IDEA @ KDD2017
14.8.2017

Interactive intent modelling

Samuel Kaski

Contents

1. Interactive intent modelling for information discovery
2. Interactive knowledge elicitation
3. Multimodal feedback
4. Inferring cognitive user models with ABC

- [Live Demo]

Glowacka et al. IUI 2013,
Ruotsalo et al. Commun ACM 2015, …

HELSINKI INSTITUTE FOR INFORMATION TECHNOLOGY

Glowacka et al. IUI 2013,
Ruotsalo et al. Commun ACM 2015, …

Some problems+solutions in information seeking

1. Underspecified, uncertain and evolving information need
   - interactive on-line-learning interfaces
2. Context bubble
   - exploration/exploitation tradeoff
3. Laziness
   - in giving relevance feedback
   - in pre-specifiying filtering criteria
   - no pain, no gain (but maximize gain/pain by making navigation more natural)

HELSINKI INSTITUTE FOR INFORMATION TECHNOLOGY

Our solution in a nutshell

- Model the user's interests on-line
- Exploration-exploitation tradeoff when suggesting new
- Interactive visualization of the estimated interests
  - for the user to navigate
  - for the system to collect "feedback"

Learning user intents/interests

Assume: Interests = keywords

Represent *i* th keyword by $\mathbf{k}_i$, where the *j*th dimension is 1 if keyword *i* occurs in document *j* ("bag of documents"; plus tf-idf)

Assume relevance feedback is a linear function,

$$\mathbb{E}[r_i] = \mathbf{k}_i^\top \mathbf{w}$$

Exploration-exploitation: Show the user keywords *i* with the highest upper confidence bound (LinRel, Auer 2002): $\hat{r}_i + \alpha\sigma_i$

Sample experiment in Information seeking

- 60,000,000 scientific abstracts
- User's task: Scientific writing scenario; collect material for an essay on a given topic (semantic search or robotics)
- Ground truth: Expert evaluations
- 30 users

