HELSINKI INSTITUTE FOR INFORMATION TECHNOLOGY

Aalto University

UNIVERSITY OF HELSINKI

Improving Controllability and Predictability of an Interactive User Model Driven Search Interface

Antti Kangasrääsiö
Dorota Głowacka
Samuel Kaski
firstname.lastname@hiit.fi

Motivation

In exploratory search, the user searches for information in a domain she is not initially familiar with. Because of this, the feedback is often uncertain.

This means that search interfaces are faced with a difficult problem: how to help the user direct the search using uncertain feedback.

Exploration / exploitation problem:
- If the feedback is certain, it can be interpreted in exploitative manned
- If the feedback is uncertain, we likely need to add in some exploration

Reinforcement learning based probabilistic user models can be used to handle the exploration / exploitation trade-off. However, they may also introduce usability problems.

These models generally assume that the user feedback are “samples from a function to be approximated”. However, the user is not a passive function, but instead trying to actively steer the system.

We propose that there needs to be a layer of interpretation between the user and the underlying model. This layer is responsible for:
- Translating user feedback into requirements for the state of the system
  → Improve the controllability of the system
- Allowing the user to predict the effects her actions will have on the system
  → Improve the predictability of the system

System Overview

This work is based on the SciNet search interface (cf. Glowacka *et al.* Directing exploratory search: Reinforcement learning from user interactions with keywords. IUI'13). In this system the user interactively gives feedback on the search intent model by moving keywords around on a radar display.

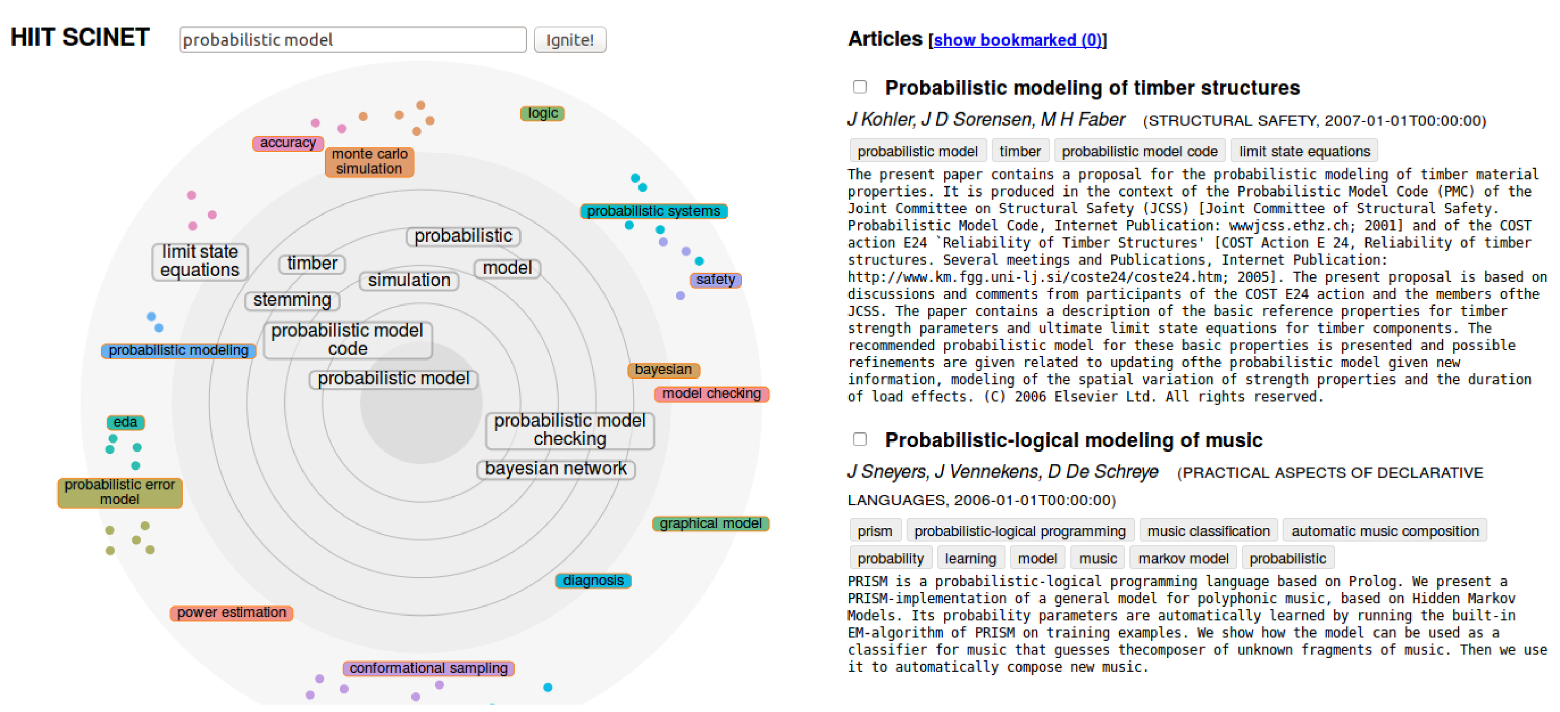

Interpreting User Feedback as Goals Instead of Just Data

Problem
Many user models treat relevance feedback data as just data points for fitting a model. However, this may lead the system to behave in a way not intended or anticipated by the user. For example, past feedback may weigh more than the new data, making it difficult for the user to cause the effects she intends to happen just by giving feedback.

