Perception-based Personalization of Hearing Aids using Gaussian Processes and Active Learning

Jens Brehm Bagger Nielsen, Jakob Nielsen, Jan Larsen, *Senior member, IEEE*

***Abstract*—Personalization of multi-parameter hearing aids involves an initial fitting followed by a manual *knowledge-based* trial-and-error fine-tuning from ambiguous verbal user feedback. The result is an often sub-optimal HA setting whereby the full potential of modern hearing aids is not utilized. This article proposes an interactive hearing-aid personalization system that obtains an optimal individual setting of the hearing aids from direct perceptual user feedback. Results obtained with ten hearing-impaired subjects show that ten to twenty pairwise user assessments between different settings—equivalent to 5-10 min.—is sufficient for personalization of up to four hearing-aid parameters. A setting obtained by the system was significantly preferred by the subject over the initial fitting, and the obtained setting could be reproduced with reasonable precision. The system may have potential for clinical usage to assist both the hearing-care professional and the user.**

***Index Terms*—Hearing Aids, Personalization, Individualization, Gaussian Process (GP), Active Learning, Pairwise Comparisons.**

I. Introduction

THE complexity of digital signal processing algorithms in hearing-aids (HAs) has increased in the past two decades due to continuous refinement of existing HA algorithms and the addition of new ones. Consequently, the number of associated algorithm parameters has increased and will continue to do so in the future. Algorithm parameters control how the incoming sound is processed by the multitude of algorithms and thereby how the sound is presented to the user. In practice, the multi-parameter adjustment—traditionally referred to as *fitting*—is done in fitting software supplied by the HA company: A restricted set of meta parameters is available, that controls the entire set of algorithm parameters. The rules defining the mapping from meta parameters in the fitting software to algorithm parameters are covered by a so-called fitting rationale or prescription. Every HA company has their own fitting rationales for their specific HAs. Typically, generic rationales, such as NAL [1] or DSL [2], are available in the software as an option as well. The overall objective of any fitting rationale is to compensate for the user's reduced ability to hear and comprehend speech. A hearing deficit is typically quantified by measuring the reduction in pure-tone hearing threshold level (HTL) in one-octave frequency bands from 500 Hz to 4 kHz relative to normal hearing (NH) [3, Chapter 10]. A user's *audiogram* refers to the pure-tone HTL difference between the user and NH. The HTL differences are specified in *dB of hearing level* (dB HL) [3, Chapter 10]. A 10 dB HL at 500 Hz indicates that the sound pressure level (SPL) at 500 Hz needs to be 10 dB louder compared to the HTL of a normal hearing subject for the user to detect the pure tone. Hence, an audiogram could directly be converted to HA gains in one-octave frequency band. However, due to loudness recruitment [4, Chapter 4-III] and reduced dynamic range [5, Chapter 1] among other factors, it is inappropriate to set gains directly matching the audiogram [5, Chapter 10]. Instead, the audiogram is used as *target gains* for the fitting, implying that the actual HA gain will not compensate fully for the reduced sensitivity. A rationale converts the target gains (or audiogram) into band-dependent and input-level-dependent (non-linear) HA gains.

HA fitting is carried out by a hearing-care professional (HCP) who measures the audiogram e.g. at, 500 Hz, 1 kHz, 2 kHz, and 4 kHz, which thus results in four target gains for the fitting. The target gains—one *set* of meta-parameters—are used to set algorithm parameters like compression ratio, gain of the linear region, and knee-points of the multi-band dynamic processor embedded in modern digital hearing aids. The goal of the fitting is to ensure audibility and optimal speech intelligibility without compromising the user's preferences.

Besides the measured target gains, there are typically additional meta-parameters that the HCP can or must adjust related to e.g. noise reduction, multi-channel beamforming, further tailoring of the dynamic compressor etc. [5, Chapter 12]. The HCP will consult the hearing-impaired (HI) client about HA use, rehabilitation, and preferences when adjusting meta-parameters; but they can only be adjusted manually based on the user's often ambiguous descriptions about the perceived sound. The HI client typically finds it difficult to explain his preferences towards sound, hence, it is very challenging to determine the best setting. Furthermore, manual fine tuning is time consuming and thus expensive to perform. In summary, this result in an imminent risk of not exploiting the full potential of modern digital HAs. This provides a great potential for new fine-tuning methods or paradigms which aim at optimal settings for individual users in robust and time-efficient manners.

In this paper, a machine-learning based *interactive HA personalization system* (IHAPS) is proposed. IHAPS optimizes

J. B. B. Nielsen is R&D engineer, Audiological Signal Processing at Widex A/S, DK-3560 Lynge, Denmark (e-mail: jeb@widex.com).

J. Nielsen is Senior Audiology Research Engineer, Audiological Signal Processing at Widex A/S, DK-3560 Lynge, Denmark (e-mail: jnl@widex.com).

J. Larsen is Associate Professor at the Section for Cognitive Systems at the Department of Applied Mathematics and Computer Science, Technical University of Denmark, DK-2800, Kongens Lyngby, Denmark (e-mail: janla@dtu.dk).

Manuscript received January, 2014; revised month data, 2013.

multiple parameters based directly on the user's perception of the sound and not based on a derived verbal ambiguous description. By the active user process of listening to and comparing HA settings, IHAPS enables the user to recognize his preference towards the sound. Active engagement also leads to greater psychological ownership, and thereby to better outcome of the entire hearing impairment therapy [6], [7].

In IHAPS, it is assumed that a user's perception is encoded by an unobserved *internal response function* (IRF). Hence, when a user compares two stimuli, the magnitudes of the IRF for the two stimuli determine which of the two stimuli the user prefers or judges to be the best. The IRF cannot be measured directly, and is assumed to be stochastic due to multiple uncontrollable factors. Furthermore, a user's judgments are not fully consistent. Consequently, a user's IRF can only be *estimated* given a set of user assessments of particular stimuli. A particular HA setting, $\mathbf{x}_i$, determines the acoustical stimulus. Hence, the IRF is a function, $f(\mathbf{x})$, of the d HA parameters, $\mathbf{x} = [x_1, ..., x_d]^\top$. Note, that IHAPS can be used both for optimization of meta parameters and of algorithm parameters directly. In the remainder of this article, no distinction between algorithm parameters and meta-parameters will be made. Instead, HA parameters will be considered, which can cover both meta- and algorithm parameters. In IHAPS, the IRF is modeled by a non-parametric Bayesian regression method, viz. a *Gaussian process* (GP) [8], which defines a distribution over flexible nonlinear functions, $f(\mathbf{x})$. Users assess settings in a pairwise-comparison paradigm, whereby users do not need to memorize previous ratings, thus resulting in a reduced cognitive load. However, to minimize the number of assessments required to estimate the user's IRF, the user does not only choose which of two particular HA settings that is preferred (forced choice), but also assesses the *degree* of which the setting is preferred over the alternative [9]. For a given set of such *degree-of-preference* assessments (observations), the distribution of a user's IRF is updated [9] and the setting associated with the largest value of the estimated (mean) IRF is suggested as the optimal setting for the user. Hence, from a modeling perspective, the task is to perform global multi-parameter optimization of the user's unobserved IRF with respect to the d HA parameters, $\mathbf{x}$. In IHAPS, global optimization is performed with minimal number of assessments by use of a sequential design in which *active learning* is used to suggest the next two settings to be compared. In summary, IHAPS sequentially loops the following three steps: (1) active learning to determine the optimal next settings to be compared given the current estimate of the user's IRF; (2) user's assessment of the degree-of-preference between the two compared settings; and (3) update of the user's estimated IRF given all past assessments—including the most recent one. When converged or stopped, the suggested optimal setting is given by the setting that maximizes the estimated IRF.

For demonstrating *solely* the potential of IHAPS, two similar studies are conducted in which HA personalization is performed in the case of two and four parameters, respectively. Preliminary results from the four-parameter study have briefly been described in [10]. Both studies considered a music scenario, because music evokes a user's immediate opinion of

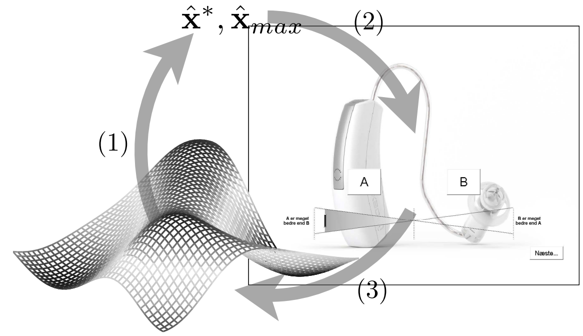

Fig. 1. A conceptual overview of the interactive system. At step (1) a new optimal setting is determined based on the current (probabilistic) estimate of the subject's IRF. Next, at step (2) the optimal setting is compared to the setting which maximizes the current estimate of the subject's IRF, and the subject assesses the *degree of preference* between the two settings using a GUI (see Fig. 5). Finally at step (3), the estimate of the subject's IRF is updated based on the recent assessment.

the quality of the HA-produced sound. Other scenarios, such as a speech scenario, could have been considered as done in [11], but to evaluate IHAPS without several external effects influencing the analysis, the music scenario was considered most suitable. For a real-life application of IHAPS, a multitude of scenarios are relevant including several different stimuli to mimic each scenario. However, these mixed conditions are irrelevant for demonstrating the potential of IHAPS.