THOMSON REUTERS

IEEE Xplore® DIGITAL LIBRARY

Springer science+business media

HELSINKI INSTITUTE FOR INFORMATION TECHNOLOGY

ACM DL DIGITAL LIBRARY

Information seeking results

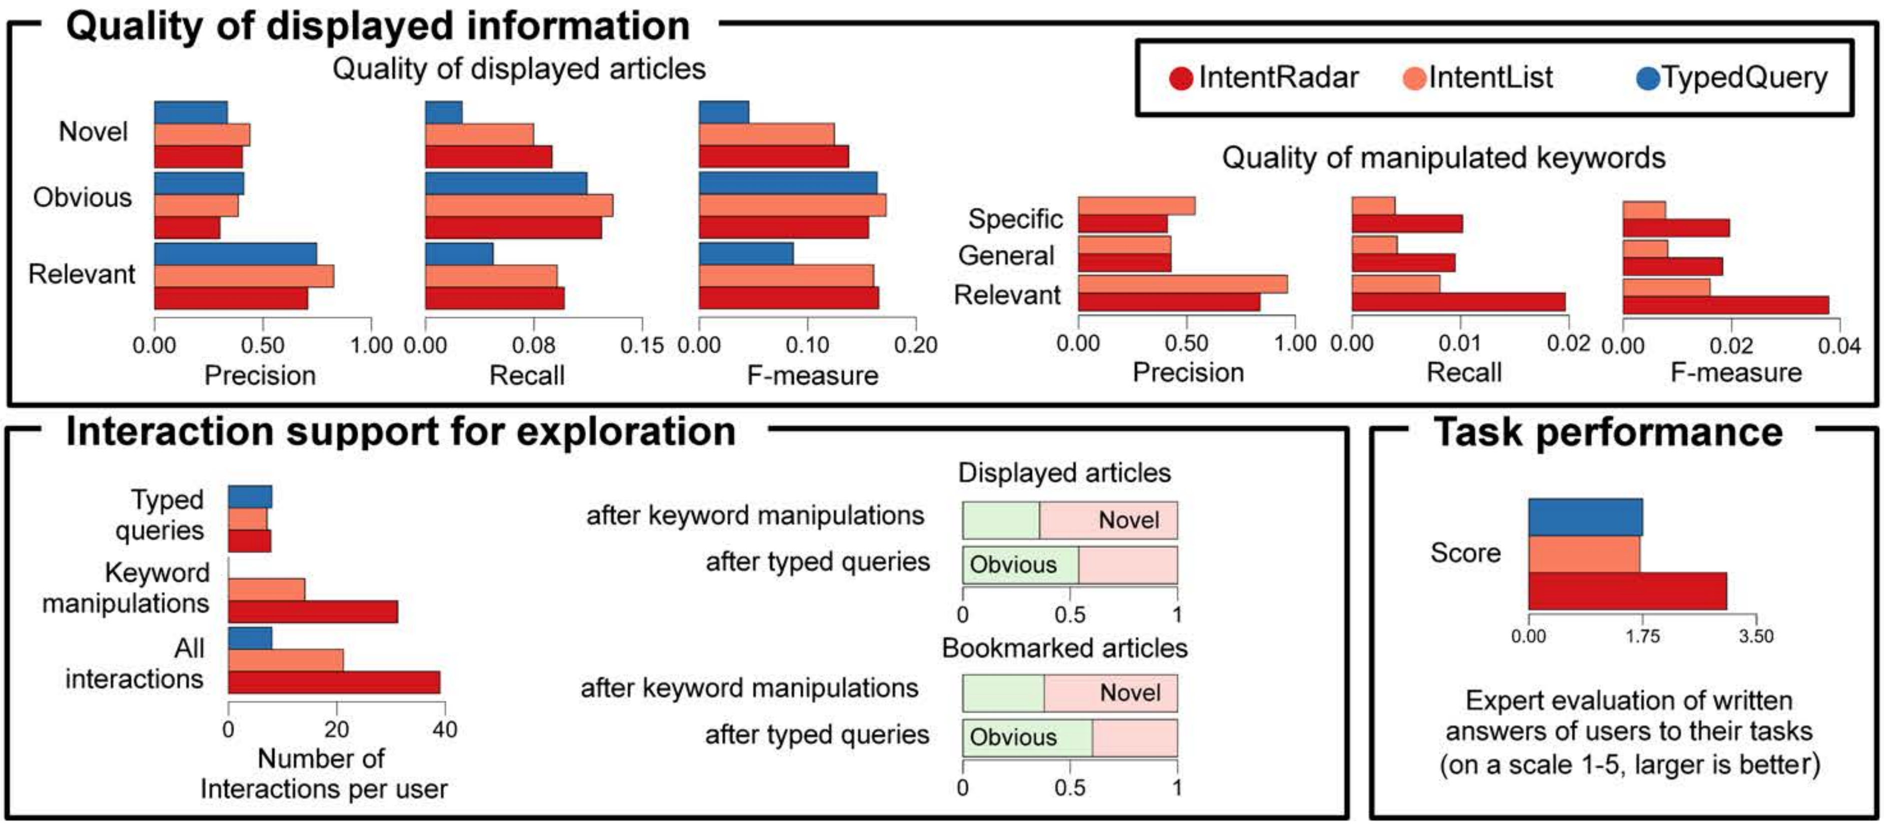

HELSINKI INSTITUTE FOR INFORMATION TECHNOLOGY

References

T. Ruotsalo, G. Jacucci, P. Myllymäki, and S. Kaski. Interactive intent modeling: Information discovery beyond search. ***Communications of the ACM***, 58(1):86–92, 2015.

T. Ruotsalo, J. Peltonen, M. J. A. Eugster, D. Glowacka, K. Konyushkova, K. Athukorala, I. Kosunen, A.Reijonen P. Myllymäki, G. Jacucci, and S. Kaski. Directing exploratory search with interactive intent modeling. In *Proceedings of* ***CIKM 2013****, the ACM International Conference of Information and Knowledge Management*. ACM.

D. Glowacka, T. Ruotsalo, K. Konyushkova, K. Athukorala, S. Kaski, and G. Jacucci. Directing exploratory search: Reinforcement learning from user interactions with keywords. In Proceedings of **IUI'13**, International Conference on Intelligent User Interfaces, pages 117-128, New York, NY, 2013. ACM. Best paper award.

T. Ruotsalo, K. Athukorala, D. Glowacka, K. Konyushkova, A. Oulasvirta, S. Kaipiainen, S. Kaski, and G. Jacucci. Supporting exploratory search tasks with interactive user modelling. In *Proceedings of* ***ASIST 2013****, the 76th ASIS&T Annual Meeting*.

+ many more recent papers

HELSINKI INSTITUTE FOR INFORMATION TECHNOLOGY

Contents

Interactive expert knowledge elicitation

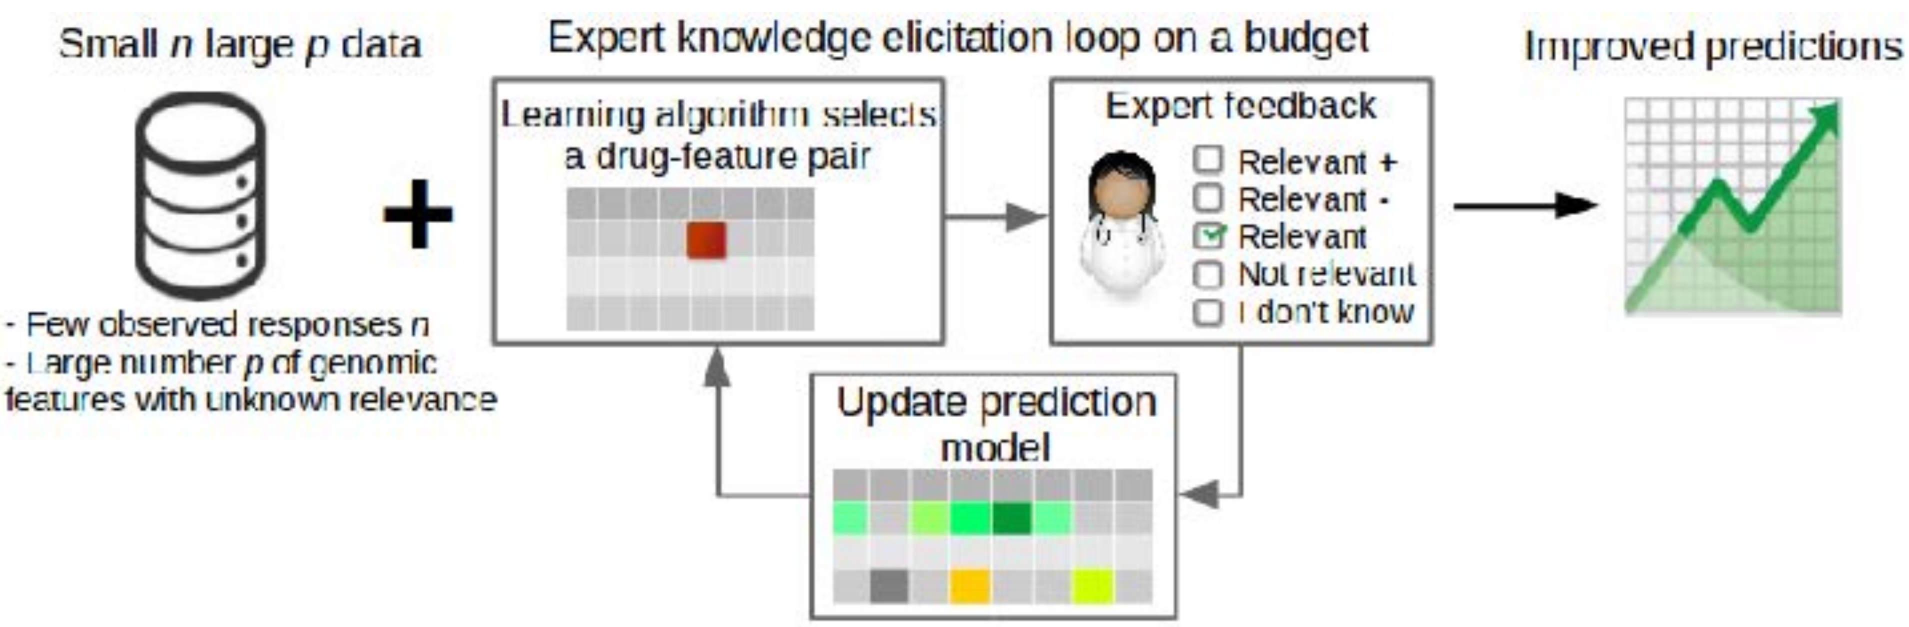

Prediction given “small n, large p”

- e.g. prediction of drug responses based on high dimensional patient profiles.

