EFFICIENT INDIVIDUALIZATION OF HEARING AID PROCESSED SOUND

Jens Brehm Nielsen[1,2] *& Jakob Nielsen*[1]

Widex A/S[1]
Nymøllevej 6
3540 Lynge, Denmark
{jeb,jnl}@widex.com

DTU Informatics[2]
Asmussens Alle B305
2800 Kgs. Lyngby, Denmark
jenb@imm.dtu.dk

ABSTRACT

Due to the large amount of options offered by the vast number of adjustable parameters in modern digital hearing aids, it is becoming increasingly daunting—even for a fine-tuning professional—to perform parameter fine tuning to satisfactorily meet the preference of the hearing aid user. In addition, the communication between the fine-tuning professional and the hearing aid user might muddle the task. In the present paper, an interactive system is proposed to ease and speed up fine tuning of hearing aids to suit the preference of the individual user. The system simultaneously makes the user conscious of his own preferences while the system itself learns the user's preference. Since the learning is based on probabilistic modeling concepts, the system handles inconsistent user feedback efficiently. Experiments with hearing impaired subjects show that the system quickly discovers individual preferred hearing-aid settings which are consistent across consecutive fine-tuning sessions for each user.

Index Terms— Hearing aid personalization, Bayesian learning, Gaussian processes, Active learning, Preference learning

1. INTRODUCTION

Modern digital hearing aids (HAs) contain a vast number of adjustable parameters that offer an almost infinite number of possible settings. Different settings make the hearing aids emphasize parts of the incoming sound to make it more or less comfortable, audible, intelligible etc. for the hearing impaired (HI). The procedure of fitting the HAs to the user is performed by skilled professionals like an audiologist.

Having fitted a set of HAs to the hearing loss of the HI user to ensure audibility and intelligibility of incoming sounds, several options are still left for the audiologist to choose from. Some of those are related to the preference of the user. Fine tuning of these parameters is normally done manually by adjusting a number of handles available in the supplied fitting software. At this point, two aspects should be considered. First, due to the large number of parameters—and thus the number of settings—a manual procedure may not be adequate for finding optimal settings for all parameters even for a fine-tuning expert like an audiologist. Secondly, the success of the fine-tuning process depends on the communication between the HA user and the audiologist. Typically, the HA user has not recognized his own preference beforehand, which may muddle the communication and result in an inadequate fine tuning.

To take full advantage of modern digital hearing aids, more sophisticated fine tuning tools are needed. These should discover the best setting for each individual in robust and time-efficient procedures to take full advantage of the flexibility of the HAs.

In this paper, an interactive system is considered that lets the HA user recognize his own preference by comparing different settings simply by listening to the resulting sounds. By letting the user report how much one setting is preferred over another in a sequence of such comparisons, the interactive system starts to learn the preference of the user. At the end, the interactive system is able to suggest which setting (or subset of settings) that is preferred by the HA user. The system builds on the assumption that each user has an *unobserved* internal representation of preference (IRP), which is a stochastic function (or process) of hearing aid settings. In the interactive system, the *mean response* of the IRP is modeled by a Gaussian process (GP) [1], which loosely speaking defines a *distribution* of functions and thus of possible mean responses of the IRP. In the remainder of the article the IRP is used to refer to the *mean response* of the IRP. The distribution of IRPs is updated iteratively each time the user compares and chooses between two HA settings using the GP framework previously proposed in [2]. To reduce the required number of comparisons needed for the system to learn the user's preference, the distribution of IRP provided by the GP is used to decide the next setting pair to compare. In the literature, this is referred to as *active learning*, and in this paper, a bivariate version of *Expected Improvement* (EI) [3] is used.

Several directions have been pursued to develop systems capable of fine-tuning settings of HAs and other devices. Some of the very first attempts used a modified simplex procedure [4], but required an unrealistic amount of preference assessments to converge. Other tournament based attempts have used genetic algorithms [5, 6], but the convergence time tend to scale badly with the number of tunable parameters. One of the most promising suggestions [7] is also probabilistic and contains at least two ideas that are similar to the ideas underlying the work presented here. Firstly, the method is also based on probabilistic modeling of the user's IRP, but does not use state-of-the-art GPs for this. These are included later in a slightly different context in for instance [8]. Secondly, the two methods also rely on probabilistic choice models that directly address the fact that humans are in general not completely consistent performing perceptual evaluations. However, the two methods are based only on forced choices (discrete decisions) using the choice model and framework from [9, 10, 11], in which subjects only select the option they prefer (discrete choice). This is in contrast to the choice model proposed in [11], in which subjects also decide how much they prefer the selected setting (continuous decision). The results in [2] give reason to believe that the additional information contained in continuous decisions reduces the number of required comparisons needed to learn a user's preferred setting. This is really the key for the application considered in this work. It is, however, beyond the scope of this work to actually compare results obtained with discrete choices to those obtained with continuous decisions. Nevertheless, this is definitely of great interest for future research. Instead, the focus is to investigate the variability between IRP and thus the preferred setting suggested by the system using continuous decisions.

