Bayesian Incremental Utility Elicitation with Application to Hearing Aids Personalization

Adriana Birlutiu[1], Tom Heskes[1], Bert de Vries[2], Alexander Ypma[2] and Tjeerd Dijkstra[2]

[1] Radboud University Nijmegen, [2] GN ReSound Eindhoven, The Netherlands

Abstract

Medical devices such as hearing aids contain many tunable parameters. The optimal setting of these parameters depends on the patient's preference (utility) function, which is unknown. This raises two questions: **(1) how should we optimize the parameters given partial information about the patient's utility?** And **(2), what questions do we ask to efficiently elicit this utility information?** We use a probabilistic decision-theoretic framework to answer these questions.

1. The Fitting Problem

- **(Algorithm)**. Consider a Hearing Aid (HA) algorithm $y = F(x, \theta)$, where x and y are input and output signals respectively and $\theta \in \Theta$ a set of tuning parameters. We assume that x is selected from an environment (i.e. a set of acoustic signals) $X = \{x_1, \ldots, x_K\}$ with probabilities $P(x_k)$.
- **(Utility model)**. Each patient has different preferences. In practice, we model "patient satisfaction" by a utility function $U(y; w_i)$, where w_i is the utility state of patient i. Our uncertainty about the "true" utility model is represented by a probability distribution $P(w|D, \alpha_i)$, where $\alpha_i \in A$ holds background information about patient i, such as his auditory profile, age, (etc.) and D refers for experimental data obtained from listening tests.
- **(Fitting)**. The fitting goal is to find the set of parameters θ_i^* that is expected to be "optimal" for patient i relative to environment X. Thus, we wish to maximize θ w.r.t. the so-called expected expected utility (EEU),

$$\theta_i^* = \arg\max_{\theta} \sum_k P(x_k) \int_W P(w|D, \alpha_i)\, U(x_k, \theta; w)\; \mathrm{d}w \tag{1}$$

Eq. 1 answers our first question (see abstract).

