Hearing Aid Personalization

Jens Brehm Nielsen[*], **Jakob Nielsen**
Widex A/S
Nymøllevej 6
3560 Lynge, Denmark
{jeb,jnl}@widex.com

Bjørn Sand Jensen, **Jan Larsen**
DTU Compute
Matematiktorvet 303B
2800 Kgs. Lyngby, Denmark
{bjje,janla}@dtu.dk

Abstract

Modern digital hearing aids require and offer a great level of personalization. Today, this personalization is not performed based directly on what the user actually perceives, but on a hearing-care professional's interpretation of what the user explains *about what is perceived*. In this paper, an interactive personalization system based on Gaussian process regression and active learning is proposed, which personalize the hearing aids based directly on what the user perceives. Preliminary results demonstrate a significant difference between a truly personalized setting obtained with the proposed system and a setting obtained by the current practice.

1 Introduction

Hearing aids (HAs) [1] are fitted by predetermined rules (prescriptions [1, Chapter 10]) given frequency-dependent hearing thresholds—called an audiogram—of the hearing-impaired user. These rules are based on years of practical experience and research of the human auditory system, however nobody knows exactly how the fitted HAs sound like, except of course, the user. From empirical studies, it is well-known [1, Chapter 12], that users with the same audiogram may benefit from—and prefer—very different HA settings. Therefore, a hearing-care professional with years of experience often needs to manually fine-tune the HAs beyond the predetermined prescription. This fine tuning is typically based on oral feedback from the user [1, Chapter 12]. In effect, this feedback is the user's oral translation of the perception using a *description* meaningful to the subject. This description, however, might not necessarily give meaning to the hearing-care professional. It is believed that HA users would benefit greatly if the HAs were adjusted and personalized based directly on how the devices sounds—and not on a poorly aligned translation thereof. In this paper, a machine-learning based personalization system is proposed, which adjusts hearing aid settings based on user feedback, which mimics what the individual actually hears. From the user's perspective, the feedback is returned as a *degree-of-preference* rating between two different hearing aid settings. This is an intuitive way of expressing what is perceived while inducing a low cognitive load compared to conveying an oral response to a single setting. The feedback is used to learn a Gaussian process regression model of the user's latent objective function—the optimum of which corresponds to the truly personalized setting. To quickly find this optimum, the GP model is repeatedly updated based on the feedback from the user and subsequently used to select the next comparison to present to the user using active learning. Fast convergence is an absolute requirement, because even quarters of an hour of careful listening is a very demanding task—especially for most hearing aid users.

2 Personalization System

The personalization system is an *interactive loop* visualized in Fig. 1. The loop essentially contains three parts: I) A modeling part where the user's objective function is modeled by a Gaussian process

[*]DTU Compute, Matematiktorvet 303B, 2800 Kgs. Lyngby, Denmark jenb@dtu.dk

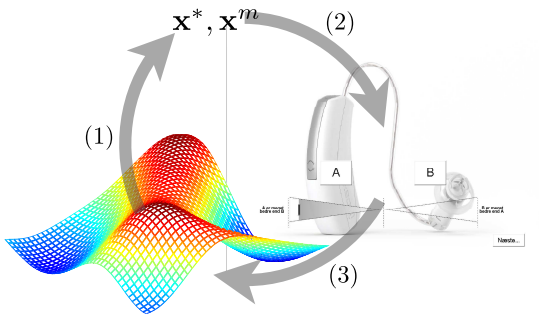

Figure 1: (1): a new optimal setting is determined based on the current (probabilistic) estimate of the subject's objective function. (2): the optimal setting is compared to the setting which maximizes the current estimate of the subject's objective function, and the subject assesses the *degree of preference* between the two settings. (3): the estimate of the subject's objective function is updated based on the recent assessment.

based on the feedback obtained, II) an active learning part setting op the next comparison based on the current state of the model, and III) an user interface part.

2.1 Part I: Modeling the User's Objective from Feedback

The modeling of the subject's objective function is performed in a Bayesian non-parametric setup based on Gaussian Processes (GPs) [2]. In the following, GP regression from *degree of preference* observations will be explained. The GP framework is based on previous work found in [3].