- Existing ways to mitigate “small n, large p”
  - strong informative modelling assumptions
  - collecting more data
  - expert prior elicitation

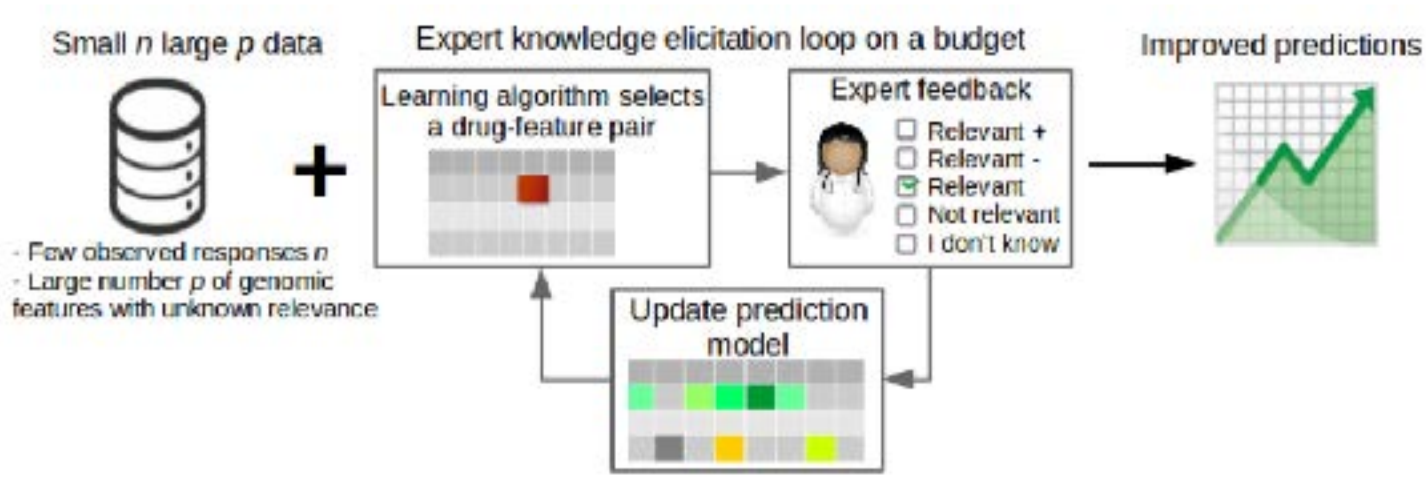

Approach 1: separate user model

- Use multi-armed bandit model as in information discovery:
  - keywords -> patient features
  - relevance for retrieval -> relevance for prediction of treatment effectiveness
- Good: explicitly aims at balancing between exploration and exploitation
- Problem: Does not directly aim at maximizing prediction accuracy

Approach 2: Sequential experimental design

Formulate knowledge elicitation as a probabilistic inference process, where expert knowledge is sequentially queried to improve predictions.

User interaction as inference

1. An observation model $p(y|x, \theta, \phi_y)$
2. A feedback model for user's knowledge $p(f|\theta, \phi_f)$
3. A prior model $p(\theta, \phi_y, \phi_f)$
4. A query algorithm that facilitates gathering f iteratively from the user.
5. Update process of the model after user interaction.