To test the fine-tuning abilities of the proposed interactive system, two adjustable parameters of a HA were fine-tuned individually to five different HI users. By comparing the results from two similar sessions with each subject, the variability of the found best setting can be investigated. The two HA parameters that were adjusted in the experiments changed how both noise reduction and speech enhancement algorithms should react to the incoming sound.

This article is organized as followed: In section 2, the interactive system is outlined and an explanation of the experiments is provided in section 3. Results are presented in section 4, and finally, section 5 contains the discussion.

2. MODELING FRAMEWORK

A user's *internal representation of preference* (IRP)—referred to as $f : \mathcal{X} \rightarrow \mathbb{R}$—is modeled by a (zero-mean) *Gaussian process* (GP) [1]. The set $\mathcal{X} = \{\mathbf{x}_i \in \mathbb{R}^d : i = 1...n\}$ is the entire set of the n possible settings of the $d = 2$ HA parameters. A GP is a non-parametric—and thus flexible—discriminative Bayesian approach, which defines a distribution of entire functions, "*any finite number of which have a joint Gaussian distribution*" [1, Def. 2.1]. This simply implies that any finite number of function values, $\mathbf{f} = [f(\mathbf{x}_1), ..., f(\mathbf{x}_n)]^\top$, have a distribution given by a multivariate Gaussian distribution as

$$p(\mathbf{f}) = \mathcal{N}(\mathbf{0}, \mathbf{K}), \tag{1}$$

with the elements of $\mathbf{K}$ given by $[\mathbf{K}]_{i,j} = k(\mathbf{x}_i, \mathbf{x}_j)$, where $k(\cdot, \cdot)$ is a covariance function (or kernel), which generally speaking defines the *smoothness* of the functions. For an introduction to kernels, see [1, Chap. 4] or [12, Chap. 6].

The fundamental benefit from the GP is that Eq. 1 can be used as a *prior* distribution of a user's IRP before any preference assessments have been performed by the user. In the Bayesian framework, the distribution of the user's IRP is recalculated conditioned on the preference assessment(s) that have been observed to give the *posterior* distribution of the user's IRP as

$$p(\mathbf{f}|\mathcal{Y}) \propto p(\mathcal{Y}|\mathbf{f})p(\mathbf{f}), \tag{2}$$

where $p(\mathcal{Y}|\mathbf{f})$ is the *likelihood* which is defined by a specific observational model (choice model). In this work, users assess their *degree* of preference (continuous decision) between two particular HA settings. To update the posterior (and *predictive*) distribution in the GP framework at any given point in the experiment with a particular number of performed preference assessments, the model proposed in [2] is used. The specific functional form of that observational model as well as details about inference and predictions are provided in [2] and will therefore not be presented here.

To reduce the number of preference assessments required to discover the optimal setting, *active learning* is used. Active learning can be formulated in several ways, but the statistics provided by the GP framework makes it possible to use a slightly modified version of *Expected Improvement* (EI) [3]. In contrast to the original formulation [3], the modification also includes the correlation between function values when calculating the (modified) EI. The added correlations are directly available from the GP framework. The EI for a possible new setting $\mathbf{x}_i$ is thus calculated in closed form as

$$EI(\mathbf{x}_i) = \sigma_{EI} \cdot \phi\left(\frac{\mu_{EI}}{\sigma_{EI}}\right) + \mu_{EI} \cdot \Phi\left(\frac{\mu_{EI}}{\sigma_{EI}}\right), \tag{3}$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ is the standard normal distribution and standard normal cumulative distribution functions, respectively, $\mu_{EI} = \mu_i - \mu_{\text{max}}$ and $\sigma_{EI}^2 = \sigma_i^2 + \sigma_{\text{max}}^2 - 2 \cdot \text{cov}_{i,\text{max}}$. Here, the $_{\text{max}}$ index refers to the point with the *current* largest predicted IRP and the notation

$$p\left(\begin{bmatrix} f_{\text{max}} \\ f_i \end{bmatrix}\right) = \mathcal{N}\left(\begin{bmatrix} \mu_{\text{max}} \\ \mu_i \end{bmatrix}, \begin{bmatrix} \sigma_i^2 & \text{cov}_{i,\text{max}} \\ \text{cov}_{i,\text{max}} & \sigma_{\text{max}}^2 \end{bmatrix}\right) \tag{4}$$

has been used for the two-variate marginal of the predictive normal distribution given by the GP framework.