Several directions have been pursued for personalization of HAs using for instance a modified simplex procedure [12] or genetic algorithms [13], [14]. However, these initial attempts require unreasonably many assessments to converge, and scale badly with the number of tunable parameters. Almost a decade ago, a probabilistic Bayesian approach was proposed [15], which reads similar to the approach proposed in the present paper. However, two fundamental aspects of the approach in [15] are different: Firstly, it assumes that the user's IRF has a known parameterized functional form, which is difficult to qualify in practice. Secondly, assessments are provided in a pairwise forced-choice paradigm using classical choice models [16], [17]. Using an artificial example, Jensen *et al* [9] show that a forced-choice paradigm requires more assessments than the degree-of-preference paradigm. The non-parametric GP approach using the forced choice paradigm [18] has been considered for instance in [19].

II. Personalization System

IHAPS is based on an interactive loop visualized in Fig. 1. The loop essentially contains three parts: (A) Modeling, (B) active learning, and (C) user-interaction.

A. Modeling of the User's Internal Response Function with Gaussian Processes

Modeling of the user's IRF is performed in a Bayesian non-parametric framework based on GPs, see e.g. [8]. In the

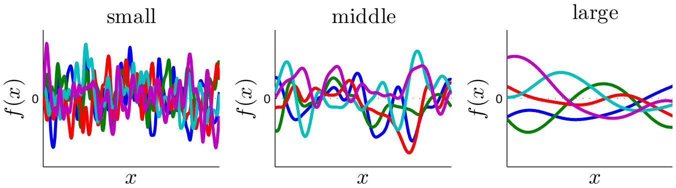

Fig. 2. Examples of sampled functions from a Gaussian process with different setting of the smoothness parameter λ, see Eq. 2. For a more thorough treatment of GP smoothness, see [8, Sec. 2.3 & 2.6, Chap. 5]

following, the different steps of the non-standard GP framework used in IHAPS to perform regression based on *degree-of-preference* assessments are described. The GP framework is based on previous work [9].

1) Gaussian Process Priors: A Gaussian process (GP) is a Bayesian non-parametric regression technique, which defines a prior over functions, $f : \mathbb{R}^d \rightarrow \mathbb{R}, \mathbf{x} \mapsto f(\mathbf{x})$, captured in the notation

$$f(\mathbf{x}) \sim \mathcal{GP}\left(0, k(\mathbf{x}, \mathbf{x}')_{\boldsymbol{\theta}_\mathcal{C}}\right), \tag{1}$$

where $k(\cdot,\cdot)_{\boldsymbol{\theta}_\mathcal{C}}$ is the *covariance function*[1] with parameters $\boldsymbol{\theta}_\mathcal{C}$. Generally speaking, the covariance function defines the smoothness of the functions. A commonly used covariance function is the isotropic *squared exponential* (SE) given by

$$k_{SE}(\mathbf{x}, \mathbf{x}') = \sigma_f \exp\left(-\frac{1}{2\lambda}(\mathbf{x} - \mathbf{x}')^\top(\mathbf{x} - \mathbf{x}')\right). \tag{2}$$

A GP is defined as a *collection of random variables, any finite number of which have a joint Gaussian distribution* [8, Definition 2.1], such that a *finite* collection of function values, $\mathbf{f} = [f(\mathbf{x}_1), ..., f(\mathbf{x}_n)]^\top$, for a corresponding set of inputs, $\mathcal{X} = \{\mathbf{x}_i \in \mathbb{R}^d | i = 1, ..., n\}$, has a distribution given by

$$p(\mathbf{f}|\mathcal{X}, \boldsymbol{\theta}_\mathcal{C}) = \mathcal{N}\left(\mathbf{f}|\mathbf{0}, \mathbf{K}\right), \tag{3}$$

where each entry in the $n \times n$ covariance matrix $\mathbf{K}$ is given by $[\mathbf{K}]_{i,j} = k(\mathbf{x}_i, \mathbf{x}_j)_{\boldsymbol{\theta}_\mathcal{C}}$ and $\mathcal{N}\left(\mathbf{z}|\boldsymbol{\mu}, \boldsymbol{\Sigma}\right)$ denotes the multi-variate normal probability density function[2]. Functions sampled from a GP prior with different settings of the *smoothness* parameter, λ, are depicted in Fig. 2. By specifying the likelihood, $p(\mathcal{Y}|\mathbf{f}, \boldsymbol{\theta}_\mathcal{L})$, of some set of observations, $\mathcal{Y}$, given the finite collection of function values, $\mathbf{f}$, the posterior distribution over the function values $\mathbf{f}$ is given by Bayes formula:

$$p(\mathbf{f}|\mathcal{Y}, \mathcal{X}, \boldsymbol{\theta}) = \frac{p(\mathcal{Y}|\mathbf{f}, \boldsymbol{\theta}_\mathcal{L})p(\mathbf{f}|\mathcal{X}, \boldsymbol{\theta}_\mathcal{C})}{p(\mathcal{Y}|\mathcal{X}, \boldsymbol{\theta})} \tag{4}$$

$$= \frac{p(\mathcal{Y}|\mathbf{f}, \boldsymbol{\theta}_\mathcal{L})p(\mathbf{f}|\mathcal{X}, \boldsymbol{\theta}_\mathcal{C})}{\int p(\mathcal{Y}|\mathbf{f}, \boldsymbol{\theta}_\mathcal{L})p(\mathbf{f}|\mathcal{X}, \boldsymbol{\theta}_\mathcal{C})d\mathbf{f}}, \tag{5}$$

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

2) Likelihood Function: In previous work [9], modeling of continuous bounded responses is performed with a likelihood function based on a re-parameterized beta distribution

[1] In the literature, several expressions are used for the covariance function, such as *kernel function* or simply *kernel*.

[2] In this paper, $\mathbf{z} \sim \mathcal{N}\left(\boldsymbol{\mu}, \boldsymbol{\Sigma}\right)$ and $p(\mathbf{z}) = \mathcal{N}\left(\mathbf{z}|\boldsymbol{\mu}, \boldsymbol{\Sigma}\right)$ are equivalent.

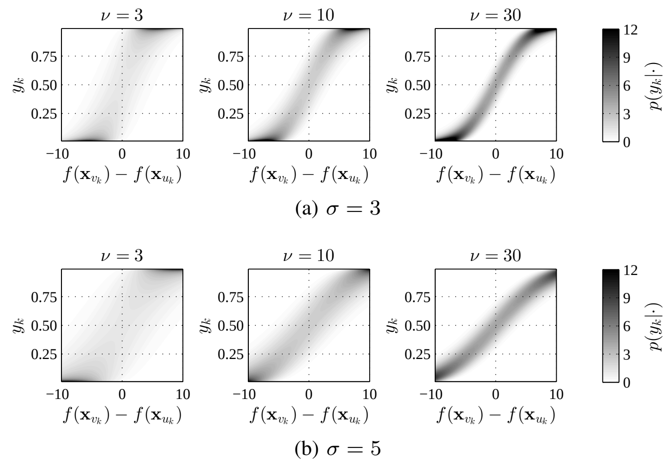

Fig. 3. Visualization of the Beta likelihood $p(y_k|\mathbf{f}_k, \boldsymbol{\theta}_\mathcal{L})$ function for three different settings of the dispersion parameter ν and two different settings of the slope parameter σ.

specifically applicable in cases where observations are pairwise degree-of-preference assessments. Thus, the framework is specifically applicable for the present work.

Progressing as in [9], consider a set of pairwise observations, $\mathcal{Y} = \{y_k \in (0, 1) | k = 1, ..., m\}$, of the *degree of preference* between any two distinct inputs, $u_k, v_k \in \{1, ..., n\}$, implying that $\mathbf{x}_{u_k}, \mathbf{x}_{v_k} \in \mathcal{X}$. An increasing preference for the first option, u_k, is reflected by $y_k \rightarrow 0$, whereas an increasing preference for the second option, v_k, is reflected by $y_k \rightarrow 1$. No preference is indicated by $y_k = 0.5$. A suitable likelihood function, $p(y_k|\mathbf{f}_k)$, is now 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

$$p(y_k|\mathbf{f}_k, \boldsymbol{\theta}_\mathcal{L}) = \text{Beta}\left(y_k|\nu\zeta(\mathbf{f}_k, \sigma), \nu(1 - \zeta(\mathbf{f}_k, \sigma))\right), \tag{6}$$

where $\boldsymbol{\theta}_\mathcal{L} = \{\nu, \sigma\}$ is the set of likelihood parameters. ν is a dispersion parameter around the mean, $\zeta(\mathbf{f}_k, \sigma)$. The mean is defined by

$$\zeta(\mathbf{f}_k, \sigma) = \Phi\left(\frac{f(\mathbf{x}_{v_k}) - f(\mathbf{x}_{u_k})}{\sqrt{2}\sigma}\right), \tag{7}$$

where $\Phi(\cdot)$ is the standard normal cumulative density function—with zero mean and unit variance—and σ is a slope parameter. The likelihood function is visualized in Fig. 3.

By assuming that observations are independent given the latent function values $\mathbf{f}$, the likelihood is written as

$$p(\mathcal{Y}|\mathbf{f}, \boldsymbol{\theta}_\mathcal{L}) = \prod_{k=1}^{m} p(y_k|\mathbf{f}_k, \boldsymbol{\theta}_\mathcal{L}), \tag{8}$$

which is plugged into Eq. (4) together with the GP prior from Eq. (3) to completely specify the Bayesian model.

