis stage

The amount of restoration for the two types of preprocessing was quantified by calculating the loudness differences between the normal hearing and the hearing impaired listener

$$\Delta L_N = L_{N|NH} - L_{N|HI}$$

This delta was then compared this with the delta between the normal hearing listener and and the hearing impaired listener *with preprocessed audio:*

$$\Delta L_{N|Proc} = L_{N|NH} - L_{N|HI-Proc}$$

In the next step, the effective results of linear processing and Mimi's non-linear preprocessing to counter hearing loss effects were analyzed. The *restoration* for each preprocessing type was quantified by computing the ratio of the two deltas above and then averaging the values across frames:

$$Restoration = \overline{Restoration}_{Frames} = \frac{\Delta L_N - \Delta L_{N|Proc}}{\Delta L_N} \times 100$$

The result for the linear preprocessing and Mimi's non-linear processing is illustrated below: For the given profile Mimi's processing can restore ~ 80% of loudness in the specific critical frequency bands for music (2kHz - 16kHz) while linear processing can only restore ~ 50% of original loudness on average in this range. Because hearing loss for the chosen hearing profile is concentrated at mid- and high frequencies, barely any loudness is lost at 1 kHz and below, and restoration is easily achieved by both processing approaches. On the other hand, any real-time DSP system runs a risk of generating undesired processing artifacts when trying to compensate for variations in individual loudness perception. Therefore, for increasing hearing losses, the processing is typically applied less aggressively.

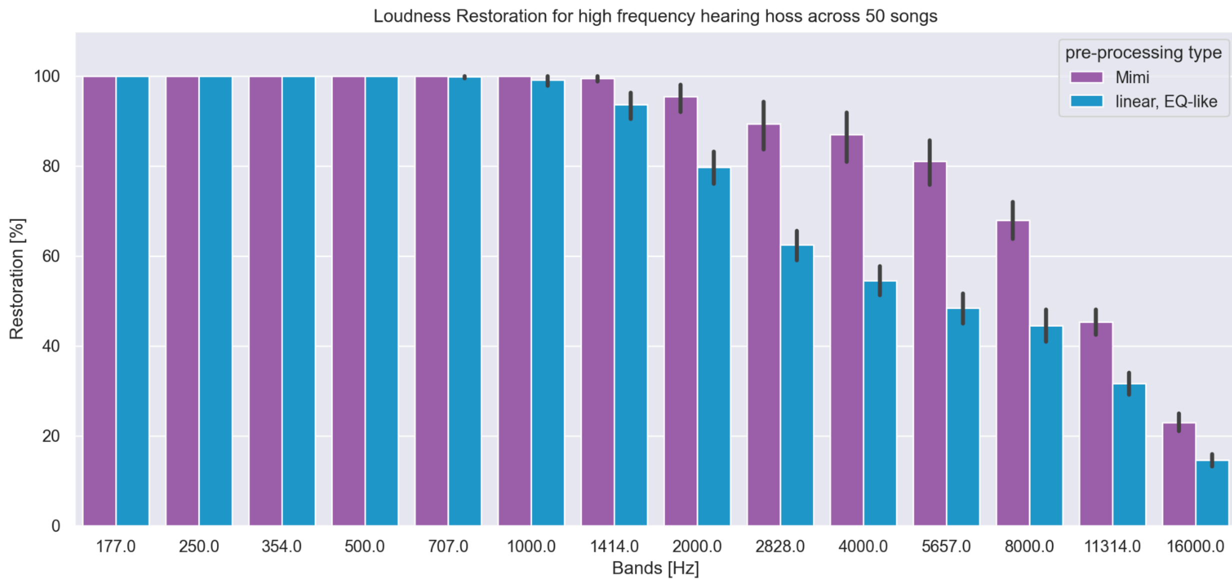

Figure 7: Restoration metric across 50 stimuli for Mimi and linear, gain-only processing

The results were validated and confirmed by calculating the overall loudness of the stimuli after pre-processing and loss simulation using a second, widely known loudness model[2].