Typically in active learning theory, an explicit trade-off between exploration (of unseen regions of input space) and exploitation (of "known" regions of input space) must be made. Generally, a system will exhibit slow convergence with too much emphasis on exploration, but will quickly get stuck in a sub-optimal solution, if too much emphasis is put on exploitation. In this work, the next proposed setting to compare with the current best one is sampled from a multinomial distribution, where the probability of a given setting is proportional to its EI given by Eq. 3. This was done to put slightly more emphasis on exploration.

3. MEASUREMENT PROCEDURE

To illustrate the behavior of the suggested interactive fine-tuning system, an experiment with five (native danish) HI subjects was conducted. To obtain an indicate of the expected variability in the proposed settings for individual HI users between consecutive fine-tuning sessions, the experiment consisted of both a test session and a re-test session. The two sessions were conducted on two separate days.

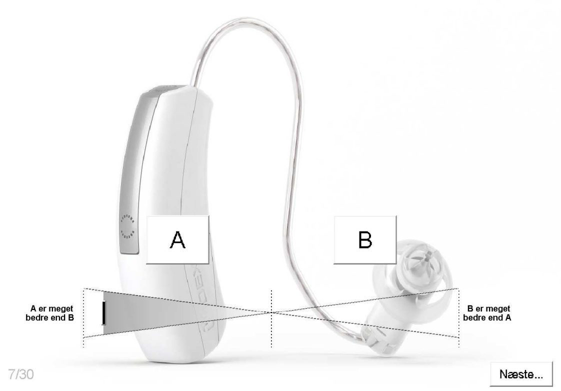

Fig. 1. The (danish) graphical user interface used in the experiments. The buttons, 'A' and 'B', were used to switch between the two current settings. The slider—currently positioned far to the left—was used to indicate the degree of preference between the two settings by how far it was positioned towards either of the two settings, 'A' or 'B'. No preference was indicated by leaving the slider at the center. After positioning the slider, the user continued to the next comparison be clicking the button in the lower-right corner.

In each of the two sessions, each subject conducted thirty comparisons between pairs of HA settings. Subjects wore (experimental) hearing aids fitted (binaurally) in advance to compensate for each individual's hearing loss, and listened to running speech in car noise played back over loudspeakers. Via a graphical user interface (see Fig. 1), the user could switch between the two current HA settings and report their degree of preference. The users were not instructed to focus on particular parts of the sound or on particular attributes,

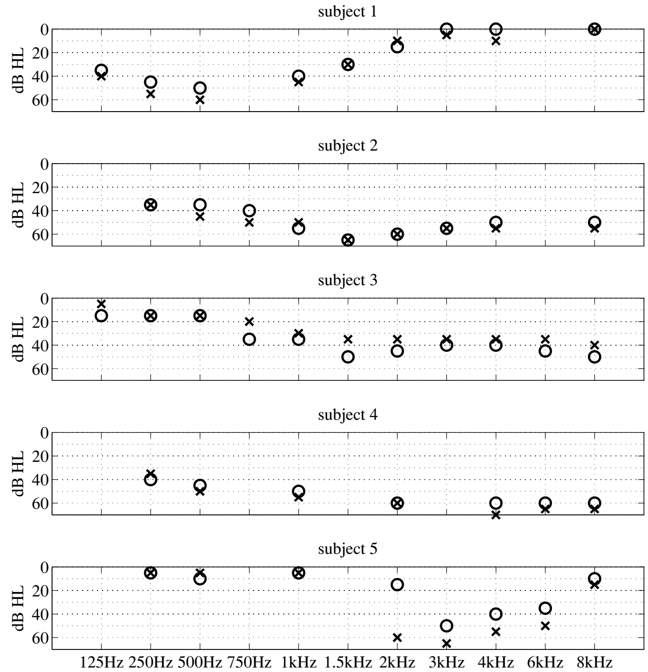

Fig. 2. Audiograms of each individual subject. Crosses and circles correspond to the left and right ear, respectively.

but were only provided with an introduction to the acoustical scene reflected in the sound file. The hearing loss of each individual subject is found in Fig. 2

4. RESULTS

In Fig. 3, the predictions of the IRP for both the test and re-test sessions for each of the five test subjects are depicted. Since the IRPs are unit free, the reader should be aware that the colors cannot be compared across subjects, and similarly, high-preference regions should in general not be interpreted as being "good", but only as being "better than" blue or green. Hence, the predicted IRP only reflects relative properties.