3) Posterior Inference and Model Training: The Gaussian process model described above is analytically intractable due to the integral in Eq. (5). Therefore, approximate inference is performed based on the Laplace approximation following [9].

The idea of the Laplace approximation [8, Section 3.4] is to approximate the intractable posterior, $p(\mathbf{f}|\mathcal{Y}, \mathcal{X}, \boldsymbol{\theta})$, from Eq. (4) with a Gaussian $q(\mathbf{f}|\mathcal{Y}, \mathcal{X}, \boldsymbol{\theta})$ of the form

$$p(\mathbf{f}|\mathcal{Y}, \mathcal{X}, \boldsymbol{\theta}) \approx q(\mathbf{f}|\mathcal{Y}, \mathcal{X}, \boldsymbol{\theta}) = \mathcal{N}\left(\mathbf{f}|\hat{\mathbf{f}}, \mathbf{A}^{-1}\right), \tag{9}$$

where $\hat{\mathbf{f}}$ is the posterior maximum (mode) and $\mathbf{A}$ is the Hessian of the negative log posterior at the mode. The mode is found by maximizing the unnormalized log-posterior given by

$$\psi\left(\mathbf{f}|\mathcal{Y},\mathcal{X},\boldsymbol{\theta}\right) = \log p\left(\mathcal{Y}|\mathbf{f},\boldsymbol{\theta}_{\mathcal{L}}\right) - \frac{1}{2}\mathbf{f}^{\top}\mathbf{K}^{-1}\mathbf{f} - \frac{1}{2}\log|\mathbf{K}| - \frac{n}{2}\log 2\pi, \quad (10)$$

with a Newton method. The Newton step is given by

$$\mathbf{f}^{new} = \left(\mathbf{K}^{-1} + \mathbf{W}\right)^{-1}\left[\mathbf{W}\mathbf{f} + \nabla\log p(\mathcal{Y}|\mathbf{f},\boldsymbol{\theta}_{\mathcal{L}})\right], \quad (11)$$

where $[\mathbf{W}]_{i,j} = -\sum_{k=1}^{m}\nabla\nabla_{i,j}\log p\left(y_k|\mathbf{f}_k,\boldsymbol{\theta}_{\mathcal{L}}\right)$ defining $\nabla\nabla_{i,j} \equiv \frac{\partial^2}{\partial f(\mathbf{x}_i)\partial f(\mathbf{x}_j)}$. Note, that unlike traditional classification and regression problems, $\mathbf{W}$ is not diagonal due to the pairwise structure. For derivatives and further details, see [20].

When Eq. (11) has converged, the approximation is simply

$$p(\mathbf{f}|\mathcal{Y},\mathcal{X},\boldsymbol{\theta}) \approx q(\mathbf{f}|\mathcal{Y},\mathcal{X},\boldsymbol{\theta}) = \mathcal{N}\left(\mathbf{f}|\hat{\mathbf{f}},\left(\mathbf{W}+\mathbf{K}^{-1}\right)^{-1}\right) \quad (12)$$

Traditionally, training of GPs are performed by optimizing the marginal likelihood, $p(\mathcal{Y}|\mathcal{X},\boldsymbol{\theta})$, from Eq.(4) with respect to the hyper parameters, $\boldsymbol{\theta}$. This is referred to as ML-II optimization [8, Chapter 5.2]. In the present paper, a slightly different scheme is used, in which the optimization is regularized by hyper priors, $p(\boldsymbol{\theta})$, over the parameters in what is a maximum-a-posterior-like (MAP-II) scheme following [8, Chapter 5.2]. More precisely, the parameters $\boldsymbol{\theta}^{\text{MAP-II}}$ in the trained GP are given by

$$\begin{aligned}\boldsymbol{\theta}^{\text{MAP-II}} &= \arg\max_{\boldsymbol{\theta}}\log q(\boldsymbol{\theta}|\mathcal{Y},\mathcal{X}) \quad (13)\\ &\approx \arg\max_{\boldsymbol{\theta}}\log p(\boldsymbol{\theta}|\mathcal{Y},\mathcal{X}) = \arg\max_{\boldsymbol{\theta}} p(\boldsymbol{\theta}|\mathcal{Y},\mathcal{X}),\end{aligned}$$

where the intractable log posterior over the parameters, $\log p(\boldsymbol{\theta}|\mathcal{Y},\mathcal{X})$, is approximated by $\log q(\boldsymbol{\theta}|\mathcal{Y},\mathcal{X})$, which is—up to a normalization constant—given by

$$\log q(\boldsymbol{\theta}|\mathcal{Y},\mathcal{X}) \propto \log q(\mathcal{Y}|\mathcal{X},\boldsymbol{\theta}) + \log p(\boldsymbol{\theta}). \quad (14)$$

In Eq. (14), $q(\mathcal{Y}|\mathcal{X},\boldsymbol{\theta})$ is the Laplace approximation to the (intractable) marginal likelihood, $p(\mathcal{Y}|\mathcal{X},\boldsymbol{\theta})$, from Eq. (4), resulting in

$$\log q(\boldsymbol{\theta}|\mathcal{Y},\mathcal{X}) \propto \log p(\mathcal{Y}|\hat{\mathbf{f}},\boldsymbol{\theta}_{\mathcal{L}}) - \frac{1}{2}\hat{\mathbf{f}}^{\top}\mathbf{K}^{-1}\hat{\mathbf{f}} - \frac{1}{2}\log|\mathbf{I}+\mathbf{K}\mathbf{W}| + \log p(\boldsymbol{\theta}). \quad (15)$$

Hence, training of the GP model consists of the following two steps which are looped until convergence[3]:

1: With fixed hyper parameters, $\boldsymbol{\theta}^{\text{MAP-II}}$, repeat Eq. (11) to find the mode of the Laplace approximation and use Eq. (12) to approximate the posterior.

2: Given the approximate posterior from Eq. (12), optimize the right hand side of Eq. (15) with respect to $\boldsymbol{\theta}$ using a BFGS gradient method to obtain the hyper parameters $\boldsymbol{\theta}^{\text{MAP-II}}$.

The specific choices of kernel, $k(\cdot,\cdot)_{\boldsymbol{\theta}_{\mathcal{C}}}$, and hyper priors, $p(\boldsymbol{\theta})$ are specific to each experiment, and are therefore explained as part of Sec. III and Sec. IV.

4) Predictions: The overall goal of almost any machine learning algorithm once trained, is to make prediction for future inputs. In a regression setting, predictions consist of predicting the function values, $\mathbf{f}_* = [f(\mathbf{x}_1^*),...,f(\mathbf{x}_o^*)]^{\top}$, at new input locations, $\mathcal{X}_* = \{\mathbf{x}_l^* \in \mathbb{R}^d | l = 1,...,o\}$. In a Bayesian framework, the variables are considered stochastic. Therefore in a Bayesian framework, predictions are formulated in terms of the predictive distribution, $p(\mathbf{f}_*|\mathcal{Y},\mathcal{X}_*)$, from which different statistics about the variables can be accessed, such as the mean value, $\boldsymbol{\mu}^*$, and (co)variance, $\boldsymbol{\Sigma}^*$.

Given the GP, the joint prior distribution between the predictive and training function values is given by

$$\begin{bmatrix}\mathbf{f}\\ \mathbf{f}_*\end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix}\mathbf{0}\\ \mathbf{0}\end{bmatrix},\begin{bmatrix}\mathbf{K} & \mathbf{K}_*\\ \mathbf{K}_*^{\top} & \mathbf{K}_{**}\end{bmatrix}\right), \quad (16)$$

where $[\mathbf{K}_{**}]_{l,r} = k(\mathbf{x}_l^*,\mathbf{x}_r^*)$ and $[\mathbf{K}_*]_{i,l} = k(\mathbf{x}_i,\mathbf{x}_l^*)$. Given Eq. (16), the conditional distribution of $\mathbf{f}_*|\mathbf{f}$ is Gaussian, hence

$$\begin{aligned}p(\mathbf{f}_*|\mathcal{Y},\mathcal{X},\mathcal{X}_*,\boldsymbol{\theta}) &= \int p(\mathbf{f}_*|\mathbf{f},\mathcal{X},\mathcal{X}_*,\boldsymbol{\theta})q(\mathbf{f}|\mathcal{Y},\mathcal{X},\boldsymbol{\theta})d\mathbf{f} \quad (17)\\ &= \mathcal{N}\left(\mathbf{f}_*|\boldsymbol{\mu}_*,\boldsymbol{\Sigma}_*\right) \quad (18)\end{aligned}$$

which is an integral over the product of two Gaussian distribution, which is again Gaussian. The solution is found in for instance [18, Eq. (17)-(18)] and is given as

$$\boldsymbol{\mu}_* = \mathbf{K}_*^{\top}\mathbf{K}^{-1}\hat{\mathbf{f}} \quad (19)$$

$$\begin{aligned}\boldsymbol{\Sigma}_* &= \mathbf{K}_{**} - \mathbf{K}_*^{\top}\left(\mathbf{K}+\mathbf{W}^{-1}\right)^{-1}\mathbf{K}_*\\ &= \mathbf{K}_{**} - \mathbf{K}_*^{\top}\left(\mathbf{I}+\mathbf{W}\mathbf{K}\right)^{-1}\mathbf{W}\mathbf{K}_* \quad (20)\end{aligned}$$

where the last expression is not given in [18], but is numerically more stable, because it avoids inverting $\mathbf{W}$. Eq. (19) is used as the estimator of the user's IRF. In addition, the covariance from Eq. (20) is utilized to formulate an active learning criterion used to select the next input $\hat{\mathbf{x}}^*$ actively, to constitute the next $(k+1)$ comparison.