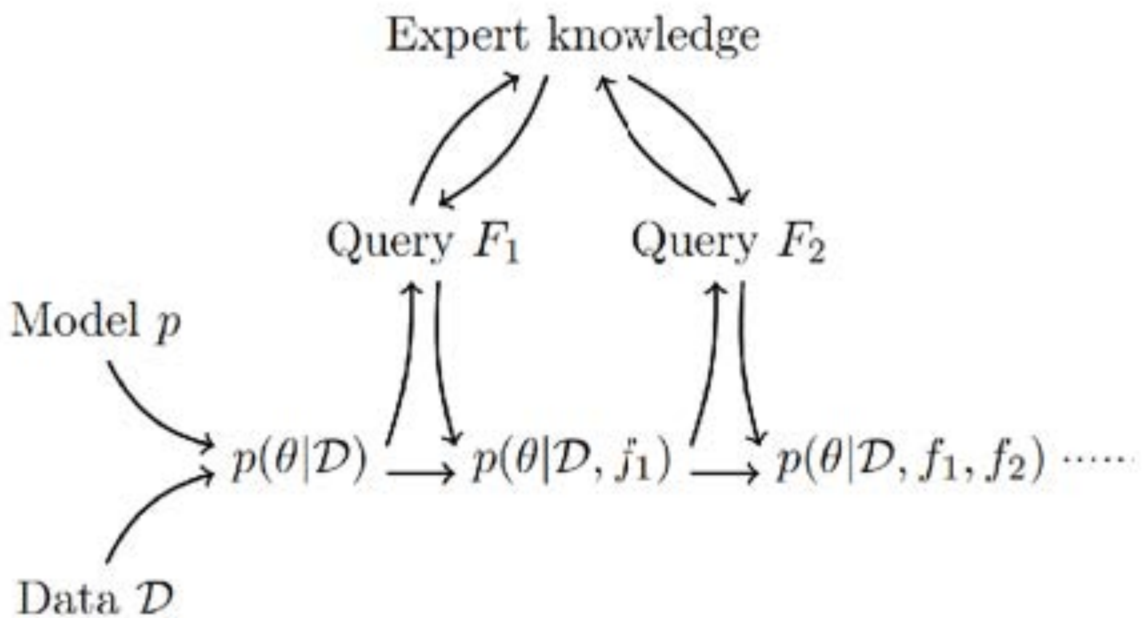

Case study: drug sensitivity predictions given genomic data

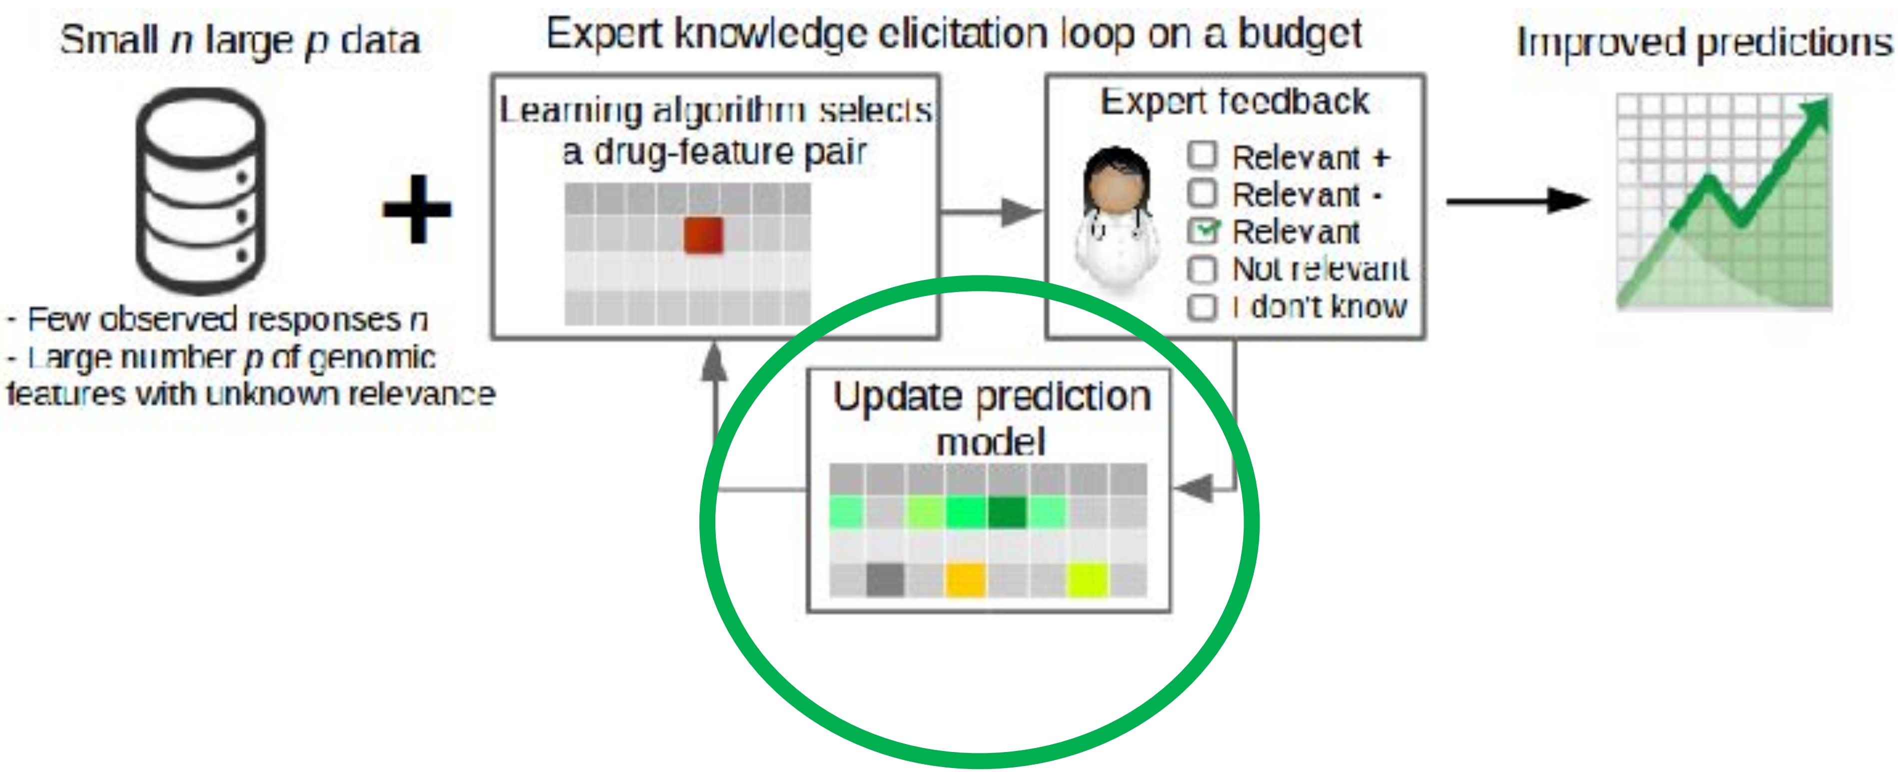

Sparse regression with feedback observation model

- $y_{n,d} \sim \mathrm{N}\left(w_d^\top x_n, \sigma_d^2\right)$

  $y_{n,d}$ Response of nth patient to drug d

  x_n Genomic features of the patient

  σ_d^2 Residual variance

- Sparsity inducing spike-and-slab prior

  $w_{d,m} \sim \gamma_{d,m} \mathrm{N}\left(0, \tau_{d,m}^2\right) + \left(1 - \gamma_{d,m}\right)\delta_0$

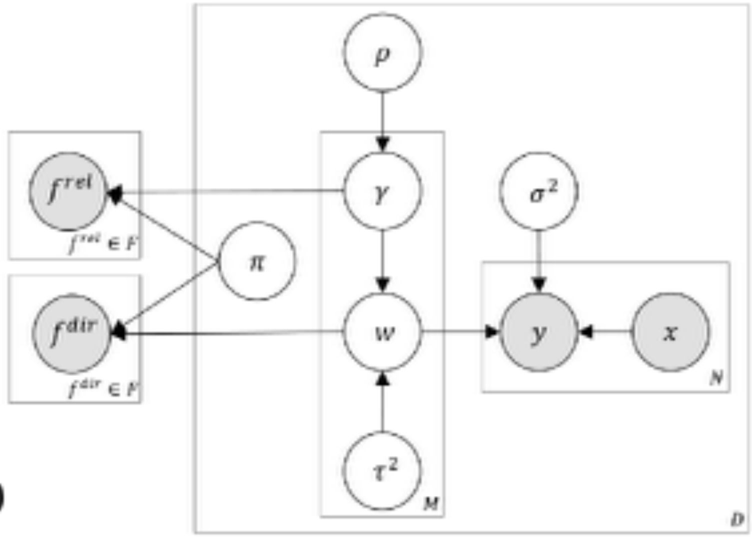

Sparse regression with feedback observation model

- Feedback observation model
  - Relevance feedback

$$f^{rel}_{d,m} \sim \gamma_{d,m}\,\mathrm{Bernoulli}(\pi^{rel}_d) + (1-\gamma_{d,m})\,\mathrm{Bernoulli}(1-\pi^{rel}_d)$$