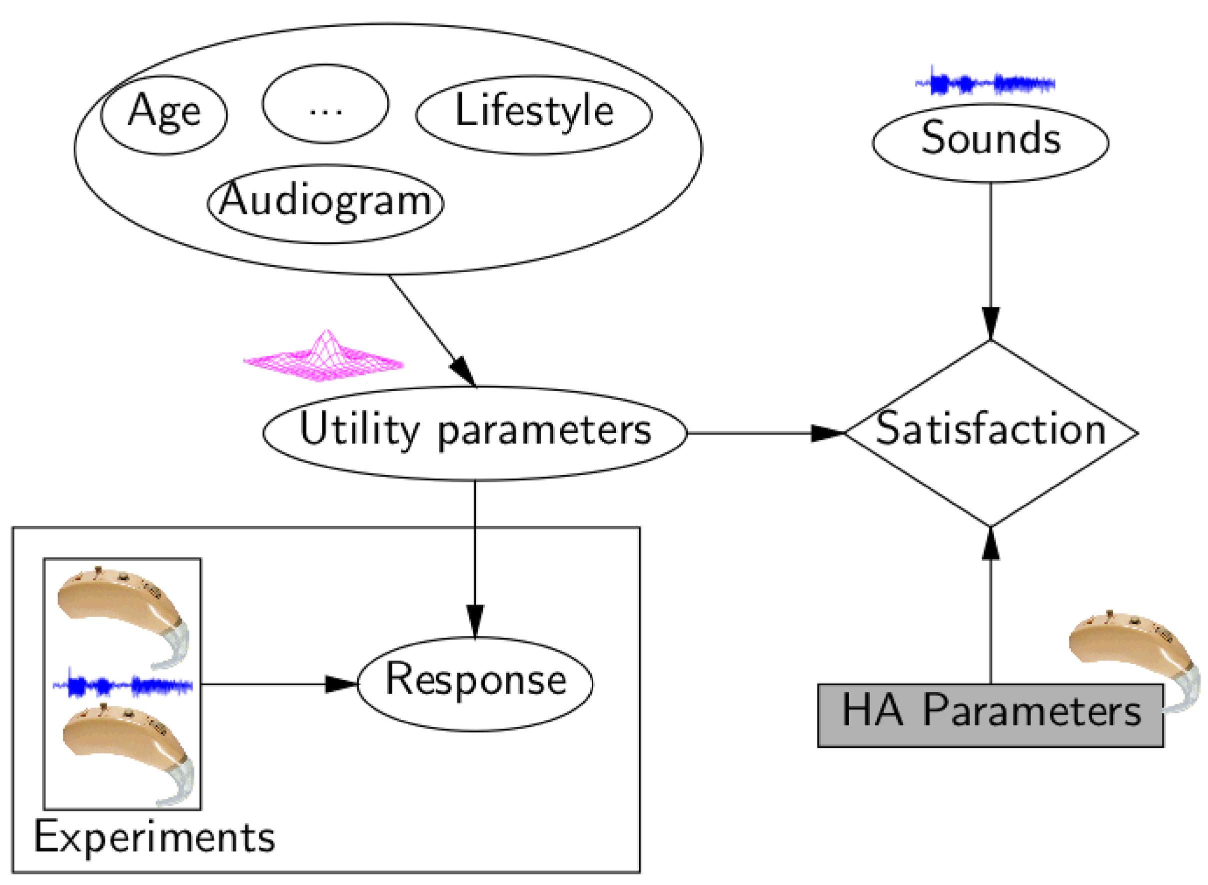

Figure 1: *Fitting a Hearing Aid*

2. Incremental Utility Elicitation

Eq. 1 leads to a satisfying fit if the utility model accurately reflects the patient's preferences. We can incrementally update our knowledge about the utility model through listening tests.

- **(Pairwise comparison)**. Let's use pairwise comparative listening experiments. The nth listening experiment e^n consists of presenting an input x from X in combination with two parameter settings θ_1 and θ_2 from Θ, i.e. $e^n = \{x^n, \theta_1^n, \theta_2^n\}$. The patient's preference decision d^n follows a Bradley-Terry (logistic regression) model:

$$p(d|e, w) = \frac{1}{1 + \exp\{-d \times [U(x, \theta_1; w) - U(x, \theta_2; w)]\}} \tag{2}$$

- **(Bayesian updating)**. Suppose $P(w|D^n, \alpha)$ denotes the PDF over utility states w after having seen the results of n experiments. Then we can use Bayes rule to incrementally absorb information from the $(n+1)$th experiment through

$$P(w|D^{n+1}, \alpha) \propto P(d^{n+1}|e^{n+1}, w)\, P(w|D^n, \alpha)$$

with $P(d^{n+1}|e^{n+1}, w)$ from Eq. 2.

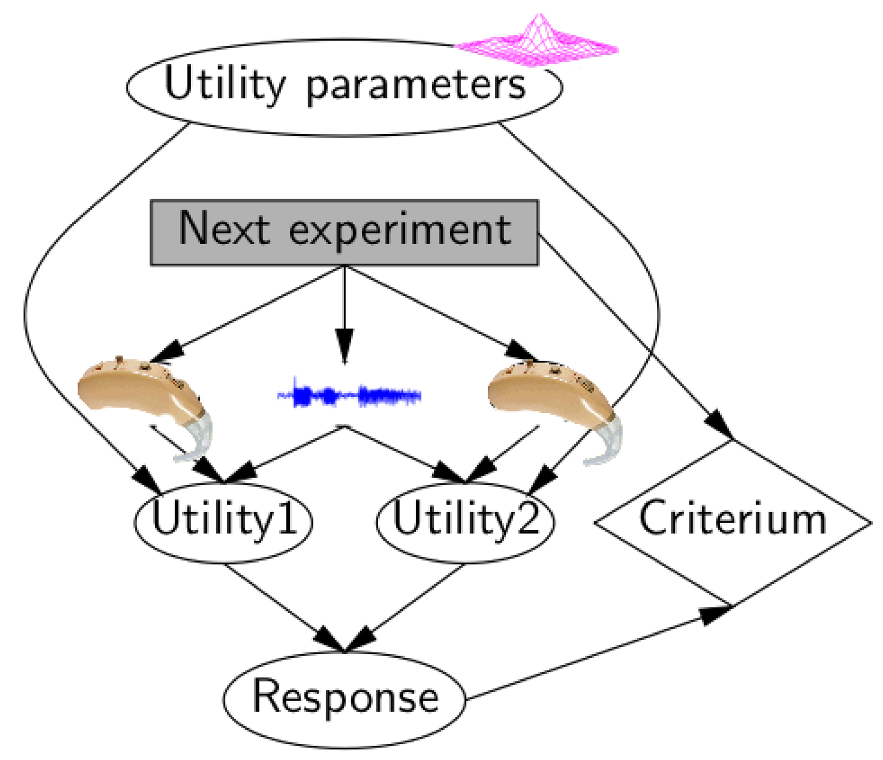

Figure 2: *Incremental Utility Elicitation*

- **(Optimal experiment selection)**. Given $P(w|D^n, \alpha)$, how do we select the "best" next listening experiment? We use Bayesian optimal experimental design to select the experiment that maximizes the Expected Value of Perfect Information (EVPI) by

$$e^* = \arg\max_e \sum_{d \in \{-1,1\}} P(d|e, D^n, \alpha) \max_{\theta} \int_w P(w|d, e, D^n, \alpha)\, \mathrm{EU}(\theta, w)\; \mathrm{d}w$$