Prediction of preference relations, y_*, can be done but is not of particular interest in the present paper, see further [20].

B. Sequential Design for Optimization

Sequential design (or active learning) is used to reduce the required number of training examples by sequentially including new *informative* training examples with respect to some criterion[4]. Traditionally, active learning is used when labeling of data is expensive and is done sequentially.

As for most machine learning algorithms, typically, sequential designs aim at maximizing the generalization performance of a model often formulated in terms of a specific criterion. A Bayesian criterion is the expected reduction in posterior Shannon entropy after inclusion of a new training example [21]. In this work, the generalization performance is not of particular importance. Instead, the aim is to find a maximum—ideally the global one—of the unknown IRF modeled by the GP. For this, a novel bivariate version of the *expected improvement* [22] is used. Expected improvement (EI) is derived by defining

[3] The similarity with the Expectation-Maximization algorithm is that step 1 can be recognized as the E-Step, and step 2 as the M-step.

[4] Some active learning methods, such as query-by-committee, do not have an explicit criterion, but this is beyond the scope of the present article to discuss.

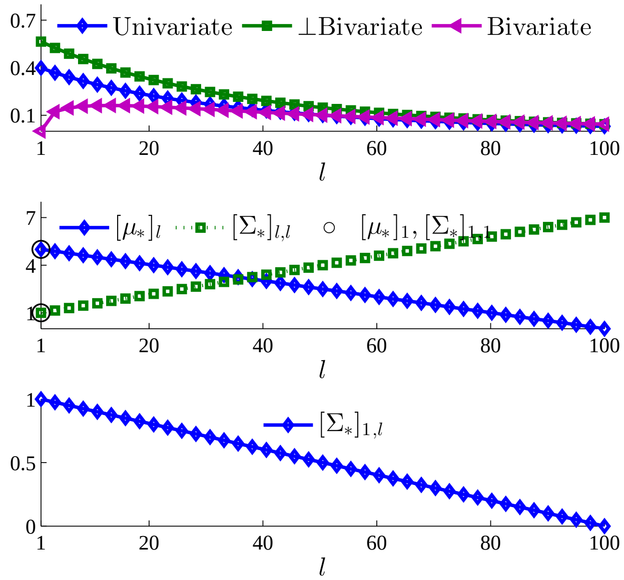

Fig. 4. Illustration of the difference between the univariate and bivariate EI. The current maximum (indicated by circles) is at $l = 1$, which is also a possible query. **Top**: EI for the standard version (Univariate), a bivariate version neglecting covariance (⊥Bivariate) and a full bivariate version incorporating covariance (Bivariate). **Middle**: mean and variance for query points, $\mathbf{x}_l^*$. **Bottom**: covariance between query $\mathbf{x}_l^*$ and maximum $\mathbf{x}_1$ used in the (full) bivariate EI. Note that the full bivariate EI is zero at the current maximum; hence, avoids querying this point again.

improvement, I, as the difference in function values between the current maximum $\hat{f} \equiv f(\hat{\mathbf{x}}_{\text{max}})$ (typically only among the training cases $\mathcal{X}$) and a query point $f_l^* \equiv f(\mathbf{x}_l^*)$

$$I \equiv f_l^* - \hat{f}. \tag{21}$$

Now, EI is the expectation of the (positive) improvement (which is normally distributed in the present model)

$$\begin{aligned} EI &\equiv \mathbb{E}_{p(I)}\{\max(I,0)\} = \int_0^\infty Ip(I)dI \\ &= \int_0^\infty I\mathcal{N}\left(I|\mu_I, \sigma_I^2\right) dI \\ &= \sigma_I \mathcal{N}\left(\frac{\mu_I}{\sigma_I}\middle| 0, 1\right) + \mu_I \Phi\left(\frac{\mu_I}{\sigma_I}\right) \end{aligned} \tag{22}$$

In standard EI, $\hat{f}$ is not considered stochastic, hence the distribution $p(I)$ is just a univariate normal with mean $\mu_I = [\boldsymbol{\mu}_*]_l - \hat{f}$ and variance $\sigma_I^2 = [\boldsymbol{\Sigma}_*]_{l,l}$ [22]. In this article, the joint distribution between the query and maximum is taken into consideration. In this case, $p(I)$ is the difference between two dependent normal distributed random variables, and is thus given by

$$\begin{aligned} p(I) &= \mathcal{N}\left(I|\mu_I, \sigma_I^2\right), \quad \text{where} \\ \mu_I &= [\boldsymbol{\mu}_*]_l - [\boldsymbol{\mu}_*]_{\text{max}} = [\boldsymbol{\mu}_*]_l - \hat{f} \\ \sigma_I^2 &= [\boldsymbol{\Sigma}_*]_{l,l} + [\boldsymbol{\Sigma}_*]_{\text{max},\text{max}} - 2 \cdot [\boldsymbol{\Sigma}_*]_{l,\text{max}}. \end{aligned} \tag{23}$$

The difference between the univariate and bivariate EI, i.e., whether to including the covariance between the query and the maximum (last term in Eq. (23)), is illustrated in a small example in Fig. 4. In this example, the current maximum point corresponds to $l = 1$ and has larger mean value than all other query points ($l \neq 1$), but smaller variance. This is a typical scenario in GP modeling. In this scenario, neglecting the covariance has the undesirable effect of querying the already observed maximum point, causing the active learning to "get stuck". The (full covariance) bivariate EI avoids this, but has the same properties when maximum and query points are independent. For GP models, predictions are typically very dependent when inputs are close to each other. Hence, the bivariate version is not as local as the standard univariate EI. In the following, EI refers to the full bivariate version.

A user's optimal setting is essentially unknown. Therefore, it is not possible to measure how close an optimal setting suggested by IHAPS is to the true optimal setting. However, the average EI over possible queries, $\mathbf{x}_l^*$, is in IHAPS used as a prediction of convergence (convergence measure). Intuitively, when the average EI is zero or close to zero, no further improvement is to be expected from another setting. Thereby, no other setting is under the predictive distribution expected to be preferred over the current optimal setting, $\hat{\mathbf{x}}_{\text{max}}$.

Fig. 5. The pairwise graphical user interface (GUI) used for the experiments. A slider is used to capture the *degree-of-preference* for either setting '1' or '2'. The user can listen to setting '1' or '2' by pressing the corresponding button. A gray button indicates that the corresponding setting is selected and thus active in the HAs. When the user is satisfied with the position of the slider, the button in the lower-right corner 'Nste' is pushed to confirm the current assessment. Next IHAPS computes the next comparison with two new settings corresponding to '1' and '2' until a prescribed number of iterations is reached.

C. Graphical User Interface

The graphical user interface (GUI) by which the users interact with IHAPS during experiments is depicted in Fig. 5. An important property of the GUI is that users shall generally find it intuitive and easy to use. Therefore, the placements of buttons and sliders are arranged to indicate the pairwise nature of the assessments. The slider is designed as a *mirrored volume control* to indicate that the preference is increasing towards the end points of the slider.

III. Study 1: Two-Dimensional Optimization

In the first study the subjectively best target gains of the HA fitting for the four basic frequency bands—500 Hz, 1 kHz, 2

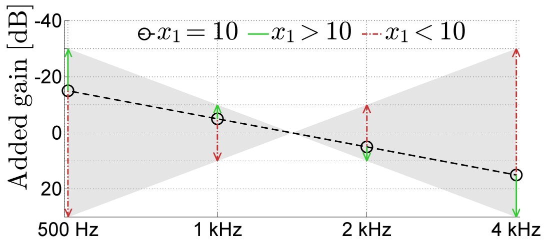

(a) Contribution of x_1

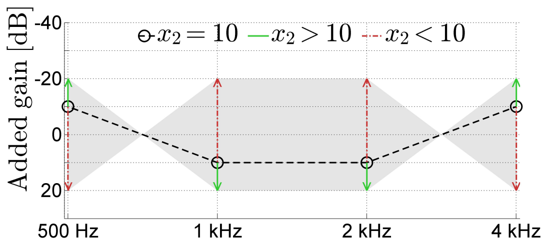

(b) Contribution of x_2

Fig. 6. Two meta-parameters, x_1 and x_2, are used to modify the target gains in the four basic frequency bands of the HA fitting, as shown in (a) and (b), respectively. Note, that a particular subjects resulting target gains (in dB) for the four frequency bands is the sum of the measured audiogram (binaural) and the added gains specified by the selection of $x_1 \in \{-20; 20\}$ and $x_2 \in \{-20; 20\}$. The grey shaded areas show the added gain limits when varying the meta-parameters in their intervals.

kHz, and 4 kHz—were learned indirectly by modifying the target gains[5] with two meta-parameters $x_1, x_2 \in \{-20; 20\}$. The meta-parameters were learned while the subject listened to a 32 sec. looping music clip[6]. The parametrization is visualized in Fig. 6. For a specific parametrization, x_1 and x_2, the resulting target gains (in dB) for the four frequency bands are computed as the sum of the two sets of added gains in Fig. 6 (a) and (b) and the measured audiogram[7]. Together x_1 and x_2 define how the audiogram is shaped for the particular piece of music, and IHAPS was used to obtain optimal shapes for the individual subject as quickly as possible.