3. An EQ is more prone to overshoot

By definition, a linear-gain approach provides a constant signal amplification. As discussed in chapter 2, we set the gain per frequency band based on a representative stimulus, the previously mentioned programme simulation noise. We primarily examined the shortcomings of linear systems in restoring signals with low intensity. However, a constant gain approach can result in additional shortcomings at different points in time in an instantaneous analysis:

- Linear processing is more prone to overshoot. This can lead to digital clipping, which either introduces artifacts or needs management with limiting, which also produces negative side effects that can impact the quality of the sound reproduction.
- In the absence of limiting or clipping, already intense segments can be over-amplified comparatively more in a linear system, which may potentially be harmful to a listener.

The problem of overshooting can be easily demonstrated through a clipping ratio analysis, which is the sum of clipped samples divided by the total number of samples. This analysis can be conducted alongside the previously mentioned case analysis. While Figure 7 demonstrated that Mimi processing significantly outperforms an equivalent EQ-like scenario in terms of restoration strength, Figure 8 reveals that the linear-gain approach simultaneously results in excessive signal amplification and clipping.

---

[2] Glasberg, B. R., Moore, B. C. J. (2002) A model of loudness applicable to time-varying sounds

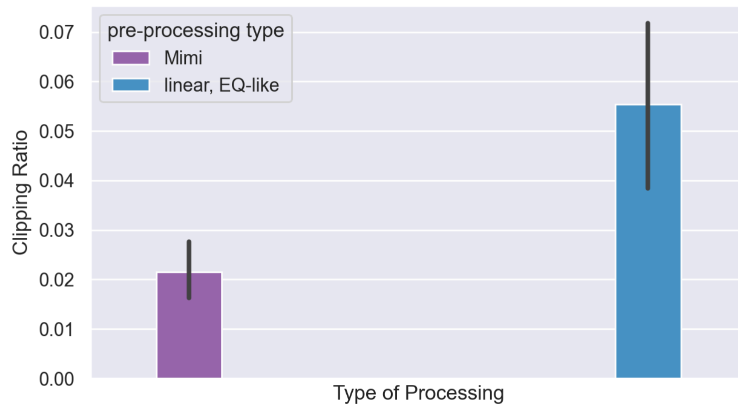

Figure 9: Sound Dose increase due to different signal processing

4. Mimi restores more and delivers a lower sound dose than an EQ

In terms of hearing health, a sound dose analysis represents another important aspect to compare Mimi processing against gain-only processing .

The WHO and the International Telecommunication Union published a Safe listening devices and systems standard (ITU H.871) that defines sound dose as the “total quantity of sound received by the human ear during a specified period. The unit of (sound) dose is Pa^2h." Thus, in the context of the WHO-ITU Global standard, it is the same as sound exposure.

The WHO and ITU Safe listening standard further describes how dose can be measured on devices and calculated for a discrete time segment-based implementation by integrating headphone sensitivity curves, human’s perception and audio content.

We used the definitions and recommendations in the WHO-ITU standard as the basis to evaluate the effect of sound processing on sound dose. For this analysis, a flat headphone frequency response was assumed and a relative comparison was conducted which makes concrete assumptions about absolute physical sound pressure levels obsolete.

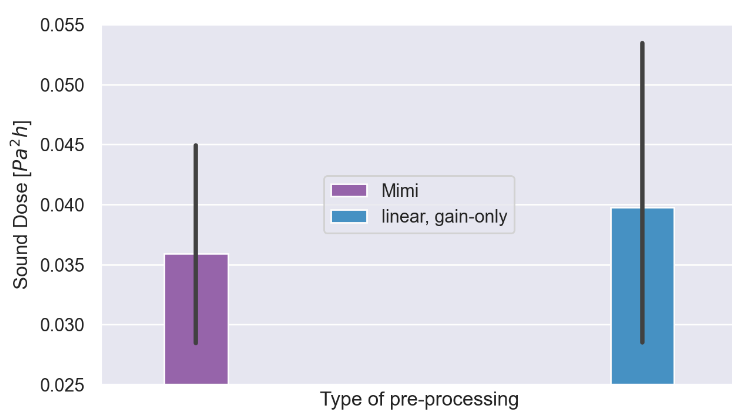

Figure 9: Sound Dose increase due to different signal processing

Figure 9 illustrates that, for the given hearing profile and fitting, Mimi's non-linear processing results in approximately 10% lower sound dose compared to a gain-only solution, while still maintaining higher restoration effectiveness in Mimi's case, as shown above.