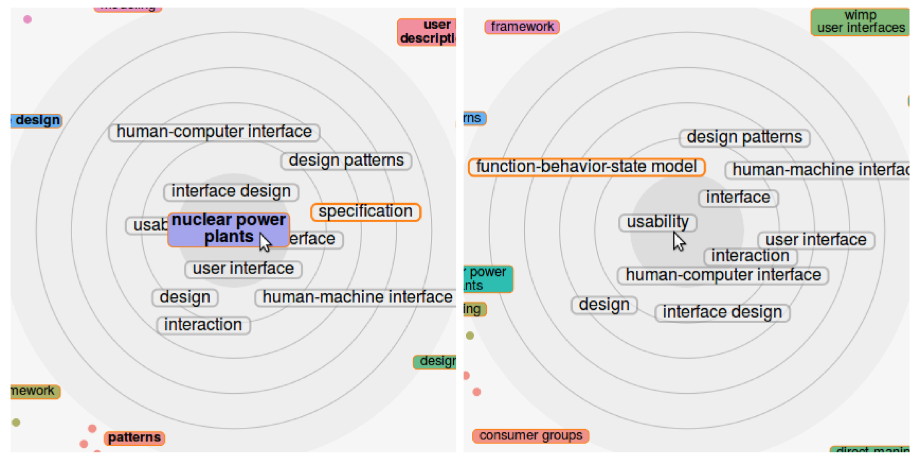

Left: giving maximal relevance feedback to the blue keyword
Right: resulting user model does not include the keyword?

Solution
The feedback given by the user is interpreted as a goal for an optimization problem regarding the next state of the system.

For example, if the user indicates that a certain keyword has relevance X to her search intent, then the optimal value of that keyword in the resulting model is X. In order to find the “optimal feedback” to make this happen from the model's point of view, the user has an automatic assistant that calculates this for her.

Enabling Predictability of Feedback Actions

Problem
Many user models are complex, and thus it may be difficult for the user to predict the actions her feedback will have on the system.

Even though the parts of the model the user gives feedback on are controlled, as mentioned above, the feedback may still cause other effects the user does not intend or can not anticipate.

Solution
If making the model simpler is not possible without sacrificing performance, one way to solve this problem is to show the user a visualization of the effects of her feedback while she is deciding on what feedback to give to the system. This way the user is able to choose the feedback based on the expected effects to the system.

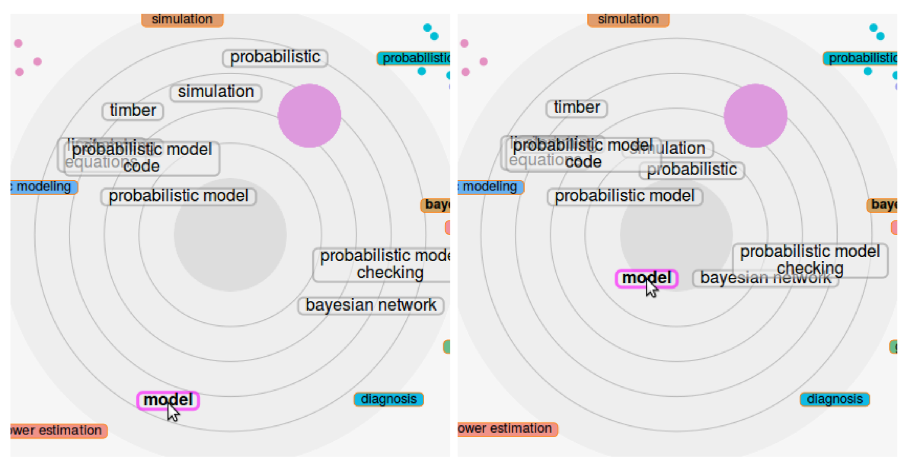

When the user is choosing what relevance feedback to give, the locations of other keywords move accordingly.

Experimental Results

A user study was conducted on 12 users, of which 2 had to be excluded as outliers. Each user performed two exploratory search tasks: one using the search engine without the improvements (baseline system) and one using the search engine with the improvements (improved system). One of the search tasks had a broader scope and the other one more focused one.

The user performance in the search tasks was graded by an expert in a 1 to 5 Likert scale. The improved system resulted in better performance in the focused task (3.1 for improved, 2.2 for baseline, p = 0.2) but worse in the broad task (3.0 for improved, 3.8 for baseline, p = 0.1).

The improved system had a better ResQue score (36.0 for improved, 32.7 for baseline, p = 0.7) and had better score in most questions (answered in 1 to 5 Likert scale).

Imp.	Bas.	Question (15-question ResQue questionnaire)
3.1	3.0	The items recommended to me matched what I was searching for
3.7	3.4	The recommender system helped me discover new items
4.2	**4.3**	The items recommended to me are diverse
3.4	3.2	The layout of the recommender interface is adequate
2.7	2.3	The recommender explains why the items are recommended to me
3.4	2.6	The information provided for the recommended items is sufficient
3.1	2.8	I found it easy to tell the system what I want / don't want to find
4.1	4.0	I became familiar with the recommender system very quickly
3.4	3.1	I found it easy to modify my search query in the recommender
3.1	2.9	I understood why the items were recommended to me
3.3	3.0	Using the recommender to find what I like is easy
3.4	3.4	The recommender gave me good suggestions
3.1	2.9	Overall, I am satisfied with the recommender
3.3	3.3	The recommender can be trusted
3.7	3.5	I would use this recommender again, given the opportunity

The users were interviewed after using the system. 7 out of 10 users reported that the visualized prediction helped them in the task. Majority of the users preferred the improved system: 5 users preferred the improved system overall, 2 had mixed preferences, 1 preferred the baseline overall and 2 had no explicit preference.

Future Work

We intend to carry out a larger user study with the next generation of the SciNet system to confirm the experimental results.

The improved control the user has on the system seems to restrict the performance in broad exploratory tasks that benefit from exploration. Could it be possible to conserve the level of exploration while still giving the user the improved power to control?