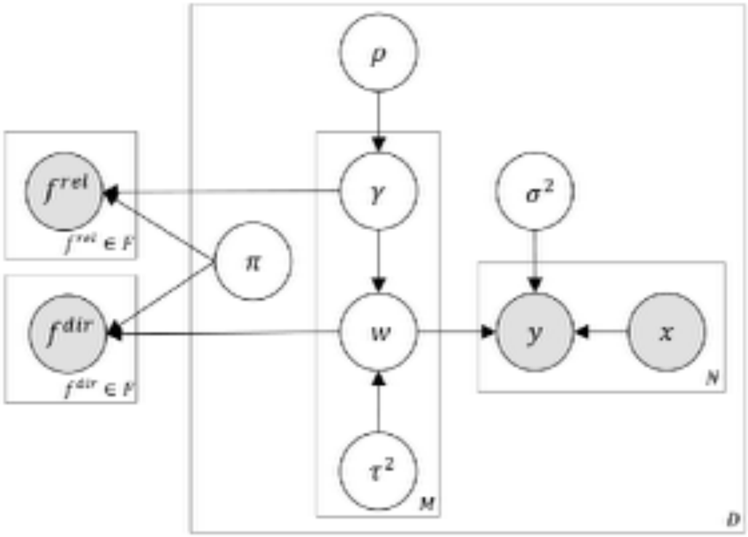

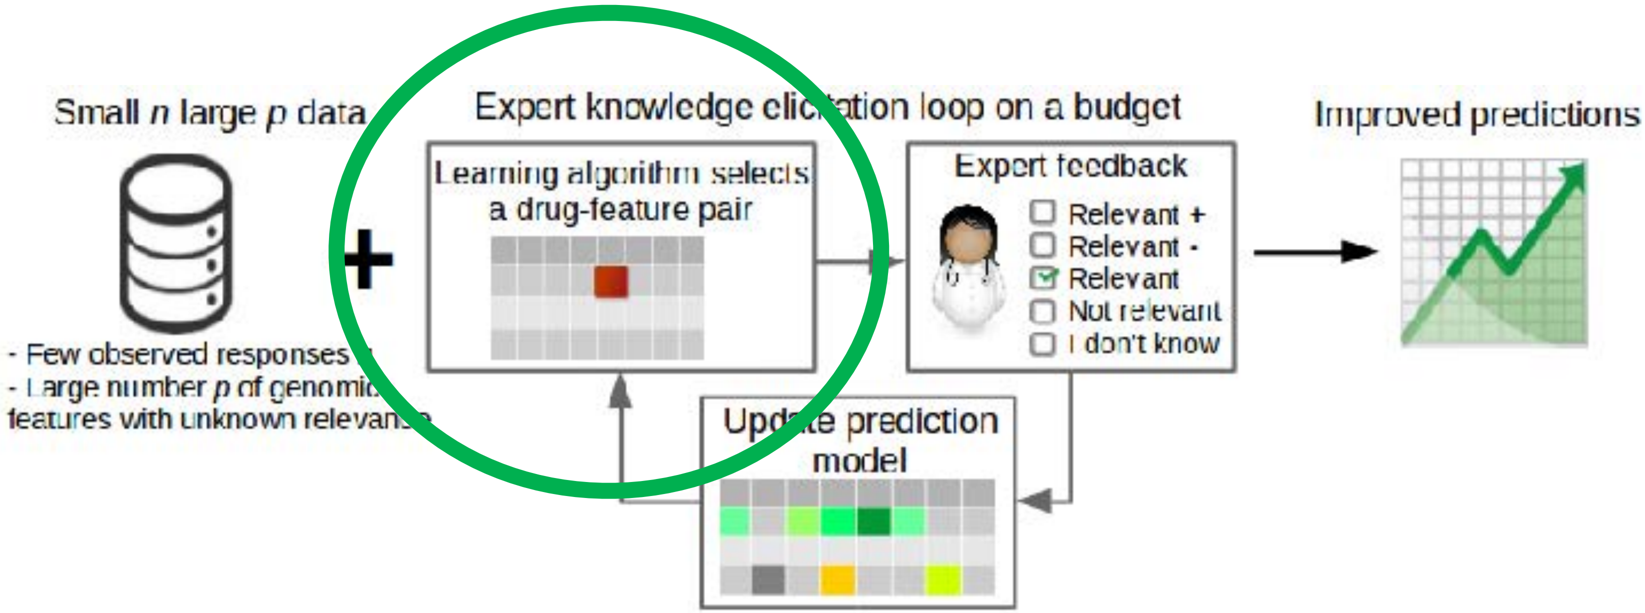

Small n large p data
- Few observed responses
- Large number p of genomic features with unknown relevance
Expert knowledge elicitation loop on a budget
Learning algorithm selects a drug-feature pair
Expert feedback
Relevant +
Relevant -
Relevant
Not relevant
I don't know
Update prediction model
Improved predictions

Query algorithm

- Formulate choosing of the query as a sequential experimental design problem. Aim at maximal expected information gain about predictions:

$$\arg\max_{(d,m)} E_{\tilde{f}_{d,m}|D_{t-1}}\left[\sum_{n=1}^{N} KL[p(\tilde{y}_d|\boldsymbol{x}_n, D_{t-1}, \tilde{f}_{d,m})||p(\tilde{y}_d|\boldsymbol{x}_n, D_{t-1})]\right]$$

Computation

Problems:

- No closed form solution is available for
  - Posterior distribution
  - Predictive distributions
  - Information gain maximization
- High dimensionality
- Needs to be fast for user interaction

Solution:

- Deterministic posterior approximations:
  - Expectation propagation to approximate the spike-and-slab prior and the feedback models (Minka 2011, Hernández-Lobato 2015)
  - Variational Bayes to approximate the residual variance
- Partial/single-step EP updates for candidate evaluation (Seeger 2008)

Simulations - synthetic data (1/2)

- 10 training data, 100 features (10 relevant, 90 zeros).

