Learning volume control for hearing aids

Jos Leenen, Almer van den Berg, Alexander Ypma, Job Geurts, and Bert de Vries
Algorithm R&D, GN ReSound A/S, Horsten 1, 5612AX, Eindhoven, the Netherlands

A Learning Volume Control (LVC) for hearing aids has been developed, tested and introduced in the market. It has the look and feel of a normal VC, the extra feature is that it gradually learns a more optimal VC setting during regular use of the hearing aid. It does so by combining features of the current input sound with past user behavior (past VC operation stored in the aid's memory). The aimed effect is that users, over time, will need less VC adjustments when being exposed to changing acoustical environments.
Like a normal VC, the LVC will always and instantly change the volume when operated, so the user will stay in immediate control of volume at all times, thus always being able to cope with wanted exceptions to the learned pattern.
Our LVC concept has been tested in a number of patient trials (Chicago, Copenhagen, and Oldenburg) with very comparable results. Average learning amounted to 2.4 dB from the default, with very large individual differences. We also found a large variability in learned volume, per patient, over different environments. This clearly shows the benefit of environmental steering in the personalization of volume.
We conclude that automatic adaptation of volume by a learning algorithm is well appreciated by users, both with respect to environmental steering and personalization.

INTRODUCTION AND BACKGROUND

Digital hearing aids have opened many new possibilities to serve the hearing impaired (Hamacher *et al.*, 2005), including highly increased possibilities for personalization (Chalupper, 2006; Dillon *et al.*, 2006). Personalization of the aid to the needs of a specific user is currently done in a clinical fitting setting. However, real-life sound environments can and will be very different from the clinical fitting environment. Furthermore, individual preferences show variation over a wide range, and show to vary considerably over environments (confirmed e.g. in this study) and time as well. Learning algorithm techniques can help to bridge the serious gap between clinic and real-life. Application in hearing aids is aimed at learning more patient-preferred parameter settings during normal use in real life.

GN ReSound regards learning algorithm techniques as an increasingly important way to add value to hearing aids (Dijkstra, 2007; Ypma, 2006) and hearing aid fitting. Bringing to market Learning Volume Control (LVC) is an important first step in the implementation of this rationale.

Personalization of volume by our LVC works during regular use of the aid in a fully transparent way. No extra information or actions are needed by the user. He/she doesn't

Auditory signal processing in hearing-impaired listeners. 1st International Symposium on Auditory and Audiological Research (ISAAR 2007). T. Dau, J. M. Buchholz, J. M. Harte, T. U. Christiansen (Eds.).
ISBN: 87-990013-1-4. Print: Centertryk A/S.

even have to know that the algorithm learns.

The only exception to this we found in one person who, by reflex, always turned his volume a fixed amount up or down depending on the environment. After unlearning this reflex, this person was extremely pleased with the LVC behavior, it was just what he needed.

The volume settings that are applied in different acoustical environments are personalized using environmental steering. The class of the incoming sound is determined by an on-line environmental classifier algorithm, see Fig. 1. Background here is the assumption that patients will have different volume preferences in different acoustical environments. This assumption has been validated in this study.

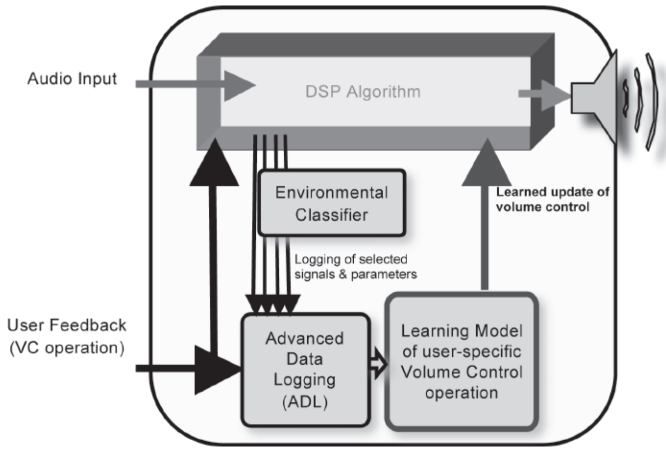

Fig. 1: Learning Volume Control Hearing Instrument.

THE LEARNING ALGORITHM

What we want to achieve is automatic and personalized volume adjustment when the acoustic environment changes significantly. An on-line (in the aid) running environmental classifier has been implemented which is able to recognize 7 different environments: quiet, soft clean speech, soft speech in noise, loud clean speech, loud speech in noise, soft noise only, and loud noise only.

