Personalization of Hearing Aids through Bayesian Preference Elicitation

Tom Heskes   Bert de Vries
Adriana Bîrluțiu   Anton Vakhrushev

Radboud University Nijmegen, TU Eindhoven
The Netherlands

December 8th, 2006

Outline

- The problem: tuning hearing aids
- The utility model
- Experimental setup
- Bayesian updating
- Optimal decision
- Optimal incremental utility elicitation
- Challenges
- Demo

The problem

1 out of 4 buyers is dissatisfied with the sound quality of the hearing aid.

- High-dimensional parameter space.
- Expected patient satisfaction is
  - unknown
  - unique for each patient
- Costly and unreliable listening tests.

How do we optimally tune the parameters such as to adapt them to the patient's (unknown...) preferences?

The utility model

Bayesian decision network:

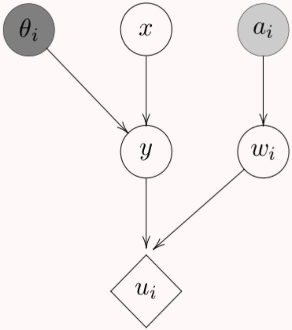

θ_i : tuning parameters
x : input signal
a_i : auditory profile
y : output signal
w_i : utility state
u_i : utility function

1. Utilities can be represented in a functional form $u_i = U(y; w_i) = U(x, \theta; w_i)$
2. The utility state w_i is treated as probabilistic variable
3. The auditory profile a_i specifies the prior over w_i

Experimental setup

Paired-comparison forced-choice:

- Experiment e consists of a sound x with two parameter settings θ_1 and θ_2, $e = \{x, \theta_1, \theta_2\}$.
- Patient is asked whether he prefers: $y_1 = F(x, \theta_1)$ over $y_2 = F(x, \theta_2)$, $y_1 \succ y_2$.

Modeling assumption is the Bradley-Terry model

$$P(y_1 \succ y_2) = \frac{1}{1 + \exp(-[U(y_1; w) - U(y_2; w)])}$$

Incremental Bayesian updating of the utility model

Let $P(w|D_n)$ be the probability density over utility state w after n experiments.

From *Bayes' rule* $\Rightarrow$ the posterior over w after the $(n+1)$th experiment

$$P(w|D_{n+1}) = P(w|d^{n+1}, e^{n+1}, D_n) = \frac{P(d^{n+1}|e^{n+1}, w)P(w|D_n)}{P(d^{n+1}|e^{n+1}, D_n)}$$

- Prior is the posterior from the previous step.
- Likelihood is the Bradley-Terry model.

Bayesian updating of the utility model in a paired-comparison experimental design

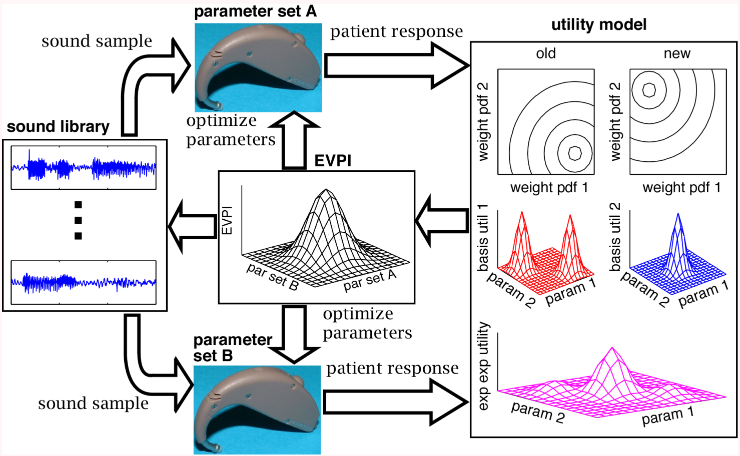

Make the optimal decision given partial utility information

- For w known, the optimal setting of parameters is the one that maximizes the *expected utility*: the utility averaged over a library of sounds

$$EU(\theta, w) = \sum_{x \in X} P(x) U(x, \theta; w)$$

$$\theta^*(w) = \arg\max_{\theta \in \Theta} EU(\theta, w)$$

- For w distributed according to $P(w|D_n)$ we average over this probability distribution as well $\Rightarrow$ *expected expected utility* (Boutilier, 2003)

$$EEU(\theta) = \sum_{x \in X} P(x) \int_W dw P(w|D_n) U(x, \theta; w)$$

$$\theta^* = \arg\max_{\theta \in \Theta} EEU(\theta)$$

Efficient utility elicitation

Bayesian optimal experimental design (Lindley, 1972).

Choose the next experiment which maximizes the expected gain in utility.

$$EV|PI(e) = \sum_{d=\pm 1} P(d|D_n, e) \max_{\theta} \int_W dw P(w|D_n, d, e) EU(\theta, w)$$

Future work

- Model choices
  - Modeling the utility function.
  - Criteria for the terminal decision.
  - Other criteria for experimental design.
  - Prior for the utility model.
- Algorithmic issues
  - Algorithms for efficient optimization.
  - Library of sounds and parameters.
- Audiological evaluation