Towards Hearing-Aid Personalization: Preference Elicitation from Audiological Data

Adriana Birlutiu[1], Perry Groot[1], Tom Heskes[1], Iman Mosavat[2], Bert de Vries[2,3], and Tjeerd Dijkstra[1,3]

[1] Radboud University Nijmegen, [2] TU Eindhoven, [3] GN ReSound Eindhoven

1. The HearClip Framework

Given a sound library, a sound sample and two parameter settings are selected to generate two hearing-aid output signals. The patient gives a preference response, whether he prefers the first or the second signal. Patient preferences are modeled by a utility function. Our uncertainty about the "true" utility model of a patient is represented by a probability distribution which is updated using Bayes' rule after each listening experiment. Using the new probability distribution, this process is repeated.

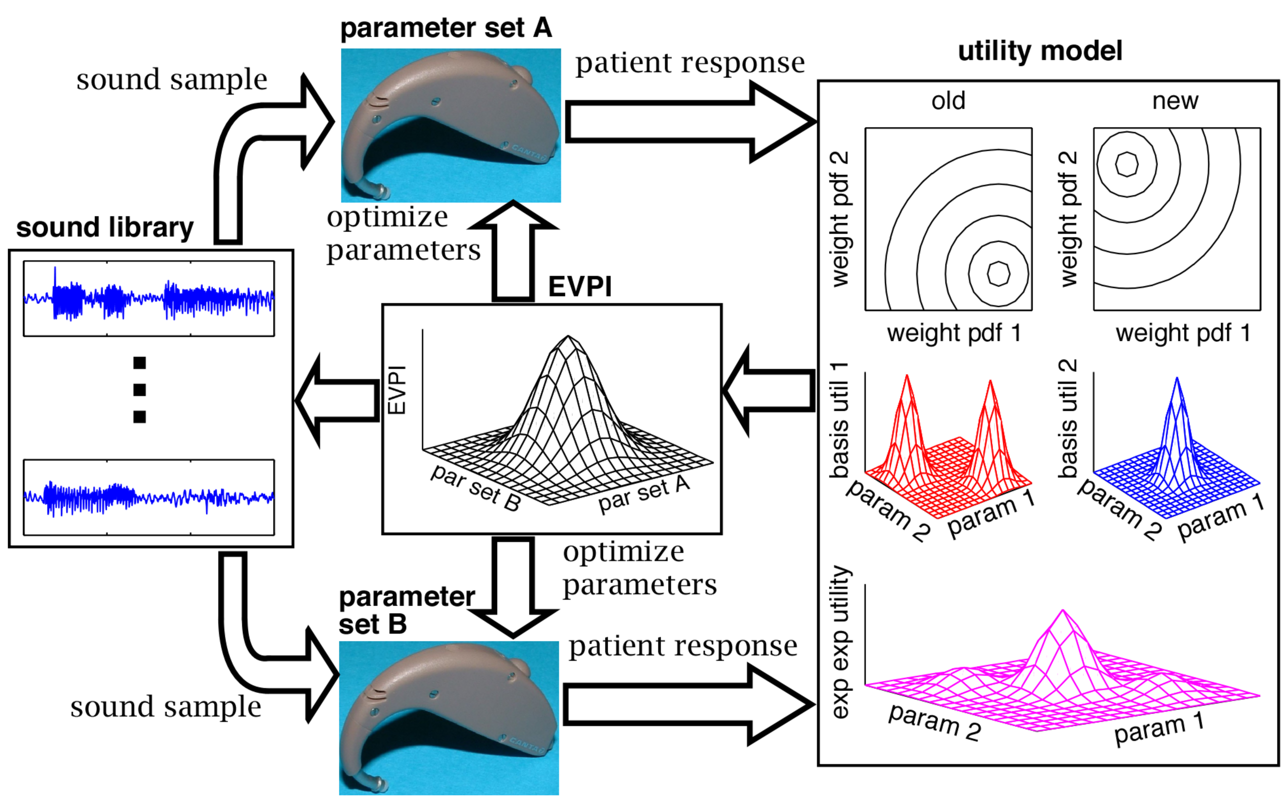

Figure 1: *Fitting a Hearing Aid [1]*

2. Pairwise Comparison Preference Data

In [2], the prediction of sound quality of 14 normal hearing and 18 hearing impaired patients was investigated. Each patient was subjected to paired comparison tests for each possible pair of 24 conditions typically found in hearing aids (noise and clipping operations). Hence, for each patient there are $24^2 = 576$ paired comparisons of the form (x_1, x_2, d) where $d \in \{1, 2\}$ denoting whether sound sample x_1 or x_2 was preferred by the patient, respectively. This data set has been made available for analysis in the HearClip project.