These 7 classes all have their own default gain value which can be interpreted as the non-learning (manually set) part in the personalization process. They serve as the starting point for the learning process that will follow during real life use. As an option these defaults can be set through the fitting software, using 14 sliders (for L + R), see that part of the fitting screen in Fig. 2. If these sliders are adapted appropriate by the dispenser for the patient at hand (e.g. using real-life sound files) the learning process can be accelerated. Our tests showed that not setting the initial defaults resulted in a somewhat slower learning process.

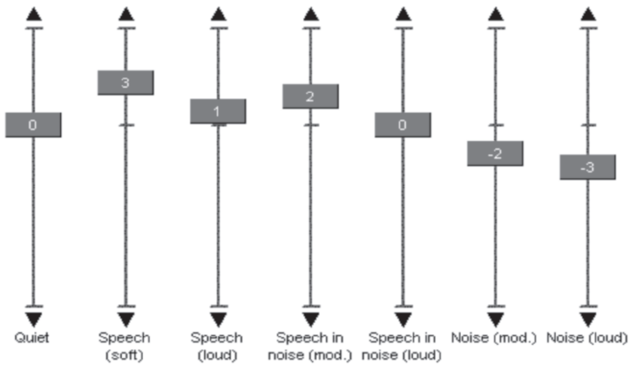

Fig. 2: Fitting Software: LVC User Interface detail (sliders for right HA).

User Consistency – Inconsistency

In any learning algorithm great care must be taken to measure an estimate of the level of consistency of its inputs, in this case user inputs. Consistency is high when the spread in user input (here the volume control setting by the user) is low, when measured over time in identical situation. During periods when consistency is low the LVC algorithm should slow down and eventually stop updating the volume. If confidence in the consistency of input data is high the update process can run full speed ahead. Thus the learning rate will depend on the evaluated level of consistency, see Fig. 3 (indicative plot), where the average number of VC operations per day is plotted against time. The slope is an indication for the speed of the process (learning rate). The algorithm works such that the amount of learning will automatically depend on the evaluated consistency level. After some time of usage the learning process will gradually saturate to a (near) optimal/preferred value, and restart updating if changes in preferences occur.

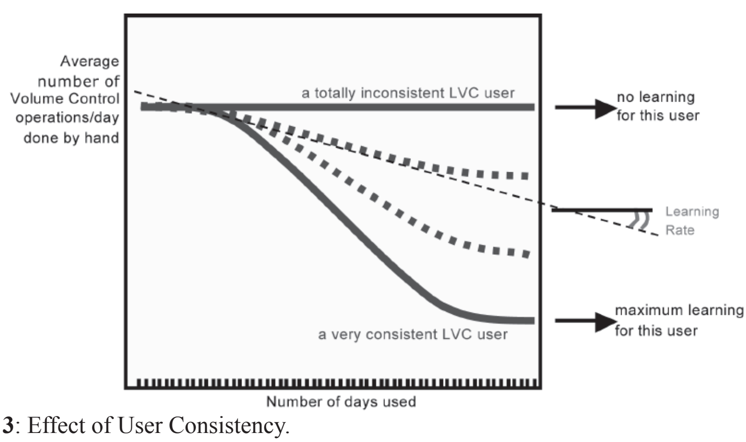

Fig. 3: Effect of User Consistency.

In practice a user often reaches a new volume setting by a combination of larger and smaller up and down turns. Therefore we will allow the user some time to change volume. After a certain time of constant VC position we interpret this value as explicit consent of the wanted new volume value in the new environment and feed it to the learning algorithm.

Learning Update Rule

The learning process can be achieved by the following (simplified) set of equations:

$$G_{vol} = G_{wheel} + G_{lvc} + G_{init} \qquad \text{(Eq. 1)}$$

$$G_{lvc} = |\theta| \bullet \mu \qquad \text{(Eq. 2)}$$

$$\mu[n] = \alpha \bullet \mu[n\text{-}1] + (1\text{-}\ \alpha)\bullet\{G_{vol}[n] - G_{init}\} \qquad \text{(Eq. 3)}$$

Where:

G_{vol} = applied volume in the hearing aid [dB]

G_{wheel} = user-set position of the VC wheel [dB]

G_{lvc} = gain provided by the learning engine [dB]

G_{init} = initial (default) gain, optionally can be set per environment in fitting software [dB]

θ = confidence parameter (high when high consistency is measured by the algorithm)

μ = running average (with smoothing factor α) [dB]

n = explicit consent event number

The explicit consent moment is the moment when the algorithm decides that the new volume is reached (after a possible series of up- and down-turn by the user).