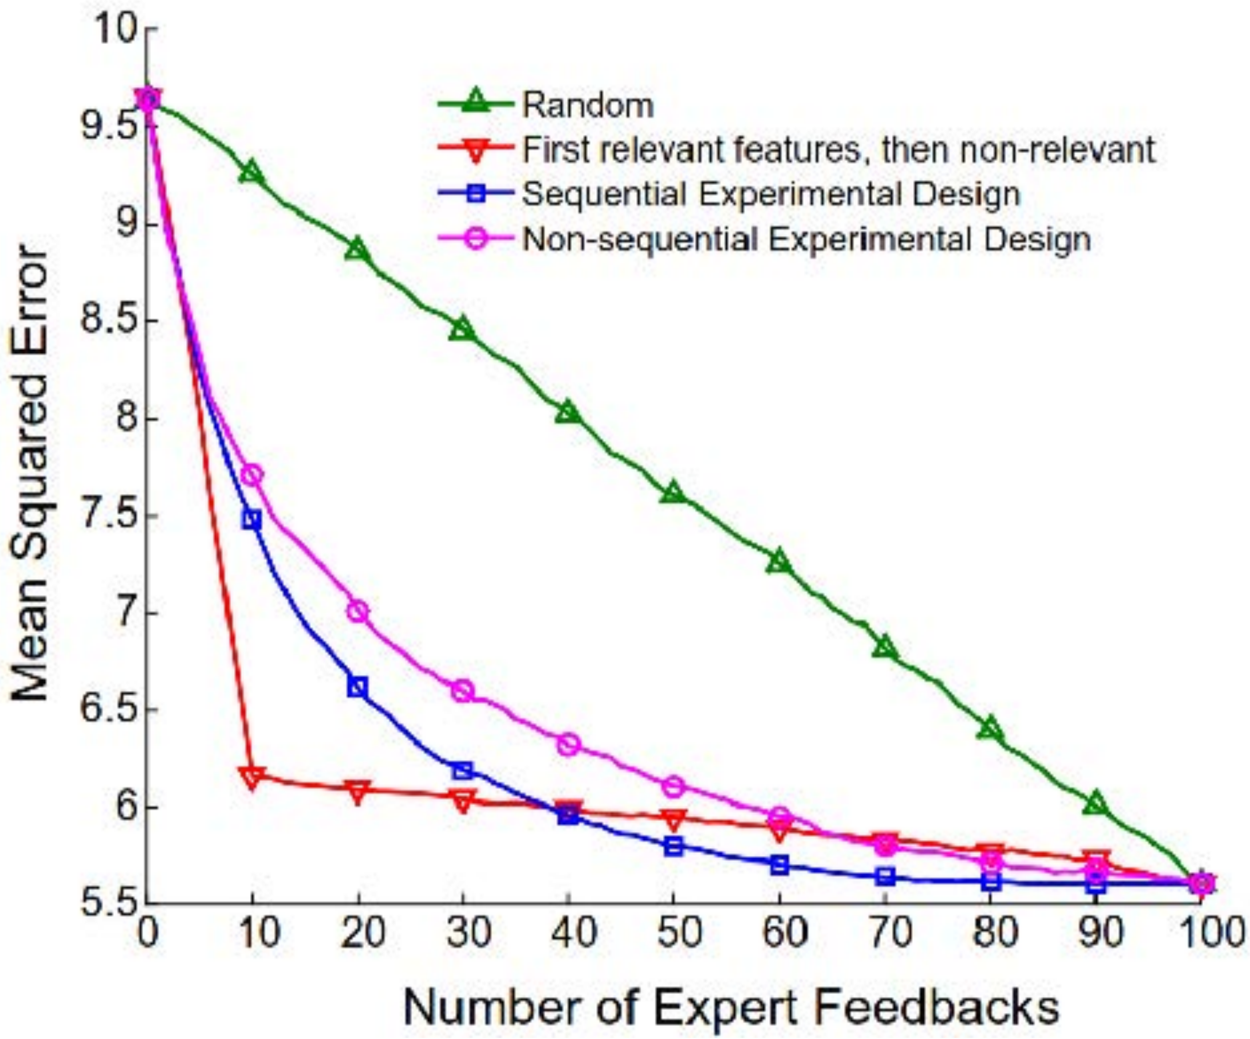

Simulations - synthetic data (2/2)

- 10 training data, 10 relevant features.
- Increasing dimensionality (hence also increasing sparsity)

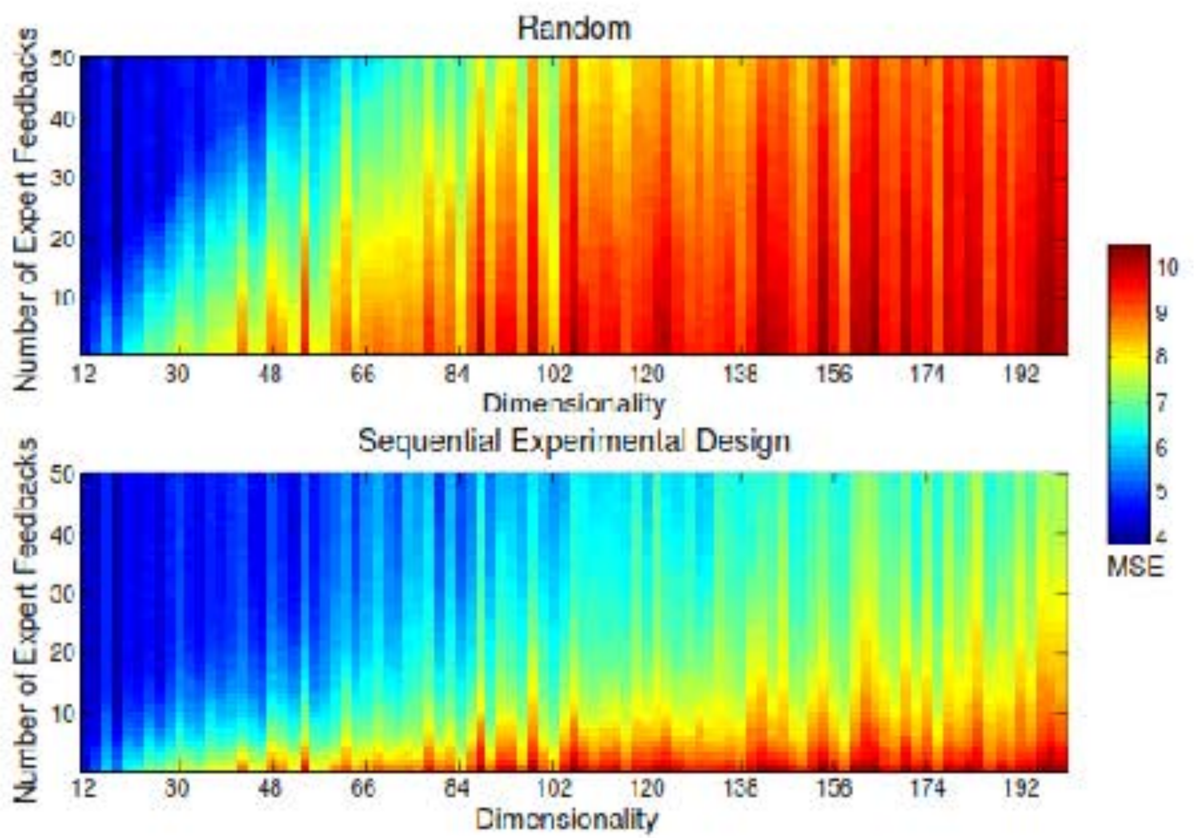

(b) Feedback on feature relevance

Results –
Sequential knowledge elicitation **reduces the number of queries** required from the expert