A. Algorithm Details

Although the described framework from Sec. II is highly generic, there are still a few properties left to define. The modeling part was set up by defining the covariance function and hyper-prior distributions:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f \exp\left(-\frac{1}{2}(\mathbf{x}-\mathbf{x}')^\top \mathbf{P}^{-1}(\mathbf{x}-\mathbf{x}')\right), \quad (24)$$

$$\text{with } \mathbf{P} = \text{diag}([\lambda_1, ..., \lambda_d]^\top), \quad (25)$$

$$\sigma_f \sim p([\boldsymbol{\theta}_\mathcal{C}]_{d+1}) = \delta(\sigma_f = 4), \quad (26)$$

$$\lambda_i \sim p([\boldsymbol{\theta}_\mathcal{C}]_i) \;= \text{half-St}(\lambda_i; 6, 10), \quad (27)$$

$$\sigma \sim p([\boldsymbol{\theta}_\mathcal{L}]_1) = \text{half-St}(\sigma; 6, 10), \quad (28)$$

$$\nu - 2 \sim p([\boldsymbol{\theta}_\mathcal{L}]_2) = \text{half-St}(\nu - 2; 6, 10), \quad (29)$$

where half-St$(z; \xi, s) \propto \left(1 + \frac{1}{\xi}\left(\frac{z}{s}\right)^2\right)^{-(\xi+1)/2}$ is the half Student's t-distribution [23], [24], [25] with ξ degrees of freedom and scale s. The above kernel and hyper-prior distribution are common choices, see e.g. [8], [23]. The hyper parameters were learned by optimizing Eq. (15) using a gradient ascend method with initial values $\sigma_f = 4, \lambda_i = 5, \sigma = \exp(1), \nu = 2 + \exp(1)$. The parameters of the hyper-prior distributions and the initialization of the hyper parameters were not tuned to perfection, but were set from a few initial experiments with normal-hearing subjects. It turns out that the framework is not overly sensitive to this tuning.

For the active learning part, the EI from Eq. (22) was not directly maximized. Instead, the EI was calculated for all possible $\mathbf{x}_l^*$ in a grid and collected in $\mathbf{EI}$ such that $[\mathbf{EI}]_l$ contained the EI for $\mathbf{x}_l^*$. The evaluation of the EI for the entire grid was computationally feasible, since $d = 2$ is small. A uniform grid from -20 to 20 with a step size of 1 was used for both x_1 and x_2, hence $\mathcal{X}_* = \{[-20:1:20]^2\}$. Now, the index $\hat{l}$ of the setting $\hat{\mathbf{x}}_l^* \in \mathcal{X}_*$ to add to the next comparison was determined by once drawing a vector $\boldsymbol{\ell}$ of length $o = 41^2$ of binary variables with exactly one nonzero component from a multinomial distribution given by

$$\boldsymbol{\ell}|\mathbf{EI} \sim \text{Mult}\left(\frac{1}{\sum_{l'=1}^{o}[\mathbf{EI}]_{l'}} \cdot \mathbf{EI}\right). \quad (30)$$

The index $\hat{l}$ was thus given by the index of the nonzero component of $\boldsymbol{\ell}$. Compared to maximization of the EI, a little randomness[8] was introduced. Recall, that the bivariate EI compared to the univariate EI avoids querying the current maximum over and over again for the entire test. The extra randomness imposed by the multinomial sampling avoids querying settings *too* close to the current maximum too often towards the end of each test.

B. Procedure

Every iteration consisted of a comparison between the actively sampled new setting, $\hat{\mathbf{x}}_l^*$, and the current best setting, $\hat{\mathbf{x}}_{\text{max}}$, among the training set, $\mathcal{X}$. To remove a possible bias effect, the two settings were randomly assigned to option 1 and 2. A single test consisted of 30 iterations/comparisons which were the desired maximal number of iterations to achieve an optimal setting. Two tests, *Test 1* and *Test 2*, were conducted to show the reproducibility of the found optima. Prior to the two tests, the subjects rated 10 comparisons between randomly chosen settings. This *training session* was used only to give the subjects an opportunity to learn how to use the setup and how the sound in the HAs varied. Following the two tests, a *significance test* was conducted to investigate if the optimal setting of IHAPS was significantly preferred by the subject over a baseline setting. The significance test used twenty repeated forced-choices between the optimal setting and the baseline setting. In each repetition they were assigned randomly to the two options presented to the subject. Significance was tested with an exact two-tailed binomial test. The optimal setting was

[5]In the WIDEX® fitting software the target gains are set in what is called a SENSOGRAM, see http://www.widex.pro/en/fitting-systems/compass/in-situ-tools/sensogram/.

[6]Teitur, "Sleeping with the Lights on", Poetry & Aeroplanes, 2003. Start at sec. 6. End at sec. 38.

[7]The audiogram is measured binaurally, hence the left and right ears are fitted individually.

[8]Typically, finding the best trade-off between exploitation (utilize the model) and exploration (reduce uncertainty) is the main challenge in active learning.

taken from *Test 2* unless this test did not converge. In that case the optimal setting was taken from *Test 1*. A natural baseline setting is the setting with target gains equal to the subject's audiogram (i.e., $x_1 = 0$, $x_2 = 0$), since this is the standard setting of the HAs without additional personalization. Settings were automatically uploaded to the HAs using proprietary WIDEX® software.

WIDEX® PASSION440 HAs equipped with RIC 1-Receivers were used in all tests with all subjects. CRET-S soft earmolds (without vent) were constructed individually to each subject to obtain a closed fitting. The HAs were fitted initially using the measured audiogram, with an omni-directional beamformer, noise reduction and speech enhancement turned off and slow-acting less-aggressive feedback cancellation (FBC) particular suitable for music[9].

To avoid placebo effects in the final significance test, subjects were not informed that the aim of the experiments was to optimize the setting of the HAs based on their feedback. Instead, they were informed that they, in a sequence of pairwise comparisons between different settings in the HAs, should judge which settings they preferred and how much. It was emphasized that the judgments should only reflect their subjective opinion. Likewise, subjects were not primed to focus or pay attention to specific things in the music.

C. Results

TABLE I
OPTIMAL PARAMETER SETTINGS (x_1, x_2) FOR TEST 1 AND TEST AND CORRESPONDING SIGNIFICANCE LEVELS.

Subject	Age	Test 1	Test 2	$p_0 <$
#1	55	NC	(−20,−20)	**0.001**
#2	58	(−20,−16)	(−16,−12)	**0.001**
#3	57	(−16,−16)	(−14,−16)	**0.001**
#4	71	(−20,−12)	(−16,−8)	**0.001**
#5	66	(−18,−14)*	(−18,−2)*	**0.001**
#6	77	NC	NC	NC
#7	45	NC	NC	NC
#8	45	(0,−14)	(−4,−12)	**0.001**
#9	35	(0,−18)	(−8,−10)*	**0.001**
#10	53	(−18,−14)	(−20,−6)	**0.001**

NC: Not converged, average EI is clearly non-zero.
*: Average EI not completely zero.

The best settings found in the two consecutive tests and the results of the significance tests are shown in Table I. NCs indicate tests that did not converge according to the average EI convergence measure. Asterisk symbols denote tests in which the convergence measure did not completely reach zero (see Fig. 7a). The optimal settings (from converged tests) transformed to actual target-gains shapings are shown in 7b.

Fig. 8 and Fig. 9 show the IRF predictions and the EI after 30 and 16 iterations, respectively, for subject 4 from Test 1 and Test 2. Additional details are found in [26].

D. Discussion

During the tests, some interesting observations were made. Firstly, subject 6 clearly was not able to consistently distinguish between different settings, which is also reflected in the convergence measure (see Fig. 7a). Secondly, subject 7 expressed that he was in conflict with himself during the experiments. Sometimes he preferred a more richer but resounding sound, while other times he preferred a more flat and neutral sound. Unfortunately, he was unable to make up his mind and switched several times between the two types of listening strategies. Consequently, IHAPS found him to be inconsistent and did not converge. Thirdly, due to a numerical issue, the active criterion was not working properly in the last part of the second test with subject 9. Therefore, this test did not completely convergence to zero. This is somewhat misleading, because effectively, the last part of the examples presented to subject 9 in the second test was chosen randomly due to the numerical issue. Indirectly, IHAPS thus refrained from *optimization* and performed *generalization* instead. Without the numerical issue, the second test with subject 9 probably would have converged completely based on her behavior from the first test (see Fig. 7a).