3. Transfer Learning

Question: Can we use the preferences from one subject to better learn the preferences of another subject?

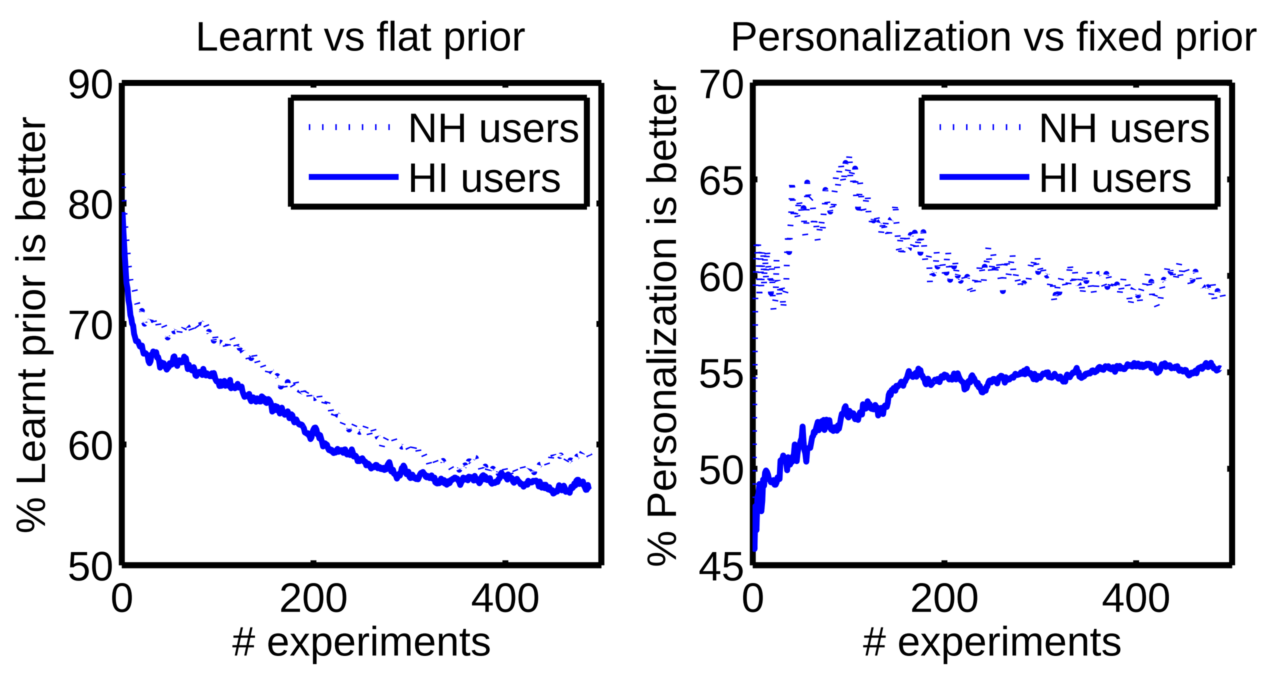

Figure 2: *Percentage of the number of times the prediction accuracy using a learnt prior is better than the prediction accuracy with a flat prior (left); and a fixed prior that does not learn from the experiments (right). With few listening experiments, using a learnt prior leads to a higher accuracy; and personalization improves the accuracy, in particular for hearing impaired patients.*

4. Optimal Experiment Selection

Question: Can we learn faster by presenting the right experiments?