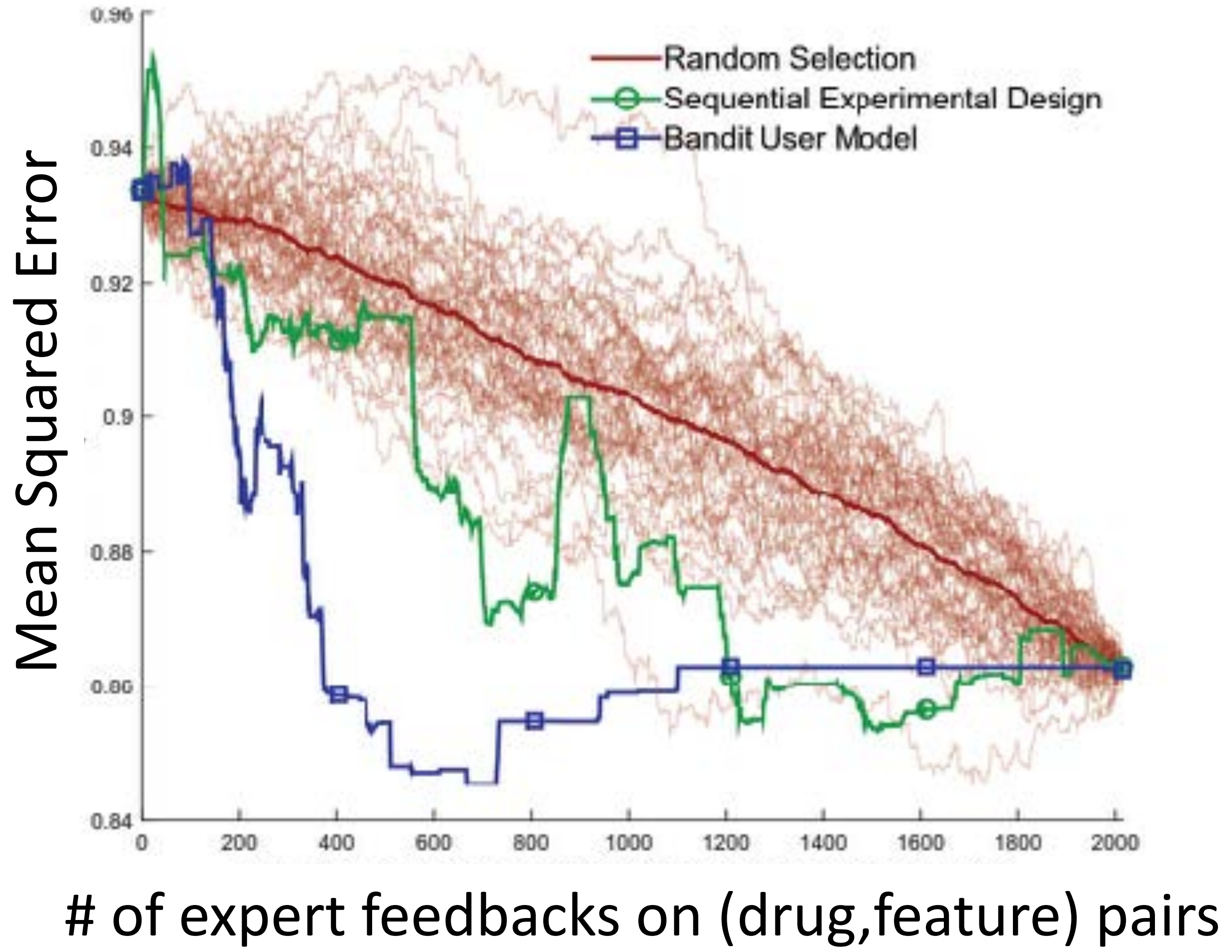

Aalto University
School of Science

arXiv:1705.03290, 2017

Improving drug sensitivity predictions in precision medicine through active expert knowledge elicitation

Iiris Sundin [1,*], Tomi Peltola [1], Muntasir Mamun Majumder [2], Pedram Daee [1], Marta Soare [1], Homayun Afrabandpey [1], Caroline Heckman [2], Samuel Kaski [1,†,*] and Pekka Marttinen [1,†,*]

Mach Learn
DOI 10.1007/s10994-017-5651-7

CrossMark

Knowledge elicitation via sequential probabilistic inference for high-dimensional prediction

Pedram Daee[1] · Tomi Peltola[1] · Marta Soare[1] · Samuel Kaski[1]

IUI 2017 • Interactive Machine Learning and Explanation March 13–16, 2017, Limassol, Cyprus

Interactive Elicitation of Knowledge on Feature Relevance Improves Predictions in Small Data Sets

Luana Micallef*[,1], Iiris Sundin*[,1], Pekka Marttinen*[,1], Muhammad Ammad-ud-din[1], Tomi Peltola[1], Marta Soare[1], Giulio Jacucci[2], and Samuel Kaski[1]

Contents

1. Interactive intent modelling for information discovery
2. Interactive knowledge elicitation
3. Multimodal feedback
4. Inferring cognitive user models with ABC

“Isn’t it trivial to infer interests? Just monitor where the user looks.”