2.1.1 Gaussian Process Prior

A Gaussian process (GP) defines a prior, $f(\mathbf{x}) \sim \mathcal{GP}\left(0, k(\mathbf{x}, \cdot)_{\boldsymbol{\theta}_\mathcal{C}}\right)$, over functions, $f : \mathbb{R}^D \rightarrow \mathbb{R}, \mathbf{x} \mapsto f(\mathbf{x})$, where $k(\cdot, \cdot)_{\boldsymbol{\theta}_\mathcal{C}}$ is a covariance function or kernel with parameters $\boldsymbol{\theta}_\mathcal{C}$. In this paper, a squared exponential (SE) kernel with individual length scales λ_d for each input dimension (ARD) will be used, hence $\boldsymbol{\theta}_\mathcal{C} = \{\sigma_f, \lambda_1, ..., \lambda_D\}$. Given a finite set of function values (random variables), $\mathbf{f} = [f(\mathbf{x}_1), ..., f(\mathbf{x}_n)]^\top$ for $\mathcal{X} = \{\mathbf{x}_i \in \mathbb{R}^D | i = 1, ..., n\}$, the GP defines a joint distribution over the function values as $p(\mathbf{f}|\mathcal{X}, \boldsymbol{\theta}_\mathcal{C}) = \mathcal{N}(\mathbf{0}, \mathbf{K})$, where $[\mathbf{K}]_{i,j} = k(\mathbf{x}_i, \mathbf{x}_j)_{\boldsymbol{\theta}_\mathcal{C}}$. By specifying the likelihood $p(\mathcal{Y}|\mathbf{f}, \boldsymbol{\theta}_\mathcal{L})$ of some set of observations $\mathcal{Y}$ given the finite set of function values $\mathbf{f}$ the posterior distribution over the function values $\mathbf{f}$ is given by Bayes formula

$$p(\mathbf{f}|\mathcal{Y}, \mathcal{X}, \boldsymbol{\theta}) = \frac{p(\mathcal{Y}|\mathbf{f}, \boldsymbol{\theta}_\mathcal{L}) p(\mathbf{f}|\mathcal{X}, \boldsymbol{\theta}_\mathcal{C})}{p(\mathcal{Y}|\mathcal{X}, \boldsymbol{\theta})} = \frac{p(\mathcal{Y}|\mathbf{f}, \boldsymbol{\theta}_\mathcal{L}) p(\mathbf{f}|\mathcal{X}, \boldsymbol{\theta}_\mathcal{C})}{\int p(\mathcal{Y}|\mathbf{f}, \boldsymbol{\theta}_\mathcal{L}) p(\mathbf{f}|\mathcal{X}, \boldsymbol{\theta}_\mathcal{C}) d\mathbf{f}}, \tag{1}$$

where the hyper-parameters $\boldsymbol{\theta} = \{\boldsymbol{\theta}_\mathcal{L}, \boldsymbol{\theta}_\mathcal{C}\}$ contain both likelihood and covariance parameters.

2.1.2 Beta likelihood

Following previous work [3], GP regression from pairwise continuous observations (*degree of preference*) is performed with a likelihood function based on a re-parameterized beta distribution. Consider a set of pairwise observations $\mathcal{Y} = \{y_k \in (0, 1) | k = 1, ..., m\}$ of the *degree of preference* between two distinct inputs $u_k, v_k \in \{1, ..., n\}$, implying that $\mathbf{x}_{u_k}, \mathbf{x}_{v_k} \in \mathcal{X}$. With this formulation, an dominant preference for the first option u_k is reflected by $y_k \rightarrow 0$, whereas an dominant preference for the second option v_k is reflected by $y_k \rightarrow 1$ reflects. No preference is indicated by $y_k = 0.5$. A suitable likelihood function $p(y_k|\mathbf{f}_k)$ can now be constructed given the function values for the two input instances $\mathbf{f}_k = [f(\mathbf{x}_{u_k}), f(\mathbf{x}_{v_k})]^\top$, by re-parameterizing the beta distribution, $\text{Beta}\left(\cdot\,; \alpha, \beta\right)$, as $p(y_k|\mathbf{f}_k, \boldsymbol{\theta}_\mathcal{L}) = \text{Beta}\left(y_k; \nu\zeta(\mathbf{f}_k, \sigma), \nu(1 - \zeta(\mathbf{f}_k, \sigma))\right)$,
where $\boldsymbol{\theta}_\mathcal{L} = \{\nu, \sigma\}$ is the set of likelihood parameters, ν is a dispersion parameter around the mean $\zeta(\mathbf{f}_k, \sigma)$, which is defined by

$$\zeta(\mathbf{f}_k, \sigma) = \Phi\left(\frac{f(\mathbf{x}_{v_k}) - f(\mathbf{x}_{u_k})}{\sqrt{2}\sigma}\right), \tag{2}$$

where $\Phi(\cdot)$ is the standard normal cumulative density function—with zero mean and unit variance—and σ is a slope parameter. By assuming that observations are independent given the latent function values $\mathbf{f}$, the likelihood can be written as $p(\mathcal{Y}|\mathbf{f}, \boldsymbol{\theta}_\mathcal{L}) = \prod_{k=1}^{m} p(y_k|\mathbf{f}_k, \boldsymbol{\theta}_\mathcal{L})$, which is plugged into Eq. 1 together with the GP prior from Eq. 2.1.1 to complete the Bayesian model.