3. Example

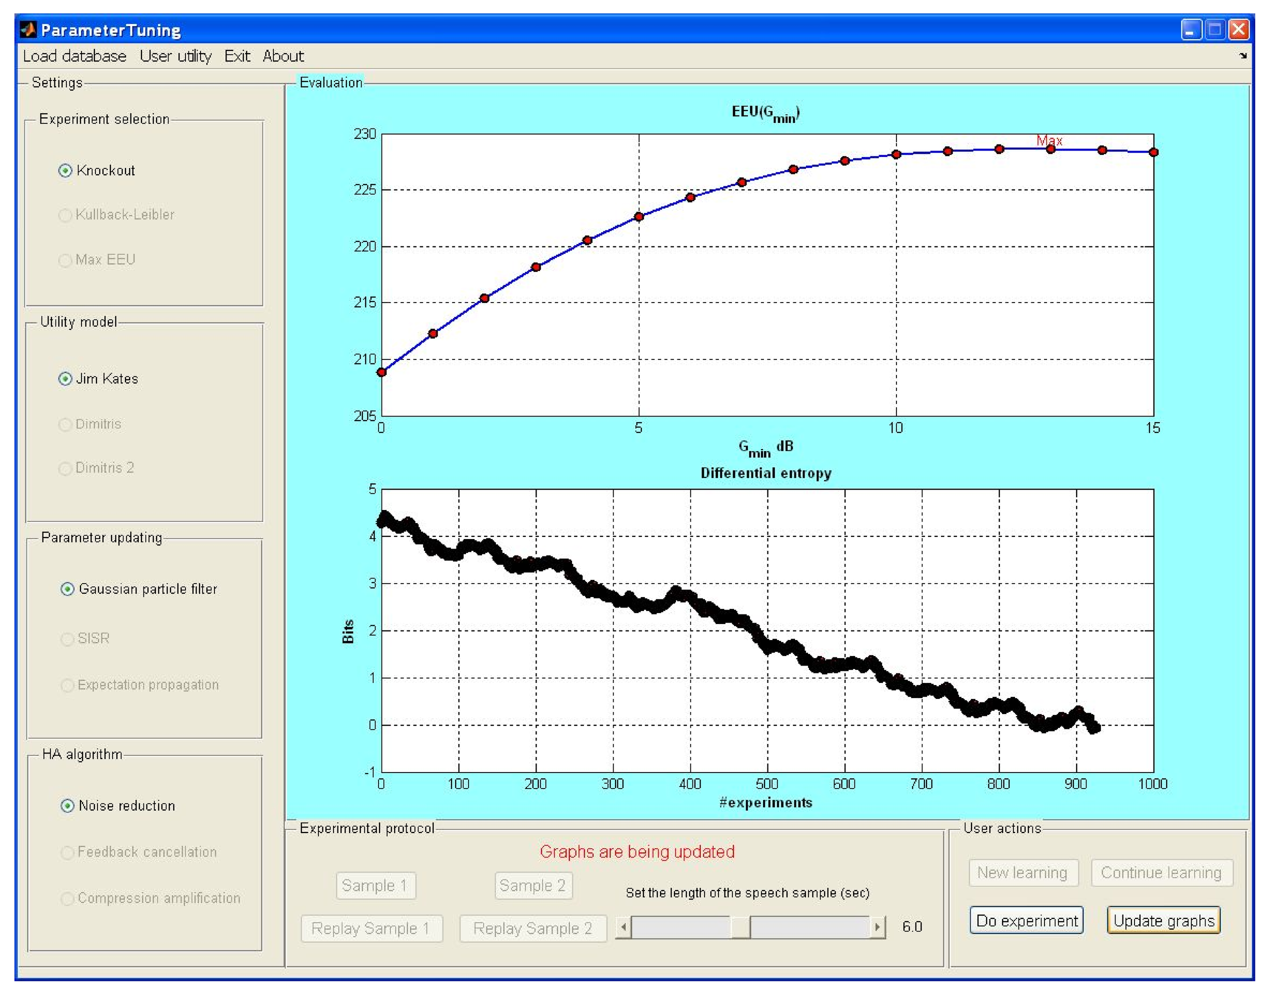

Figure 3: *Snapshot Incremental Fitting GUI during a noise suppression fitting session. The top graph shows expected expected utility (EEU) vs hearing aid parameter θ. The bottom graph shows the uncertainty in the model (entropy of $P(w|D^n, \alpha_i)$) vs listening experiment index.*

4. Challenges

- Deal with large computational complexity
- Design of appropriate utility models
- User ($=$ dispenser) interface

References

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

[2] U. Chajewska, D. Koller, and R. Parr. Making rational decisions using adaptive utility elicitation. In AAAI/IAAI, 2000