Optimal Experimental Design in a Hierarchical Setting for Probabilistic Choice Models

Adriana Birlutiu and Tom Heskes

Institute for Computing and Information Sciences
Radboud University Nijmegen
The Netherlands

Outline

Outline

Preliminaries

Active experiment selection

Experiments

Conclusions and future work

Learning user's preferences

- Personalization of hearing aids: learning user's preferences over sound quality
- Optimal experimental design, *D-optimal criterion*
- Active learning, *Query-by-Committee*
- Criterion that makes use of the judgement of other users

Probabilistic choice models

- Probability that a subject with parameters $\boldsymbol{\theta}$ prefers option k when given input $\mathbf{x}$

$$p(k; \mathbf{x}, \boldsymbol{\theta}) = \frac{\exp\left[\sum_{j=1}^{n} A_{kj}(\mathbf{x})\theta_j\right]}{Z(\boldsymbol{\theta}, \mathbf{x})} \,,$$

  $Z(\boldsymbol{\theta}, \mathbf{x}) =$ normalization constant, $A =$ function which extracts features of the input $\mathbf{x}$ related to option k

- $\boldsymbol{\theta}$ treated as a random variable and updated using Bayes' rule

Hierarchical modeling (1)

- Use Expectation-Maximization algorithm to gather data from the other users in a prior distribution for the new user
- $P(\boldsymbol{\theta}_i) = G(\boldsymbol{\theta}_i; \boldsymbol{\mu}, \Sigma)$ Gaussian prior, the same $\boldsymbol{\mu}$ and Σ for all subjects

$$\boldsymbol{\mu} = \frac{1}{M} \sum_{i=1}^{M} \boldsymbol{\theta}_i^*$$

$$\Sigma = \frac{1}{M} \sum_{i=1}^{M} (\boldsymbol{\theta}_i^* - \boldsymbol{\mu})(\boldsymbol{\theta}_i^* - \boldsymbol{\mu})^T + \frac{1}{M} \sum_{i=1}^{M} V_i$$

$\boldsymbol{\theta}_i^*$ and V_i are the posterior mean and variance for subject i computed based on the previous prior mean and variance

Hierarchical modeling (2)

- Make use of the data available from a group of users for which preferences were already learnt, when learning preferences of a new user

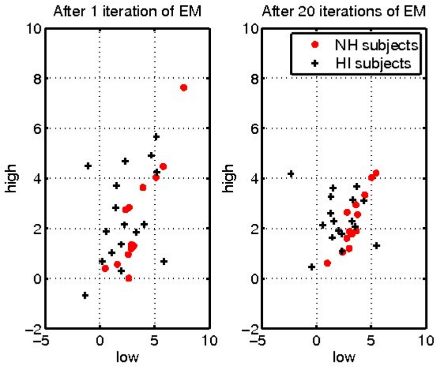

Outline

Criteria for selecting experiments

Make use of the data available from other subjects when selecting the experiments to perform with a new subject

- Select those experiments for which the other subjects disagree the most, according to their responses
  - Experimentally $\Rightarrow$ not good enough

- Select those experiments with the highest disagreement between the other subjects, induced by the uncertainty in their distribution (committee criterion)

Committee criterion

$$I_{\text{committee}}(\mathbf{x}) = \frac{1}{M-1} \sum_{j \neq i} \sum_{k} \bar{p}_{\setminus i}(k;\mathbf{x}) \log \left[\frac{\bar{p}_{\setminus i}(k;\mathbf{x})}{p_j(k;\mathbf{x})} \right] - \sum_{k} \bar{p}_{\setminus i}(k;\mathbf{x}) \log \left[\frac{\bar{p}_{\setminus i}(k;\mathbf{x})}{p_i(k;\mathbf{x})} \right]$$

- $p_j(k;\mathbf{x}) = p(k;\mathbf{x},\boldsymbol{\theta}_j^*)$
  $\boldsymbol{\theta}_j^* \equiv$ maximum posterior solution for subject j
- $\bar{p}_{\setminus i}(k;\mathbf{x}) =$ logarithmic average over all subjects $j \neq i$

$$\bar{p}(k;\mathbf{x}) = \frac{1}{\Gamma(\mathbf{x})} \exp \left[\int d\boldsymbol{\theta}\, P(\boldsymbol{\theta}) \log p(k;\mathbf{x},\boldsymbol{\theta}) \right]$$

$\Gamma(\mathbf{x}) =$ normalization constant

Committee criterion - example

committee criterion vs. disagreement

	exp #10		exp #28		exp #35	
	response	prediction	response	prediction	response	prediction
subjects						
	0	0.1	0	0.4	1	0.95
	0	0.05	0	0.45	1	0.87
	0	0.2	1	0.6	0	0.01
	0	0.1	0	0.4	0	0.1
	0	0.02	1	0.55	0	0.05
	0	0.08	1	0.6	1	0.9
	1	0.87	1	0.65	0	0.18
	0	0.07	1	0.6	0	0.1
	1	0.95	0	0.35	1	0.9
	0	0.1	0	0.4	1	0.97
	disagree *low*	opinion *low*	disagree *high*	opinion *low*	disagree *high*	opinion *high*

