Dealing with Concept Drift in Exploratory Search: An Interactive Bayesian Approach

Antti Kangasrääsiö
Yi Chen
Dorota Głowacka
Samuel Kaski
first.last@hiit.fi

Problem Setting

Exploratory search

- The user initially has some knowledge of the search topic but not enough to reduce the task into simple fact retrieval tasks
- The user has to learn while searching, iteratively reformulating a hypothesis of what information would satisfy her information need and where to find it
- This setting makes it difficult for the user to directly formulate good search queries

A recently developed search system called SciNet (IUI 2013, 2015) helps the user by allowing interactive formulation of the search query through relevance feedback on a visualized user model. However, the user model makes many implicit assumptions:

- All user feedback is assumed equally accurate
- The user is assumed to make no mistakes in giving feedback
- No learning or change in search interests is assumed to occur

Our Contribution

Hypothesis: it might be useful for the user to be able to see the feedback history as a whole and to get suggestions on which feedback are likely in need of adjustment.

In this paper we present

- Timeline interface that visualizes the user's feedback history
- User model that estimate both user's current interest and the accuracy of feedback

We demonstrate that

- The new user model improves performance over a baseline model
- Users report that the new interface provides usability improvements
- Users give more feedback and interact more with the new interface

User Model

We assume that the user's interest can be approximately described with a linear Gaussian model, where the accuracy of feedback given by the user may be different for each observation. This gives us the following model:

$$y_i \sim Normal(x_i\phi, \sigma^2/w_i),$$

where *y* is relevance, *x* is item feature vector, *phi* are linear model parameters, *sigma* is variance and *w* is feedback accuracy. We further assume priors for the parameters:

$$\phi_j \sim Normal(\mu_\phi, \lambda_\phi),$$
$$\sigma^2 \sim InverseGamma(\alpha_{\sigma^2}, \beta_{\sigma^2}),$$
$$w_i \sim Gamma(\alpha_{w_i}, \beta_{w_i}), w_i^{fix} \sim Delta(1.0),$$

where the distribution of *w* depends on whether the user has marked that feedback as accurate. We approximate the posterior distribution of the model parameters using mean-field variational inference.

User Interface

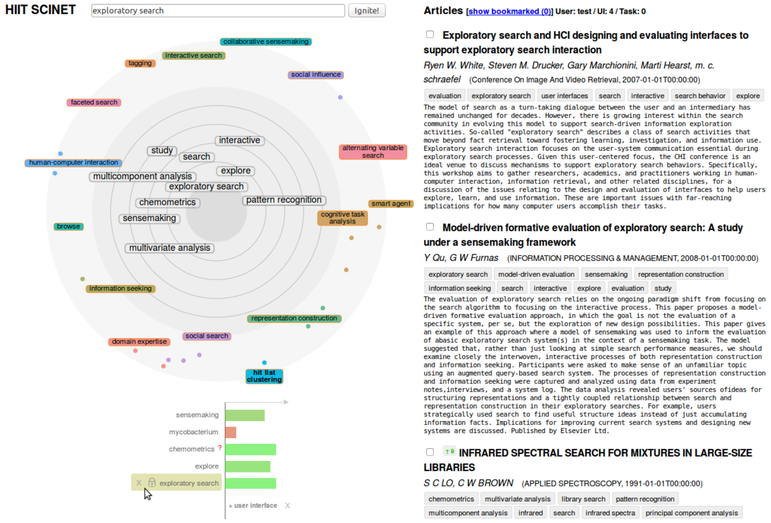

We extended the SciNet search interface with a timeline interface function. The timeline:

- Visualizes the history of relevance feedback
- Provides suggestions on what feedback to adjust (based on estimated accuracy)
- Allows user interaction with the past feedback
  - Delete a feedback from the model
  - Revise relevance value given to a keyword
  - Access keywords from past search sessions
  - Mark feedback as accurate (indicate false positive suggestions)

Simulation Study

We conducted an experiment with a simulated user. The user was searching for newsgroup articles using noisy relevance feedback. We compared the performance of the new model to an oracle and to a baseline that did not estimate accuracy of feedback.

Left graph

- User did not make corrections to past feedback

Center graph

- Search engine made suggestions on what feedback to correct
- User corrected feedback values is case of true positive suggestions
- User marked feedback accurate in case of false positive suggestions

Right graph

- As above, but user did not mark feedback accurate in case of false positives

User Study

We ran a user study where we compared the new interface with a baseline where the timeline was hidden. The study had 4 participants who performed 4 tasks each.

Average number of different user actions per task (20 min) during the experiment:

Interface	Keyword queries	Relevance feedbacks on radar	Relevance feedbacks on timeline	Feedbacks deleted	Feedbacks marked as accurate	Total keyword interactions
Baseline	4.0	5.4	N/A	N/A	N/A	5.4
Timeline	4.0	5.6	1.0	1.6	0.6	8.9

After each task set, we conducted a semi-structured interview with the user

- Main benefits of the timeline interface:
  - It is easier to understand what feedback affects the results
  - It is easy to re-find keywords both from current and past search sessions
  - It helps in the search process as it is easy to "go back" by deleting feedback
- The users also reported the following drawbacks:
  - The red question mark icon made the user feel as if she had made an error
  - The user felt like the system was not fully under her control when the accuracy of feedback was changed automatically
  - The task time limit made them avoid functionality they were not familiar with, as their focus was on performing the task well

Conclusion

- We have presented a user model that is able to take into account concept drift when a search intent model is refined iteratively
- We have presented a timeline interface that offers the user suggestions on what past feedback is most likely in need of adjustment

Overall, the initial results from the user study look encouraging and as the next step we are planning to conduct a larger study with an improved version of the system.

To the best of our knowledge, this is the first time a system has been presented that both models the accuracy of individual user feedback in a search setting and allows the user to directly interact with this model.

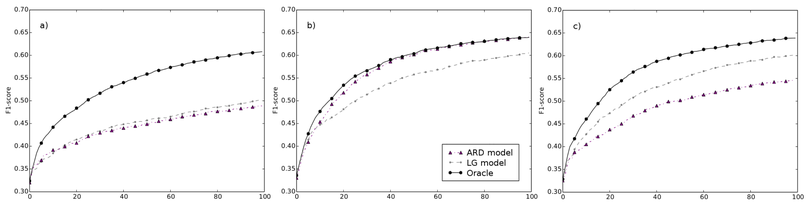

Figure: Results from simulation study. ARD is the new user model, LG is the baseline.