Considering that a parameter change from one end of the space to the other is extremely subtle, the predicted high- and low-preference regions between test and re-tests within each subject are consistent, except for subject 5. The results with subject 5 do, however, coincide with what subject 5 expressed after the sessions, namely that the subject was unable to hear any differences between any of the pairs of presented settings. For this reason, subject 5 chose only occasionally to move the slider, and when the subject did, the subject moved it as little as possible.

A statistic significance test using forced choices was performed to prove significance between the most and least preferred settings discovered by the system in the two sessions. Options 1 and 2 were mixed randomly in eleven trials, yet this only proved significance ($p < 0.005$) for subject 3.

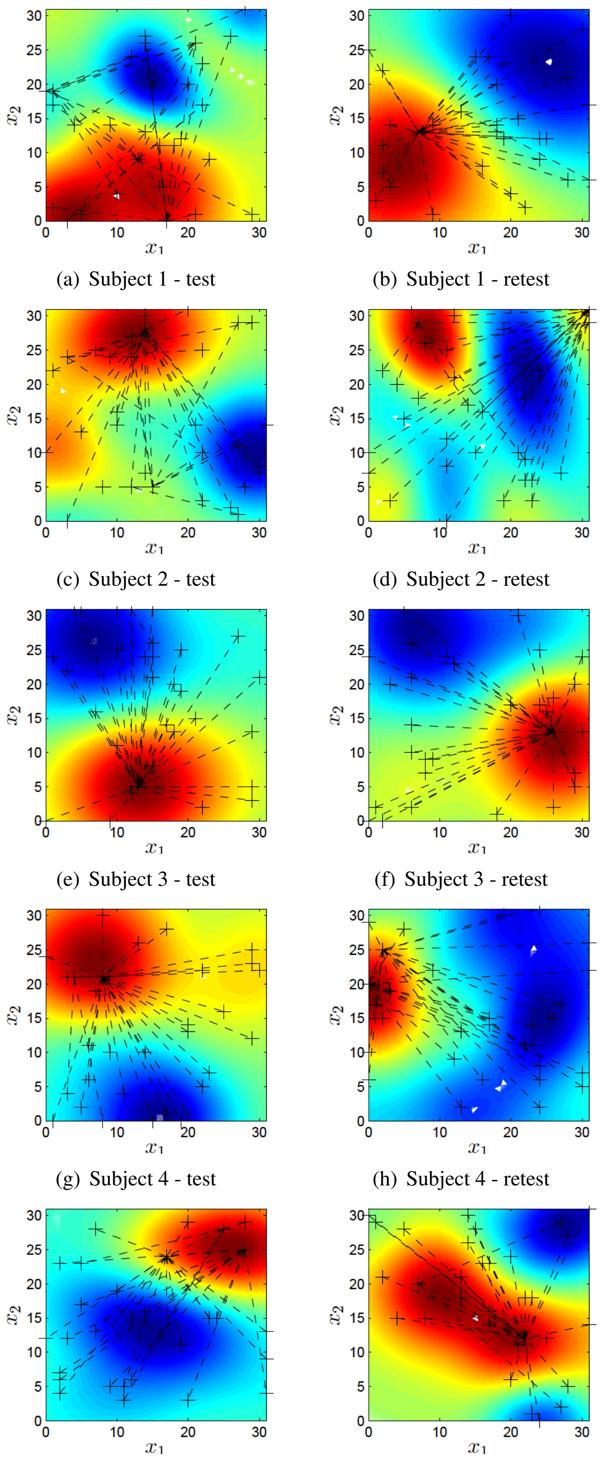

(a) Subject 1 - test (b) Subject 1 - retest

(c) Subject 2 - test (d) Subject 2 - retest

(e) Subject 3 - test (f) Subject 3 - retest

(g) Subject 4 - test (h) Subject 4 - retest

(i) Subject 5 - test (j) Subject 5 - retest

Fig. 3. IRP as a function of the two HA parameters, x_1 and x_2, predicted by the fine-tuning algorithm after 30 comparisons for the test (left column) and re-test (right column) sessions. Red and blue colors indicate high and low preference regions, respectively. Crosses connected with a dashed line indicate comparisons. Note, the IRPs are unit-free.