2.1.3 Inference and Prediction

The Gaussian process model outlined above is not analytical tractable due to the Beta-like likelihood function from Eq. 2.1.2. Instead, approximate inference based on the Laplace approximation [2, Section 3.4] is performed as in [3], giving

$$p(\mathbf{f}|\mathcal{Y},\mathcal{X},\boldsymbol{\theta}) \approx q(\mathbf{f}|\mathcal{Y},\mathcal{X},\boldsymbol{\theta}) = \mathcal{N}\left(\hat{\mathbf{f}}, \left(\mathbf{W}+\mathbf{K}^{-1}\right)^{-1}\right) \tag{3}$$

where $\hat{\mathbf{f}}$ is the maximum of the posterior (mode) and $[\mathbf{W}]_{i,j} = -\sum_{k=1}^{m} \frac{\partial^2 \log p(y_k|\mathbf{f}_k\boldsymbol{\theta}_\mathcal{L})}{\partial f(\mathbf{x}_i)\partial f(\mathbf{x}_j)}$. Note, that unlike traditional classification and regression problems, $\mathbf{W}$ does not become diagonal due to the pairwise structure. For further details, see [4].

For (hyper) parameter optimization, traditional ML-II optimization [2, Chapter 5.2] results in large length scales with few observations (< 20) available. This is an undesirable property in combination with active learning. Therefore, a half-student's-t prior is placed on critical hyper-parameters, resulting in the evidence $q(\mathcal{Y}|\mathcal{X},\boldsymbol{\theta})$ of the Laplace approximation being augmented with a extra term (see [5] for similar use). The resulting MAP-II scheme for hyper parameter optimization is therefore:

$$\boldsymbol{\theta}^{\text{MAP-II}} = \arg\max_{\boldsymbol{\theta}} \left\{\log q(\mathcal{Y}|\mathcal{X},\boldsymbol{\theta}) + \log p(\boldsymbol{\theta})\right\}, \tag{4}$$

where $\sigma_f \sim \delta(\sigma_f = 4)$, $\lambda_d \sim \text{half-St}(\cdot|6,100)$ and $\sigma, (\nu-2) \sim \text{half-St}(\cdot|6,10)$ with

$$\text{half-St}(z;\xi,s) \propto \left(1+\frac{1}{\xi}\left(\frac{z}{s}\right)^2\right)^{-(\xi+1)/2}. \tag{5}$$

The predictive distribution $p(\mathbf{f}_*|\mathcal{Y},\mathcal{X},\mathcal{X}_*,\boldsymbol{\theta})$ of the function values $\mathbf{f}_* = [f(\mathbf{x}_1^*),...,f(\mathbf{x}_o^*)]^\top$ at new input locations $\mathcal{X}_* = \{\mathbf{x}_l^* \in \mathbb{R}^D | l = 1,...,o\}$ is given by

$$p(\mathbf{f}_*|\mathcal{Y},\mathcal{X},\mathcal{X}_*,\boldsymbol{\theta}) = \mathcal{N}\left(\boldsymbol{\mu}_*, \boldsymbol{\Sigma}_*\right), \tag{6}$$

$$\boldsymbol{\mu}_* = \mathbf{k}_*^\top \left(\mathbf{W}+\mathbf{K}^{-1}\right)\hat{\mathbf{f}} \tag{7}$$

$$\boldsymbol{\Sigma}_* = \mathbf{K}_* - \mathbf{k}_*^\top \left(\mathbf{I}+\mathbf{W}\mathbf{K}\right)^{-1}\mathbf{W}\mathbf{k}_*. \tag{8}$$

Predicting preference relations y_* are not of interest in the present paper, but are considered in [4].

2.2 Part II: Efficient Sequential Design for Faster personalization

In most machine learning algorithms sequential design (or active learning) aims at maximizing the generalization performance of a model in terms of a specific measure of performance. In this work, the generalization performance is not of particular importance. Instead, the aim is to find the maximum—ideally the global one—of the unknown objective function. For this, a bivariate version of the *expected improvement* [6] (EI) is used given by

$$EI = \sigma_I \phi\left(\frac{\mu_I}{\sigma_I}\right) + \mu_I \Phi\left(\frac{\mu_I}{\sigma_I}\right) \tag{9}$$

with $\mu_I = [\boldsymbol{\mu}_*]_l - [\boldsymbol{\mu}_*]_{\max}$, and $\sigma_I^2 = [\boldsymbol{\Sigma}_*]_{l,l} + [\boldsymbol{\Sigma}_*]_{\max,\max} - 2\cdot[\boldsymbol{\Sigma}_*]_{l,\max}$. The EI is optimized with a gradient descent method with 5 random initializations. By using only 5 random initializations, a little more exploration is build into the sequential designs for robustness.