These formulas are evaluated for each of the 7 environmental sound classes.

The learning rate will go up if the use of the VC in combination with the incoming sound is such that the confidence parameter θ will go up, which will only happen for a more consistent user. During more inconsistent VC operation the algorithm will detect only a weak pattern. Then volume learning will be less, or zero in the extreme case (no pattern, meaning random operation, see top line in Fig. 3).

RESULTS FROM PATIENT TRIALS

The LVC concept has been tested in a number of patient trials (Chicago, Copenhagen, and Oldenburg) with very comparable results. In total 41 ears of 22 hearing impaired were tested in a field trial. The patient group was a good representation of the 'mid region' of the hearing impaired population. Patients were told that the test considered a new type of volume control, however they were not told that the algorithm learned from their operations. Initial gains were set to zero dB. The patients used our LVC prototype hearing aids for 2 weeks during daily life. At the end of this period we read out the learned volume per classified environment. The results are shown in Fig. 4 and 5.

The deviation from the default gains, averaged over all patients and all environments, was only 0.6 dB, indicating that the used (company proprietary) fitting rule resulted in the correct gain for 'the average user' in the 'average environment'. However, from Fig. 4 we also learn that there is a large spread in learned volume over patients as well as over environments. These spreads are important reasons why we think LVC is a useful algorithm and patient comments greatly confirmed the usefulness of LVC.

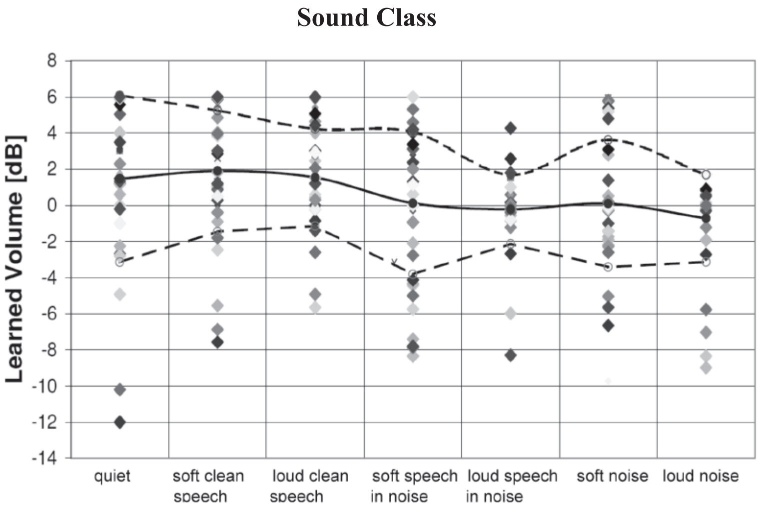

Fig. 4: Learned Volume per Environment in 41 Ears.

There is a small trend in the averages per sound category (solid line) towards lower selected volumes for noisy signals and towards higher selected volumes in quiet and clean speech. This trend can be used in the preset definitions for the 14 sliders in Fig. 2. The dashed lines indicate one standard deviation. Over all 41 x 7 = 287 data points the average absolute value of learning was 2.4 dB, however there were very large individual differences. Even if we, per patient, average over environments large individual differences remain, see Fig. 5.

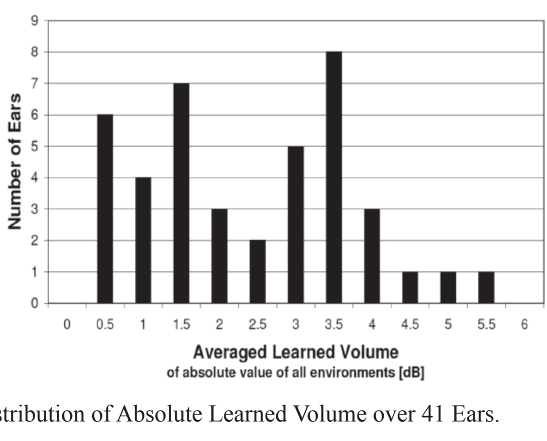

Fig. 5: Distribution of Absolute Learned Volume over 41 Ears.