Generally, IHAPS was able to obtain a personalized setting of the two meta-HA parameters for eight of the ten hearing-impaired subjects. Obviously, if a user does not have a consistent preference or is unable to distinguish any settings from each other, IHAPS cannot and shall not obtain an optimal setting. Ideally, IHAPS should be able to identify when it has not

[9]This FBC setting is obtained in the WIDEX® software with the "Super-Gain Music" setting of the FBC

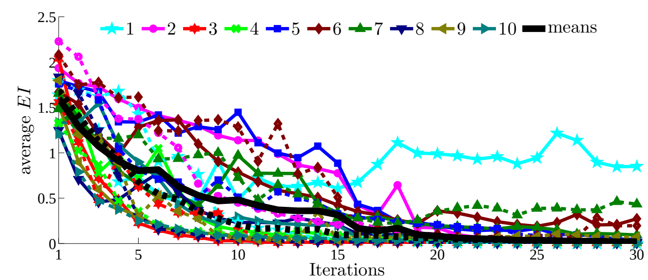

(a) Convergence measure

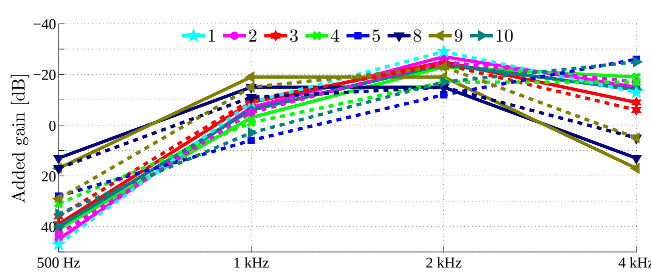

(b) Optimal target-gains shapings

Fig. 7. Solid lines correspond to Test 1 and dotted lines correspond to Test 2. (a) The convergence measure for each subject calculated as $\sum_{l'=1}^{o} [\mathbf{EI}]_{l'}$ plotted as a function of iterations and means across subjects that converged (see Tab. I). (b) Added target-gains shapings given the optimal parameters for the tests that converged (see Table. I).

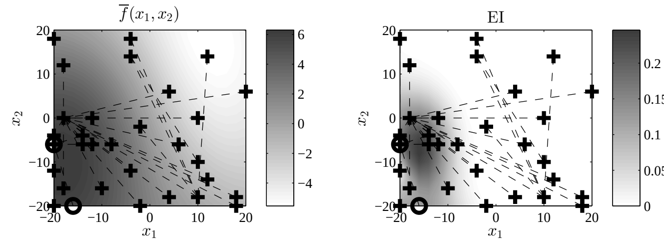

(a) Test 1

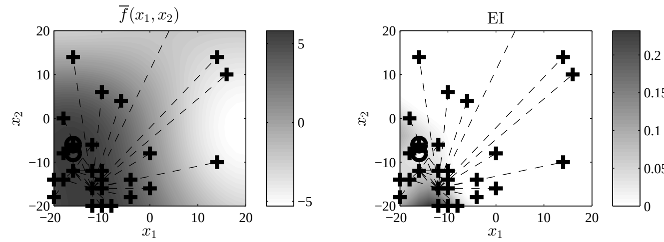

(b) Test 2

Fig. 8. **Reproducibility:** Predictions of the IRF (left figures) and the EI (right figures) for subject 4 after the 30'th (final) iteration from (a): test 1 and from: (b) test 2. Crosses indicate observations, dotted lines indicate comparisons and circles show the suggested next comparison (although the test stops at this point).

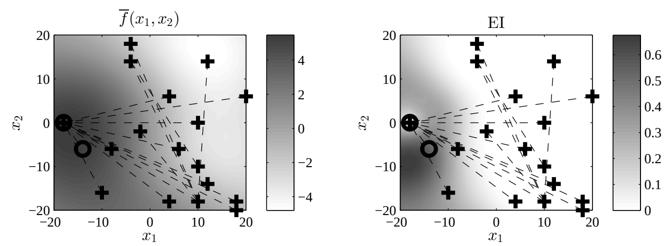

(a) Test 1

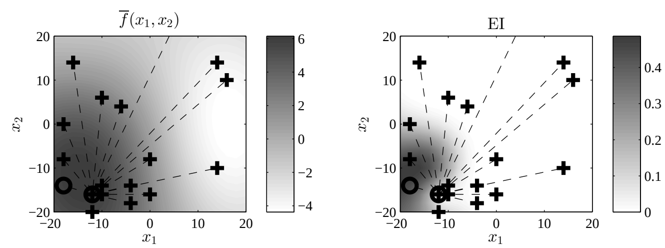

(b) Test 2

Fig. 9. **Convergence:** Predictions of the IRF (left figures) and the EI (right figures) for subject 4 after the 16'th iteration from (a): test 1 and from: (b) test 2. Crosses indicate observations, dotted lines indicate comparisons and circles show the suggested next comparison.

obtained an optimal setting for the user, i.e., that a session does not converge. Since, each user's true optimum is unknown, only well-founded speculations can be made. Nevertheless, for all tests that converged according to the average EI, IHAPS suggested a setting that the subject preferred significantly over the prescribed setting. Furthermore, the session of subject 7 did not converge and he did not prefer the suggested setting over the baseline. The comments given after the test by subject 7 explain why the system was unable to obtain an optimal setting for subject 7, as the system cannot deal with subjects that change their opinion during a test. It is speculated that if subject 7 had indicated that two (or even a couple) of settings were equally good, the test would have been successful. This might have been achieved with more thorough instructions about what the test was actually about. For instance, subject 7 could have been instructed to intentionally stick with only one objective at the time, instead of switching between them during a single test. However, this would have biased the results. Nevertheless, it is desirable that IHAPS by the average EI seems to indicate if a test successfully obtains a (near) optimal setting, even though the average EI appears to be somewhat conservative (indicated by asterixes in Fig. 7a).

The reproducibility is actually better than what can be concluded from inspection of the suggested optimal settings only, indicated by studying the close resemblance between the predicted IRF of the two tests (shown for subject 4 in Fig. 8a and Fig. 8b). Furthermore, by comparing Fig. 9 and Fig. 8 it is observed that the IRF for subject 4 was already quite well captured halfway through both tests. Generally, this was common to all successful tests and illustrates that the average EI is somewhat conservative for predicting of convergence.

IV. Study 2: Four-Dimensional Optimization

The setup for the second study was similar to the first study described in Sec. III. However, instead of modifying the four target gains by two meta parameters, all four target gains were in both hearing aids defined directly ranging from 0 to 80 dB HTL (hearing threshold level) in 5 dB HTL steps, hence $\mathbf{x} = [x_1, ..., x_4]^\top$ with $x_i \in \{0, 5, 10, ..., 80\}$. Note, that this setup did not account for any differences between a subject's two ears, but all subjects had a similar hearing loss on both ears. The main purpose of this study was to compare the performance of IHAPS in a four-dimensional scenario with that of the two-dimensional scenario in terms of reproducibility and convergence.

A. Algorithm details

The model was defined similar to the model used in Sec. III, except that the scale of the half Student's t-distribution for the length-scale parameters, λ_i, was changed since the range of each dimension x_i was different. The model was defined as

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f \exp\left(-\frac{1}{2}(\mathbf{x}-\mathbf{x}')^\top \mathbf{P}^{-1}(\mathbf{x}-\mathbf{x}')\right), \quad (31)$$

$$\text{with } \mathbf{P} = \mathrm{diag}([\lambda_1, ..., \lambda_d]^\top), \quad (32)$$

$$\sigma_f \sim p([\boldsymbol{\theta}_\mathcal{C}]_{d+1}) = \delta(\sigma_f = 4), \quad (33)$$

$$\lambda_i \sim p([\boldsymbol{\theta}_\mathcal{C}]_i) = \text{half-St}(\lambda_i; 6, 100), \quad (34)$$

$$\sigma \sim p([\boldsymbol{\theta}_\mathcal{L}]_1) = \text{half-St}(\sigma; 6, 10), \quad (35)$$

$$\nu - 2 \sim p([\boldsymbol{\theta}_\mathcal{L}]_2) = \text{half-St}(\nu - 2; 6, 10). \quad (36)$$

Again, the hyper parameters were learned by optimizing Eq. (15) using a gradient ascent method with initial values $\sigma_f = 4, \lambda_i = 30, \sigma = \exp(1), \nu = 2 + \exp(1)$.

For the active learning part, evaluating the EI for all possible input values was computational intractable in this four-dimensional scenario. Instead, the setting, $\hat{\mathbf{x}}^*$, to constitute the next comparison was found directly by maximizing the EI with respect to the input, $\mathbf{x}_l^*$, with a BFGS gradient ascent method [27]. Five uniformly-random starts of the initial value of $\mathbf{x}_l^*$ were used for the gradient ascent method. With only five random starts, the global maximum of the EI is generally not discovered. This creates a similar effect as in Sec. III, although it was achieved differently. Likewise, the average of the EI could not be computed in a reasonable four-dimensional grid. Instead, the average EI along the path of the gradient ascend method was used as an estimate of the true average. A single estimate can be very different from the true average. If for instance the initialization of the gradient ascent method is close to the maximum, the estimate is much larger than the true average. To remove some of this variance, a 4-block running average was used to smooth the convergence.

B. Procedure

The procedure was identical to the two-dimensional study described in Sec. III-B, except that the baseline setting was directly the measured audiogram.

Due to practical circumstances not all ten subjects from the first study were able to participate in the second study. Therefore, four new test subjects participated. Only subject 6 was deliberately not considered for the second study, since she was clearly unable to distinguish between different settings in the first study. It was considered not to include subject 7 either, due to the results in the first study. However, apparently subject 7 did not have difficulties distinguishing settings, but was only in doubt of what he preferred. Hence as such, subject 7 constitutes an interesting case.