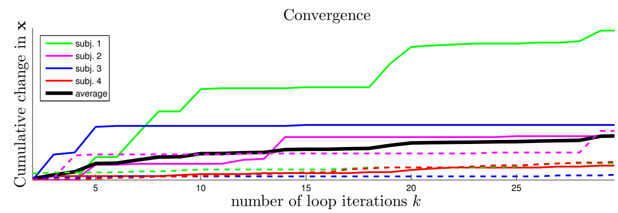

Fig. 4. The cumulative euclidean change in the location of the maximum point of the predicted IPR after a each new assessment as a function of the number of assessments.

The number of assessments that each subject needs to perform before the algorithm discovers a steady preferred setting is visualized in Fig. 4 (see caption for an explanation). Note, that subject 5 has not been included, since subject 5 did not prefer any setting over others and hence did not convergence.

5. SUMMARY AND DISCUSSION

Overall, the reproducibility of the found preferred settings is satisfactory given the subtle differences between parameter settings and is found well before the 30th assessment. However, since the perceptual differences between settings are very subtle, it was not possible to prove or to reject significance of the preferred settings overall. However, the apparent good reproducibility indicates that the found preferred settings are actually a result of the subjects' individual preferences and not a results of a random effect.

The variability in the preferred settings across users from the results in Fig. 3 corroborates previous findings in the literature [13] that individual preferences among HA users do exist, and the system proposed here discovers such preferences before the subjects have performed twenty comparisons in a worst case scenario (see Fig. 4). In case of parameter settings that are perceptually easier to distinguish, the required number of comparisons would presumably be even smaller.

The results presented here are preliminary and serve merely to visualize how the system works. In future work, especially the scaling issue with respect to the number of required comparisons in relation to the number of adjustable parameters is of interest. Also, a similar experiment should be conducted in the future, but with parameter settings that are easier to distinguish from each other, to verify that settings that are suggested by the system to be preferred are significantly different from settings that are not suggested to be preferred. Next, better convergence measures based on the actual statistics provided by the probabilistic modeling framework should be studied. One possible suggestion could be the mean of the EI across settings. Finally, investigation of suitable metrics for expressing the similarity between test/re-test results would be interesting. One (Bayesian) suggestion could be based on the likelihood of the test data given the re-test data or vice versa.

6. REFERENCES

[1] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*, MIT Press, 2006.

[2] B. S. Jensen, J. B. Nielsen, and J. Larsen, "Efficient Preference Learning with Pairwise Continuous Observations and Gaussian Processes," *IEEE Workshop MLSP, Beijing*, September 2011.

[3] D. R. Jones, "A Taxonomy of Global Optimization Methods Based on Response Surfaces," *Journal of Global Optimization*, vol. 21, no. 4, pp. 345–383, 2001.

[4] F. K. Kuk and N. M. C. Pape, "The reliability of a modified simplex procedure in hearing aid frequency-response selection," *J Speech Hear Res*, vol. 35, no. 2, pp. 418–429, 1992.

[5] D. Baskent, C. L. Eiler, and B. Edwards, "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. 370–80, June 2007.

[6] E. A. Durant, G. H. Wakefield, D. J. Van Tasell, and M. E. Rickert, "Efficient perceptual tuning of hearing aids with genetic algorithms," *Speech and Audio Processing, IEEE Transactions on*, vol. 12, no. 2, pp. 144–155, 2004.

[7] T. Heskes and B. de Vries, "Incremental utility elicitation for adaptive personalization," in *Proceedings of the 17th Belgium-Netherlands Conference on Artificial Intelligence, Brussels*. 2005, pp. 127–134, Citeseer.

[8] Adriana Birlutiu, Perry Groot, and Tom Heskes, "Multi-task preference learning with Gaussian processes," in *Proceedings of the 17th European Symposium on Artificial Neural Networks (ESANN)*, 2009, pp. 123–128.

[9] L. L. Thurstone, "A Law of Comparative Judgement," *Psychological Review*, vol. 34, 1927.

[10] R. D. Bock and J. V. Jones, "The Measurement and Prediction of Judgment and Choice," 1968.

[11] W. Chu and Z. Ghahramani, "Preference Learning with Gaussian Processes," *Proceedings of the 22nd International Conference on Machine Learning (ICML)*, pp. 137–144, 2005.

[12] C. M. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2006.

[13] K. H. Arehart, J. M. Kates, M. C. Anderson, and L. O. Harvey, "Effects of noise and distortion on speech quality judgments in normal-hearing and hearing-impaired listeners.," *The Journal of the Acoustical Society of America*, vol. 122, no. 2, pp. 1150–64, Aug. 2007.