Other design criteria

- D-optimal criterion $\equiv$ the accuracy with which the parameters of the model can be estimated
- in the Bayesian context $\Leftrightarrow$ reduction in the entropy of the posterior distribution over model parameters

$$I_{\text{det}}(\mathbf{x}) = -\sum_k p(k;\mathbf{x}) \log \det V(k,\mathbf{x}) + \log \det V$$

  - $V(k,\mathbf{x}) =$ new variance after presenting $\mathbf{x}$ and observing k
  - $p(k;\mathbf{x}) =$ probability that the subject prefers alternative k when presented $\mathbf{x}$

Connections between design criteria

- Assuming that $V(k, \mathbf{x})$ is close to V

$$I_{\text{det}}(\mathbf{x}) \approx \sum_k p(k; \mathbf{x}, \boldsymbol{\theta}^*) \mathbf{g}(k; \mathbf{x}, \boldsymbol{\theta}^*)^T V \mathbf{g}(k; \mathbf{x}, \boldsymbol{\theta}^*)$$

$$\boldsymbol{\theta}^* = \text{maximum posterior solution}$$

- Connection with the standard D-optimal criterion

$$I_{\text{committee}}(\mathbf{x}) = \frac{1}{2} \sum_k p(k; \mathbf{x}, \boldsymbol{\mu}) \mathbf{g}(k; \mathbf{x}, \boldsymbol{\mu})^T \tilde{V} \mathbf{g}(k; \mathbf{x}, \boldsymbol{\mu})$$

$\boldsymbol{\mu}$ = prior mean learned from all other subjects and

$$\tilde{V} \equiv \frac{1}{M-1} \sum_{j \neq i} (\boldsymbol{\theta}_j^* - \boldsymbol{\mu})(\boldsymbol{\theta}_j^* - \boldsymbol{\mu})^T - (\boldsymbol{\theta}_i^* - \boldsymbol{\mu})(\boldsymbol{\theta}_i^* - \boldsymbol{\mu})^T$$

Outline

Experiments

We addressed the following two questions:

1. Can we use the already learned preferences of other subjects to better learn the preferences of the current subject?

2. Can we learn faster by optimally selecting the experiments to present to a subject?

Experiments (1)

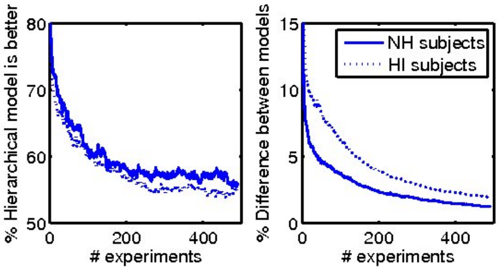

- Left: percentage of the number of times the prediction accuracy using the hierarchical prior is better than the prediction accuracy with a flat prior.
- Right: percentage of the number of predictions on which the two models (with the hierarchical and with a flat prior) disagree.

Experiments (2)

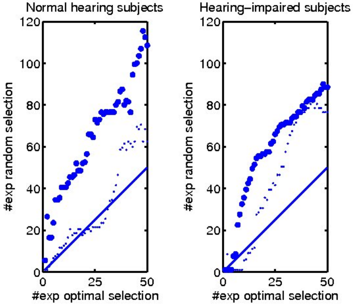

- The number of listening experiments needed using random selection (on the *y*-axis) to get the same accuracy as with the optimal selection (on the *x*-axis).
- The optimal experiment selection is implemented by presenting those experiments which are hard to predict according to other subjects responses (the small dots) and actively using the opinion pool criterion (the large dots).

Related work

Committee criterion, proposed here, is similar to the ones used in

- A. McCallum and K. Nigam *Employing EM and Pool-Based Active Learning for Text Classification (ICML'98)*
- P. Melville et al. *Active learning for probability estimation using Jensen-Shannon divergence (ECML'05)*

with the difference that

- different setting, preference learning
- real subjects as members of the committee

Outline

Conclusions

- We proposed a new criterion for experimental design that makes direct use of the judgements of other subjects
- Advantage of this new criterion:
  - interpretation and simplicity
  - easy to compute
- Future work: apply active experiment selection to other types of models, in particular to Gaussian Processes

Conclusions

- We proposed a new criterion for experimental design that makes direct use of the judgements of other subjects
- Advantage of this new criterion:
  - interpretation and simplicity
  - easy to compute
- Future work: apply active experiment selection to other types of models, in particular to Gaussian Processes

Thanks for your attention!