C. Results

The found best settings in the two consecutive tests and the significance-test results are shown in Fig 10a. The convergence is shown in Fig. 10b. Generally, nine of ten subjects obtained a setting that was significantly preferred over the baseline setting given by the user's audiogram. Subject 7 neither preferred the obtained setting nor the baseline setting significantly. The setting resulting from Test 1 - instead of Test 2 - was used for the significance tests for subject 11 and 13. The reason is that the two obtained settings were found to be very different from each other. Furthermore, the two subjects reported, on their own initiative, that the settings presented to them in the second test were in general noticeably worse than the settings from the first test —even at the end of the session, where at least one of the settings should have been good.

From Fig. 10b, two runs—test 1 with subject 7 and test 2 with subject 11—are seen not to have converged by the 30th iteration. Overall, the estimation of average EI is more noisy, which makes it more unclear if particular test converged.

In Fig. 11, the long-term power spectra of the SPL at the eardrum generated by the (left) HA are shown for the three different settings (Test 1, Test 2, Audiogram) for five subjects (see sub-figure caption). The measurements were made on a KEMAR through a GRAS IEC711 coupler.

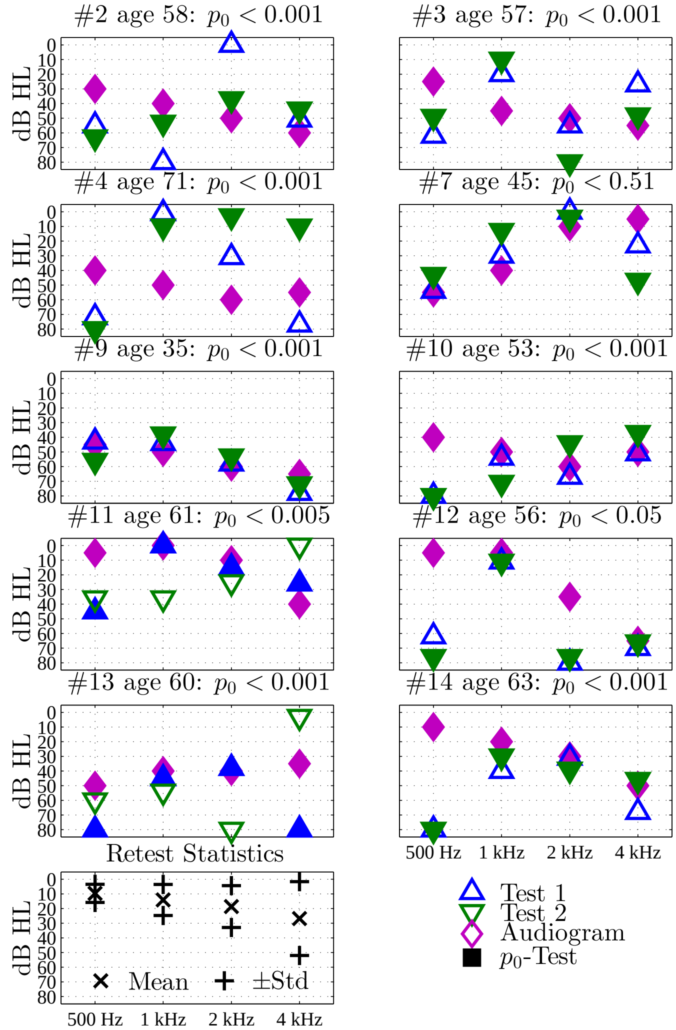

(a) Optimal settings

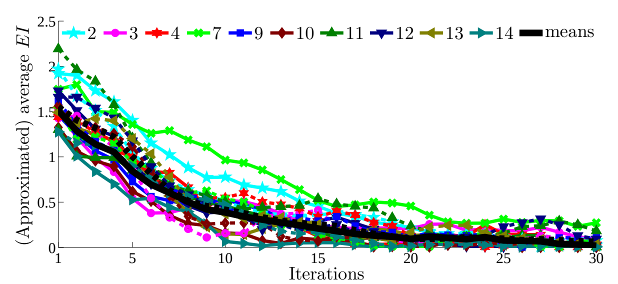

(b) Convergence measure

Fig. 10. (a) Optimal target-gains settings found in Test 1 (△) and Test 2 (▽) together with the measured audiogram (◊). Filled markers indicate the settings used in the significance tests (■). The bottom left plot shows the mean (×) and standard deviation (+) of the found parameter difference between test 2 and test 1. (b) Estimated convergence measure for each subject and the mean convergence over subjects using only the tests that converged (i.e., excluding Test 1 with subject 7 and Test 2 with subject 11).

D. Discussion

Generally, it is satisfying that the only subject (subject 7), that did not have a significant preference for the setting obtained with IHAPS, actually obtained two settings which are almost identical to the baseline—both as regards the param-

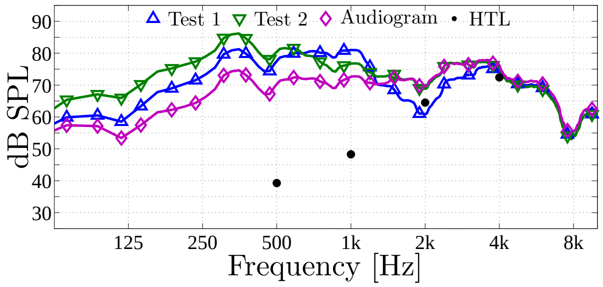

(a) Subject 2

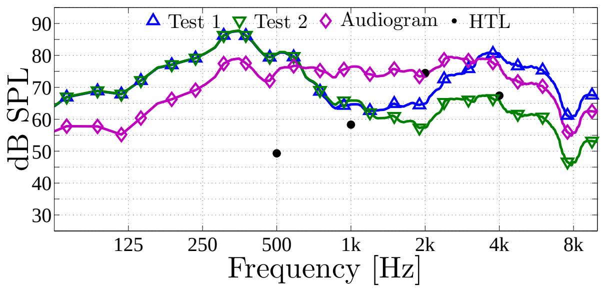

(b) Subject 4

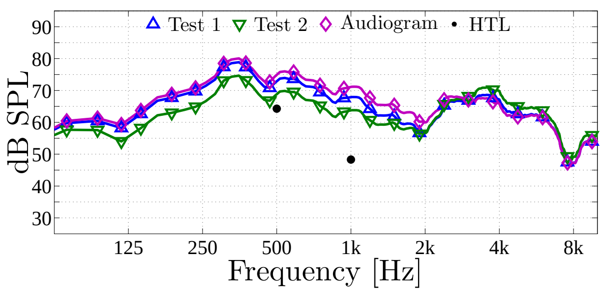

(c) Subject 7

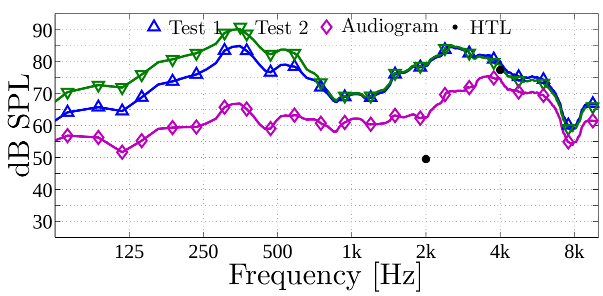

(d) Subject 12

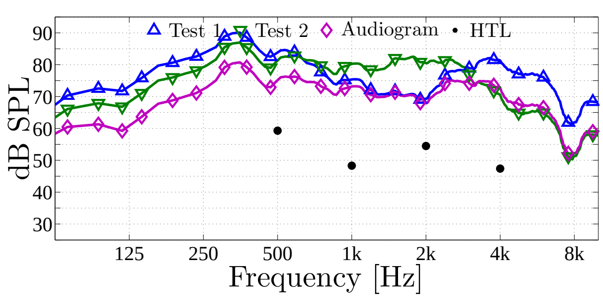

(e) Subject 13

Fig. 11. Power spectra of the SPL at the eardrum generated by the (left) HA with the obtained optimal setting from Test 1 (△) and Test 2 (▽) together with the audiogram (♢) for (a) subject 2, (b) subject 4, (c) subject 7, (d) subject 12, and (e) subject 13. The HTL of each subject at the four basis frequencies—500 Hz, 1 kHz, 2 kHz and 4 kHz—are indicated by black dots (if above 25 dB SPL). The A-weighted SPL at the location of the KEMAR/subject was measured to be 69.4 dB SPL. The peaks around 300 Hz are due to a Helmholtz resonance caused by a little leakage in the earplug of the KEMAR.

eters (see Fig. 10) and the output (see Fig. 11c). Remember, that this is the subject that could not decide what type of sound he preferred in the first study (Sec. III). Before the experiment, this subject actually remembered that he was not able to decide between two types of HA sound in the first experiment, and ensured that he would not behave similarly in the second experiment. This bias is a plausible explanation of why this subject suddenly was able to obtain a similar HA sound in the two tests.