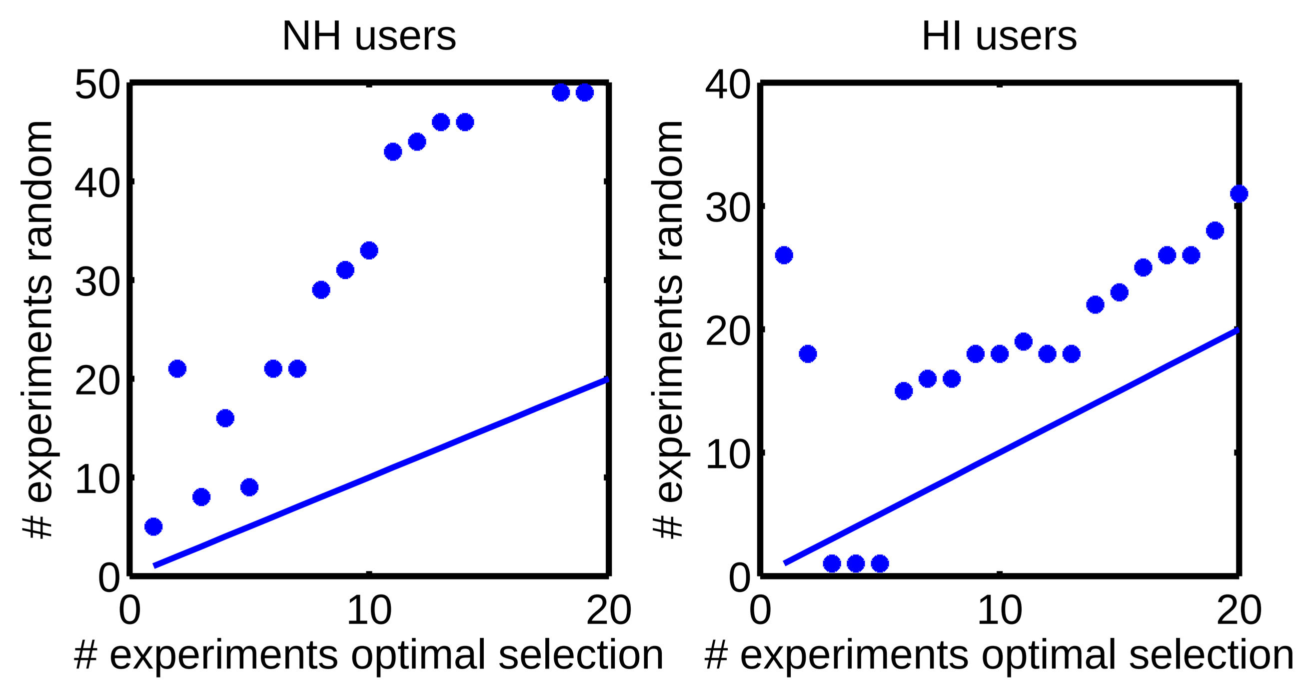

Figure 3: *The number of listening experiments needed using random selection to get the same accuracy as optimal selection. The optimal experiment selection is implemented by actively choosing the listening experiment that gives the most information about patient's preferences. Left: normal hearing patients. Right: hearing impaired patients.*

5. Linear versus Non-Linear Utility Functions

Question: Can we stick to a simple linear utility model or do we have to consider nonlinear dependencies?