However, as an inferior alternative to sophisticated pre-processing, individuals with hearing impairments may resort to simply increasing the device's volume. To demonstrate the adverse effects of this behavior, another analysis was conducted. Loss simulation[3] was applied to the audio that was pre-processed with Mimi processing. The loudness model[4] mentioned above was used to calculate the overall loudness for each pre-processed and loudness loss deteriorated stimulus. Here, Mimi's processing served as the target and stimuli with linear pre-processing were iteratively increased in levelvolume until they matched the equivalent overall loudness. This was done for both the linear gain-only (EQ) scenario and a very simple volume increase per device control (higher volume per device control is considered the most simple form of pre-processing resulting in a flat level increase).

The following figure demonstrates that in comparison to Mimi, loudness equivalent linear processing results in significant increases in sound dose for both the EQ-like gain-only and the simple flat volume increase. Here, linear processing results in 20% higher sound dose on average.

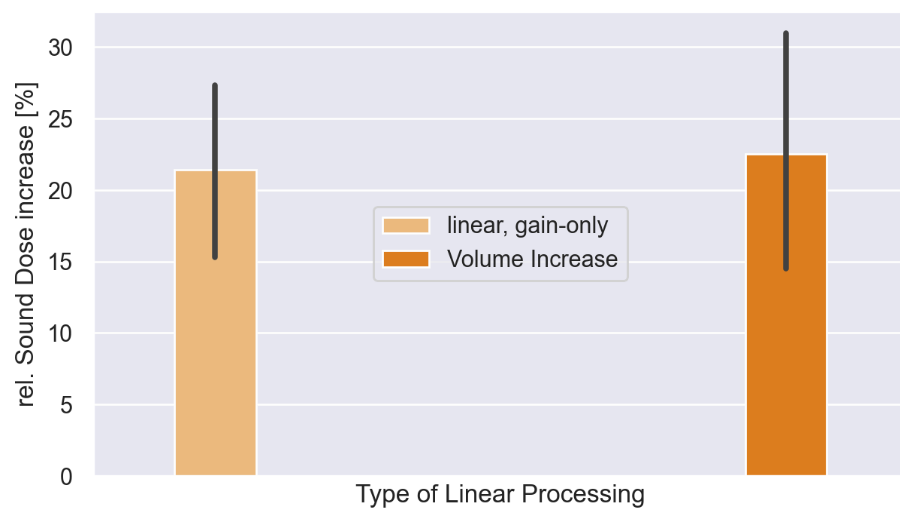

Figure 10: Sound Dose increase for loudness-matched scenarios relative to Mimi processing - with similar loudness for HI listeners for gain-only and volume increase pre-processing. Linear solutions show >20% dose increase compared to Mimi processing on average.

[3] Loss simulation was applied to each audio stimulus by applying the frame-by-frame level attenuation in dBFS that describes the distance between the modeled loudness of the normal hearing and the hearing impaired listener. The loudness model described in chapter 2 was used.
[4] Glasberg, B. R., Moore, B. C. J. (2002) A model of loudness applicable to time-varying sounds

Summary

Due to the non-linear nature of hearing loss, an adequate compensation mechanism must also act in non-linear fashion. Mimi's non-linear processing was compared to a linear gain-only EQ profile for an exemplary high-frequency hearing loss. A restoration metric was employed to quantify the restoration strength for a processed set of music stimuli. The results demonstrated that Mimi's processing restores signal loudness significantly better than an equivalent EQ for a given hearing profile. Furthermore, we showed that while Mimi processing achieves higher restoration strength, it also yields lower clipping and sound dose compared to an equivalent EQ.

Typically in consumer electronic devices, the only vector of control available to the user to compensate for hearing loss is the volume control. Here, we showed that the risk of dose-related harm is actually comparable between EQ and volume-control solutions. Conversely, the risk is reduced by employing a non-linear solution.

Published: June 2024

Mimi Hearing Technologies

Boxhagener Str. 82
10245 Berlin, Germany

info@mimi.io
www.mimi.io