The reproducibility of the found settings is not perfect. However, the processing in HAs for very different target-gains is not necessarily very different; it depends entirely on the fitting rationale. This is the case for the two settings obtained for subject 2 and 3. To realize this, compare the actual HA output in Fig. 11a with the difference in the obtained settings for subject 2 in Fig. 10a. The settings at 1 kHz suggests a gain difference of nearly 30 dB between Test 1 and Test 2. The difference, however, results in less than 5 dB (long-term) SPL difference at 1 kHz. Similarly, at 2 kHz the difference in the obtained setting is around 35 dB, but results in around 7-9 dB difference in the output at 2 kHz. Other effects also occur between the two settings due to the presence of dynamic compression in the HAs, but the long-term power spectra show that the HA outputs for the two target gains obtained for subject 2 were not too different after all. Nevertheless, at least the HA-output difference of 7-9 dB around 2 kHz must have been perceptually distinguishable. The reason why IHAPS apparently failed to obtain a similar parameter setting at 2 kHz in the two tests for subject 2, may be because a parameter change at 2 kHz given the other parameters results only in subtle changes of the HA output around the hearing threshold level (HTL) of subject 2 as seen in Fig. 11a. This demonstrate that internal dependencies among parameters in the HAs obviously need to be included in the analysis of the reproducibility. Some subjects apparently had an IRF with large regions with nearly identical responses as a results of their HTL in combination with the HA processing for different parameter settings. This was actually observed in the two dimensional case for several subjects including the example shown in Fig. 8. Furthermore, one might speculate that the HTL and the HA processing might have imposed multiple optima of the IRF with equal responses, such that the corresponding settings would have equally been preferred. With all the above in mind, it is fair to conclude that the reproducibility is acceptable overall, with subjects 4, 11 and 13 being the exceptions.

An interesting effect is observed from the bottom left figure of Fig. 10. The mean and standard deviation of the difference in the obtained optimal settings between test 1 and test 2 show a linear increasing tendency as a function of frequency. Three possible explanations (and likely a combination) are: First, after the first test the majority of subjects explained that they primarily preferred a full-bodied and soft sound as opposed to a thin and metallic sound. This indicates that at least in the beginning - i.e., a large part of test 1 - subjects tended to focus on the low-mid frequency regions, but might have been less aware of the subtle details at higher frequencies. Apparently, subjects were not aware of these details until later and possibly not until the second test. This suggests a training effect. Secondly, several subjects appeared to get tired and thus distracted during the second test, whereby they might not have noticed the subtle details at higher frequencies. From the results, the latter seems to be the case for subject 11 and 13. A third and indeed possible explanation is that the majority of the subjects had high-frequency sloping hearing losses. As a consequence, the SPL at higher-frequency was below the HTL for the majority of the settings. As a result, IHAPS might have learned that the high-frequency parameters had no influence on the IRF. IHAPS should eventually be able to learn that actually a limited range of these parameters (the settings above the HTL) imposes a perceptual differences. However, this requires that the active learning criterion queries settings within this limited range. This is not its main priority in the beginning with no assessments indicating that these parameters are important. The effect may explain the output difference around 2 kHz for subject 2 (see Fig. 11a). Obviously, this emphasizes the importance of restricting the parameter range of the HA devices to a reasonable range, where settings are perceptually different.

V. Conclusion and Future Directions

An interactive hearing-aid personalization system (IHAPS) based on a flexible non-parametric Gaussian process model and on an efficient sequential design is proposed. For ten HI subjects it was demonstrated that the system obtained a successful individual setting of a set of HAs controlled by either two or four parameters within ten to twenty user assessment—equivalent to a 5-10 min. session length. The subjects significantly preferred their individual setting provided that they could distinguish between the different settings. An obvious pitfall occurs if no perceptual difference exists for a large range of settings. Furthermore, listener fatigue and training effects appeared to noticeably influence the consistency of subjects and should be investigated more systematically.

In time, IHAPS may potentially be applicable in clinics to help both the hearing-care professional and the client to fine-tune hearing aids more efficiently and precisely to the client's preferences. To get there, the reproducibility of an individual setting should be further studied. Furthermore, the stimulus (music) was kept constant during the experiments; hence, the obtained settings may not generalize to other similar stimuli (music pieces). In a more realistic scenario, the stimulus used in each assessments could be randomly chosen from a library of music, speech and other sound types. Thereby, additional uncertainty is introduced, but an individual setting obtained with IHAPS in this manner has a better chance to generalize to for instance the music context in general.

Acknowledgment

This work was conducted in collaboration between the Technical University of Denmark and Widex A/S through the industrial PhD program (case number 10-096365) and supported in part by The Danish Agency for Science, Technology, and Innovation.

References

[1] National Acoustic Laboratories, "NAL-NL2," http://www.nal.gov.au/nal-software_tab_nal-nl-2.shtml, accessed: 28-07-2014.

[2] National Centre for Audiology, "DSL v5.0a," http://www.dslio.com/page/en/pubs_downloads.html, accessed: 28-07-2014.

[3] William A. Yost, *Fundamentals of Hearing - An Introduction*, Academic Press, 3rd edition, 1994.

[4] Brian C. J. Moore, *Cochlear Hearing Loss*, John Wiley & Sons, Ltd, 2nd edition, 2007.

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

[6] E. Convery, G. Keidser, H. Dillon, and L Hartley, "A self-fitting hearing aid: Need and concept," *Trends in Amplification*, vol. 15, no. (4), pp. 157–166, 2011.

[7] H. Dillon, J.A. Zakis, H McDermott, G. Keidser, W. Dreschler, and E. Convery, "The trainable hearing aid: What will it do for clients and clinicians?," *The Hearing Journal*, vol. 59, no. 4, 2006.

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

[9] 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.

[10] J. B. Nielsen, B. S. Jensen, J. Nielsen, and J. Larsen, "Hearing Aid Personalization," in *NIPS 2013 Workshop on Personalization*, Lake Tahoe, USA, 2013, pp. 2–5.

[11] J. B. Nielsen and Jakob Nielsen, "Efficient individualization of hearing aid processed sound," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2013.

[12] 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.

[13] 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.

[14] 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.

[15] 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.

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

[17] R. D. Bock and J. V. Jones, *The Measurement and Prediction of Judgment and Choice*, Holden-day, 1968.

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

[19] 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.

[20] B. S. Jensen and J. B. Nielsen, "Pairwise Judgements and Absolute Ratings with Gaussian Process Priors," Tech. Rep., DTU, IMM, November 2011. http://www2.compute.dtu.dk/pubdb/p.php?6151.

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

[22] 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.

[23] A. Gelman, "Prior distributions for variance parameters in hierarchical models," *Bayesian Analysis*, vol. 1, no. 3, pp. 515–533, 2006.

[24] 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.

[25] J. Vanhatalo and A. Vehtari, "Speeding up the binary Gaussian process classification," in *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence*, 2010.

[26] J. B. Nielsen, "Supplementary material - perception based personalization of hearing aid using gaussian processes and active learning," Jan 2014. http://www2.compute.dtu.dk/pubdb/p.php?6726.

[27] Hans Bruun Nielsen, "A matlab toolbox for optimization and data fitting," 2010. http://www2.compute.dtu.dk/~hbni/immoptibox/Manual.pdf.

Jens Brehm Bagger Nielsen received the M.Sc. degree in electrical engineering from the Technical University of Denmark (DTU) in 2010 with a background in auditory signal processing, acoustics and non-linear modeling. In 2014 he was awarded the Ph.D. degree from the Department of Applied Mathematics and Computer Science at DTU following an industrial Ph.D. program carried out between 2011 and 2014 in collaboration with Widex A/S. Over the past three years, Dr. Nielsen has published several papers in the audio signal processing and machine learning community. Currently, he is employed at Widex as R&D engineer, working with audiological signal processing, machine learning, psychoacoustics and audiology.

Jakob Nielsen received the M.Sc. degree in Multi Media Technology from the IT-University of Copenhagen in 2005. He also holds a B.Sc. degree in Mathematics and BA in Musicology from the University of Copenhagen. Since 2005, Jakob Nielsen has been employed at Widex A/S working with research and innovation in the intersection between signal processing, psychoacoustics and audiology.

Jan Larsen received the M.Sc. and Ph.D. degrees in electrical engineering from the Technical University of Denmark (DTU) in 1989 and 1994, respectively. Dr. Larsen is Associate Professor of Digital Signal Processing at Department of Applied Mathematics and Computer Science, DTU. Jan Larsen has authored and co-authored around 150 papers and book chapters within the areas of nonlinear statistical signal processing, machine learning, neural networks and datamining with applications to biomedicine, monitoring systems, multimedia, audio, and web-mining. According to Google Scholar the no. of citations is 3063, h-index is 30 and i-10 index 76 as of 27.11.2014. He has participated in more than ten national and international research and innovations programs and has served as reviewer for many international journals, conferences, publishing companies and research funding organizations. Currently he coordinates the strategic the research project CoSound (2012-2016) and participate in five other research projects. As regards synergistic activities, he took part in conference organizations, among others the IEEE Workshop on Machine Learning for Signal Processing (formerly Neural Networks for Signal Processing) 1999-2012. He is director of the Danish Sound Innovation Network (2009-), past chair of the IEEE Machine Learning for Signal Processing Technical Committee of the IEEE Signal Processing Society (2005-2007), and chair of IEEE Denmark Section's Signal Processing Chapter (2002-). He is a senior member of The Institute of Electrical and Electronics Engineers. Other professional committee participation includes Steering committee member of the Audio Signal Processing Network in Denmark, 2006-. Editorial Board Member of Signal Processing, Elsevier, 2006-2007; and guest editorships involves IEEE Transactions on Neural Networks; Journal of VLSI Signal Processing Systems; and Neurocomputing.