2.3 Part III: Interface

The system relies on the *degree-of-preference* paradigm discussed earlier, and the user interface (PC screen) presents two options to the user, A and B, as illustrated in Fig. 1. The user can now listen to both options, and finally select to which degree A or B is preferences by dragging the sliders to either side.

3 Preliminary Results

The feasibility of the system was evaluated in an experiment where the personalization system was used to find the preferred settings of HAs with several HA users. The preliminary results [1] in Fig. 2

[1] A full analysis of the results is currently in preparation

show a long-term spectra of the sound pressure level (SPL) at the eardrum of a HA user wearing HAs while listening to a piece of music. Each spectrum corresponds to a particular four parameter setting of the HAs. The spectra labeled *test 1* and *test 2* correspond to two HA settings obtained for the user with the personalized system. The spectrum labeled "prescription" corresponds to the setting resulting from current practice using the user's audiogram and the prescription. In a separate test, it was validated that the setting of "Test 2" is significantly ($p_0 < 0.05$) preferred over the setting resulting from the prescription. The system takes about 10 minutes to discover the preferred setting.

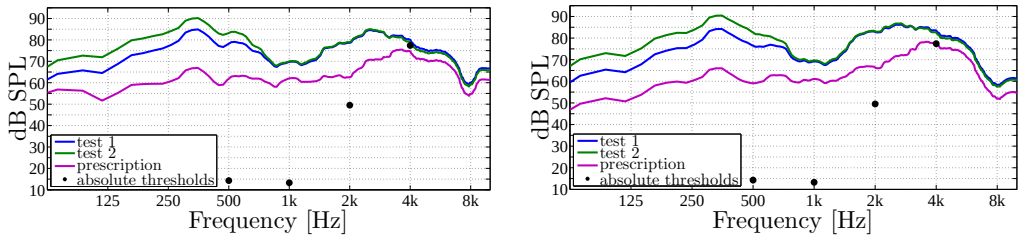

(a) Left Ear (b) Right Ear

Figure 2: KEMAR measurements of long-term power spectra of the sound pressure level at the eardrum of a HA user wearing HAs while listening to a piece of music. The user's thresholds at four distinct frequencies are marked with black dotes.

4 Discussion & Conclusion

In this paper, a machine learning based personalization system has been proposed directly addressing a fundamental issue of hearing aid personalization, namely, that the fine-tuning process should be based *directly* an what the hearing impaired perceives. The proposed personalization system appears to be both fast and robust in finding personalized HA settings, that are significantly preferred over standard prescription based first-fit settings. Hence, the system could possibly be a useful fine-tuning supplement in clinics. The system could easily be extended to support other types of user feedback with the Gaussian process framework, such as rankings [7] or absolute scorings [8], instead of pairwise comparisons, although in general, the latter is probably preferable due to its low cognitive load. The proposed Gaussian process based framework is applicable for other than personalization. By changing the active learning criterion to for instance BALD [9], the framework could be used to *generalize*—in constrast to *optimize*—the latent objective function over all settings.

References

[1] H. Dillon, *Hearing Aids*, Boomerang Press, 2nd edition, 2012.

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

[3] 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.

[4] B. S. Jensen and J. B. Nielsen, "Pairwise Judgements and Absolute Ratings with Gaussian Process Priors," Tech. Rep., November 2011.

[5] J. Vanhatalo and A. Vehtari, "Sparse log gaussian processes via mcmc for spatial epidemiology," in *JMLR Workshop and Conference Proceedings*, 2007, vol. 1, pp. 73–89.

[6] 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.

[7] Wei Chu and Zoubin Ghahramani, "Extensions of gaussian processes for ranking: semi-supervised and active learning," in *NIPS workshop on Learning to Rank*. 2005, pp. 29–34, Citeseer.

[8] B. S. Jensen, J. B. Nielsen, and J. Larsen, "Bounded gaussian process regression," in *IEEE International Workshop on Machine Learning for Signal Processing*, sep 2013.

[9] N. Houlsby, F. Huszár, Z. Ghahramani, and M. Lengyel, "Bayesian Active Learning for Classification and Preference Learning," *ArXiv e-prints*, Dec. 2011.