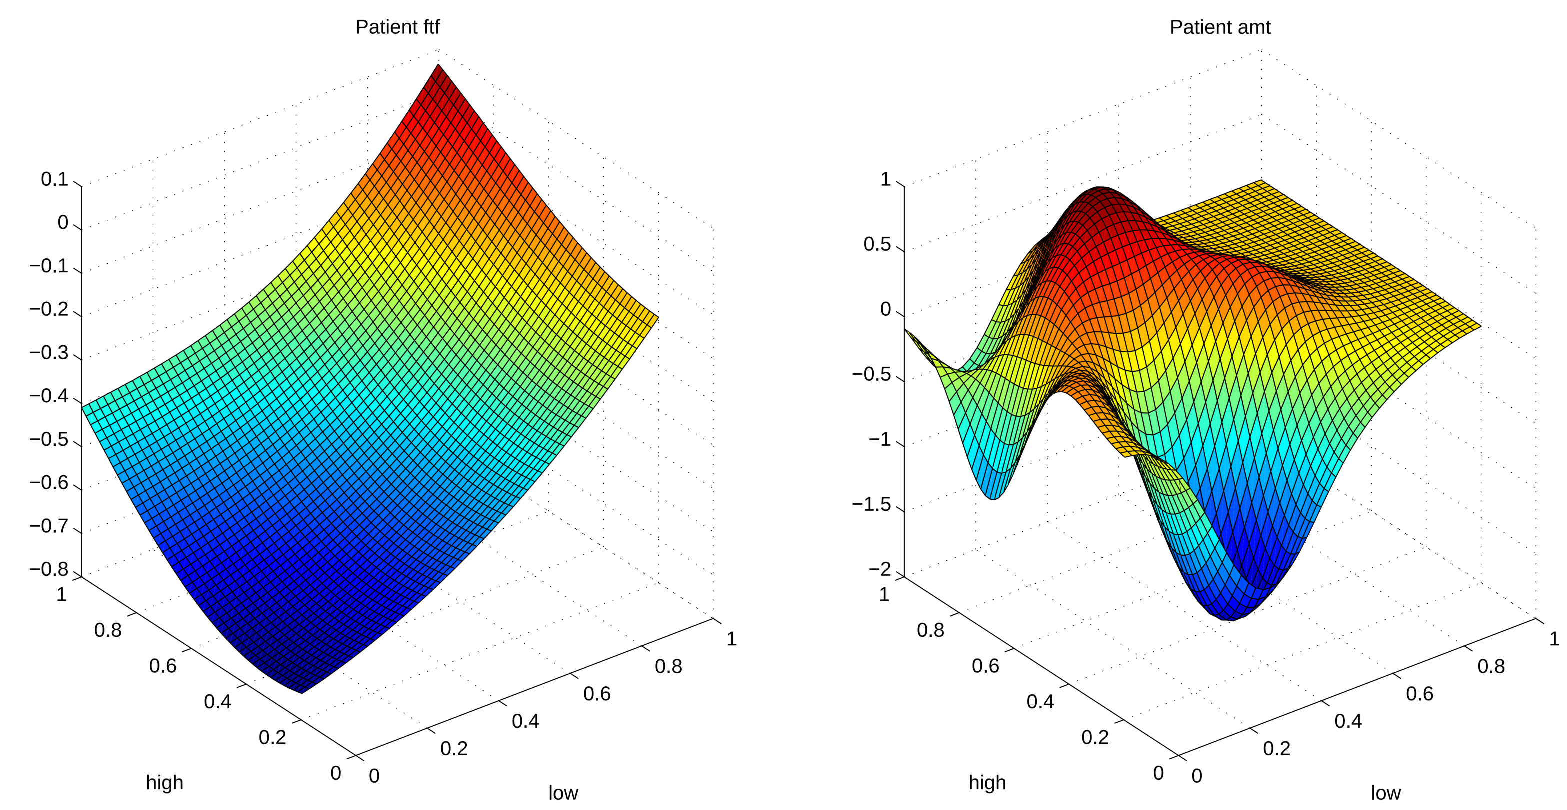

Figure 4: *Two utility functions. Left: normal hearing patient. Right: hearing impaired patient. We have investigated earlier linearity assumptions of patient utility functions using the pairwise comparison audiology data of [2]. We have used a non-parametric modeling approach with Gaussian processes and have found significant improvements with respect to the linear model for hearing impaired patients (based on a 2-sided McNemar test).*

References

[1] Tom Heskes and Bert de Vries, Incremental Utility Elicitation for Adaptive Personalization, *BNAIC*, Brussels, October 2005

[2] K.H. Arehart, J.M. Kates, M.C. Anderson, and L.O. Harvey jr. Effects of noise and distortion on speech quality judgements in normal-hearing and hearing-impaired listeners. *J Acoust Soc Am*, 122, 1150-1164.

Netherlands Organisation for Scientific Research

STW project 07605 and NWO VICI 639.023.604.