In a group of 4 users, that continued wearing the device for a second period of 2 weeks, the additional average absolute learning declined to 0.95 dB, indicating a saturation of the learning process within 2 to 3 weeks of continuous use.

Average Environmental Variability

We define, for each ear, the Average Environmental Variability (AEV) as:

$$\mathrm{AEV} = \frac{\sum_{i,j}^{N} |L_i - L_j|}{N^2 - N} \qquad \text{(Eq. 4)}$$

Where:

L_i = learned volume value for sound class i, per ear

N = number of sound classes

The AEV is, for each ear, the average mutual distance of the learned volume in the different sound classes. By this definition, a patient showing bigger AEV is more actively changing volume when changing his/her acoustical environment.

In Fig. 6 the AEV is plotted against the averaged value (over all environments) of the absolute learned volume for the same ear. We conclude that the AEV generally scales with averaged absolute learning, indicating that people with a higher mean learning score also show more variability (over different environments) in learned gain.

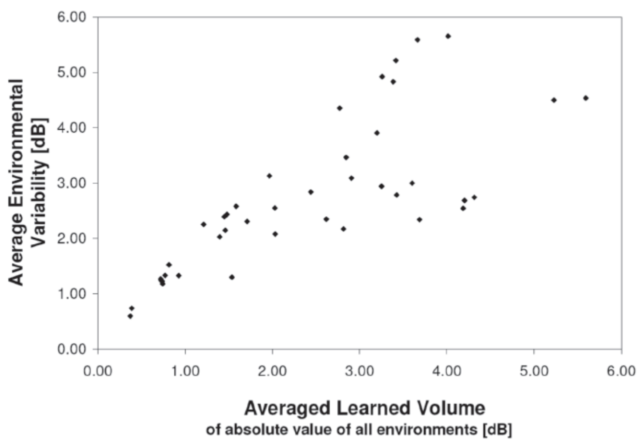

Fig. 6: AEV against Averaged Absolute Learned Volume for 41 Ears.

The fact that far most AEV scores were rather deviant from 0 dB is evidence that environmental steering is crucial in an LVC algorithm. Without environmental steering only points on the x-axis in Fig. 6 can be reached, which is apparently not what users want, thus from the measured AEV scores we conclude that users did benefit from environmental steering.

CONCLUSION

An algorithm for Learning Volume Control in hearing aids has been designed and evaluated in a field trial with 22 patients. Average learning amounted to 2.4 dB from the default in the first 2 weeks. A smaller group, followed for 2 weeks more, showed considerably less learning in the second period, indicating a saturation of the learning process, on average, within 2 to 3 weeks of continuous use.

We found a big spread in learned volume over patients as well as over environments. There is a clear trend in a higher variability of learning (over the different environments) for patients showing a higher environmentally averaged value of absolute learned gain. The relatively high values of environmental variability found, clearly show the need for environmental steering in personalization of volume.

The great majority of the test-patients reported that the LVC had worked fully transparent for them and that they find LVC a useful feature in their hearing aid.

REFERENCES

Chalupper, J. (**2006**). "Changing how gain is selected: the benefits of combining datalogging and a learning VC," The Hearing Review, Dec 2006.

Dijkstra, T. M. H., Ypma, A., de Vries, B., and Leenen, J. R. G. M (**2007**). "The Bayesian Approach to Hearing Aid Fitting: An Example with Common Sense Reasoning," The Hearing Review, Oct 2007.

Dillon, H., Zakis, J. A., McDermott, H., Keidser, G., Dreschler, W, Convery, E. (**2006**). "The trainable hearing aid: What will it do for clients and clinicians?," Hear Jour. **59**(4) 31-36.

Hamacher, V., Chalupper, J., Eggers, J., Fischer, E., Kornagel, U., Puder, H., and Rass, U. (**2005**). "Signal Processing in High-End Hearing Aids: State of the Art, Challenges, and Future Trends," EURASIP Journal on Applied Signal Processing **18**, 2915–2929.

Ypma, A., de Vries, B., and Geurts, J. (**2006-1**). "Robust Volume Control Personalisation From On-Line Preference Feedback," IEEE Int. Workshop on Machine Learning for Signal Processing, Maynooth, Ireland, 2006.

Ypma, A., de Vries, B., and Geurts, J. (**2006-2**). "A learning volume control that is robust to user inconsistency," The second annual IEEE Benelux/DSP Valley Signal Processing Symposium, Antwerp, March 2006.