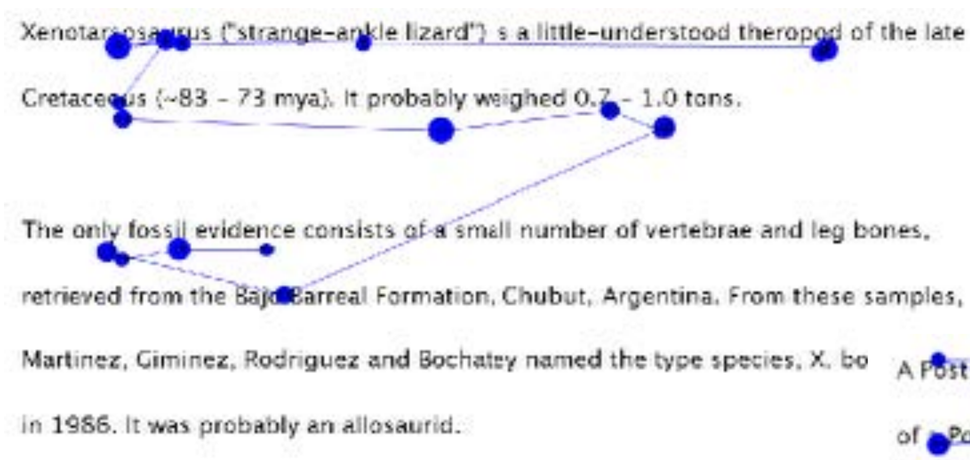

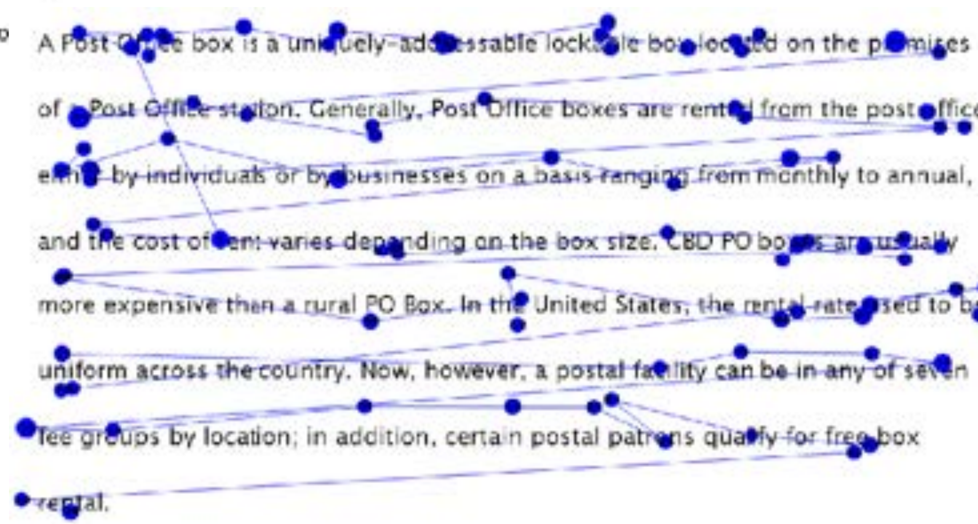

- Accuracy of inferring which titles were relevant: 73% (naive model: 67%)
- Combined with collaborative filtering: 85%

Puolamäki et al., SIGIR 2005

Natural brain-information interfaces

Natural brain-information interfaces - *Recommending information by relevance inferred from human brain signals*

Manuel J. A. Eugster* , Tuukka Ruotsalo*, Michiel M. Spapé*, Oswald Barral, Niklas Ravaja, Giulio Jacucci and Samuel Kaski. * equal contributions

Article published in:
Nature Scientific Reports. 6, 38580; doi: 10.1038/srep38580 (2016)

Aalto University
School of Science

UNIVERSITY OF HELSINKI

HELSINKI INSTITUTE FOR INFORMATION TECHNOLOGY

Eugster et al., SIGIR 2014, Scientific Reports, 2016; Kauppi et al., NeuroImage 2015

Other examples of Augmented Research @ HIIT

Visual Re-Ranking for Multi-Aspect Information Retrieval

Khalil Klouche[1,3], Tuukka Ruotsalo[2], Luana Micallef[2]
Salvatore Andolina[2], Giulio Jacucci[1,2]

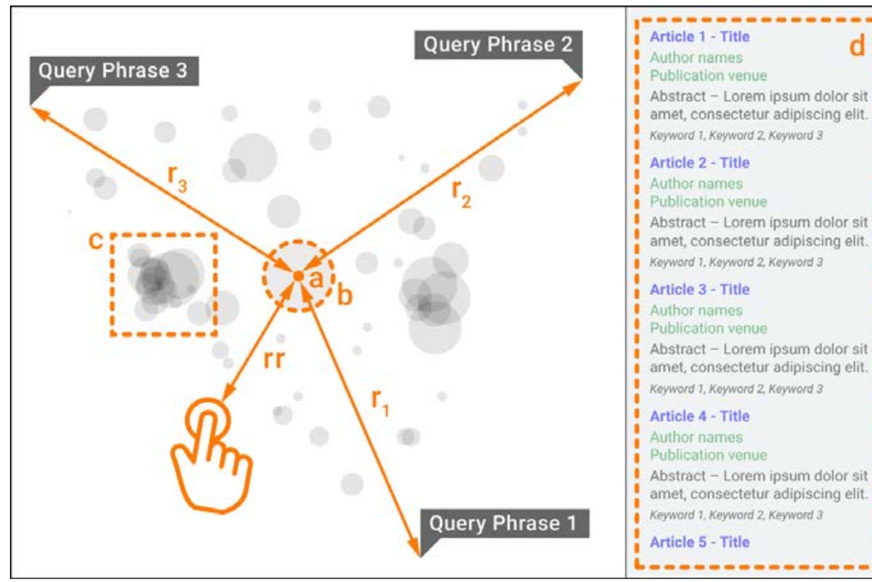

ACM CHIIR 2017

Crowdboard: Augmenting In-Person Idea Generation with Real-Time Crowds

Salvatore Andolina[1], Hendrik Schneider[2,3], Joel Chan[4], Khalil Klouche[2]
Giulio Jacucci[1,2], Steven Dow[5]

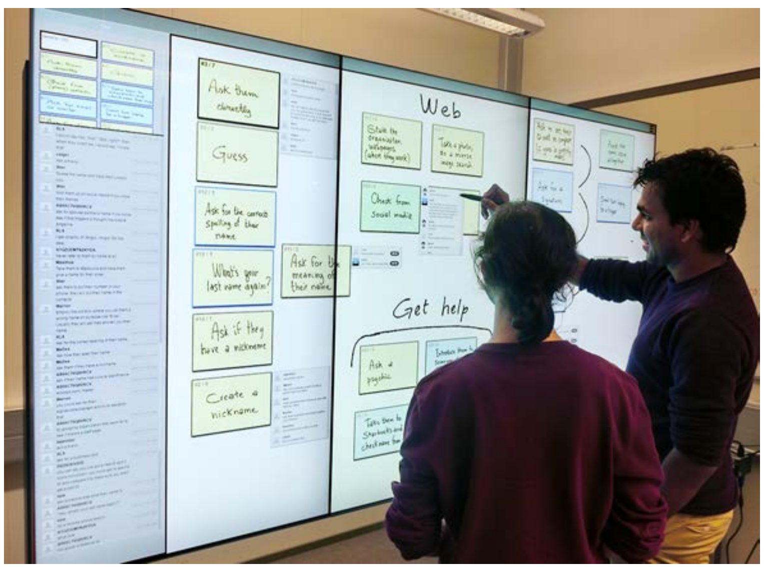

ACM Creativity and Cognition 2017

http://augmentedresearch.hiit.fi

QueryWall: Flexible Entity Search

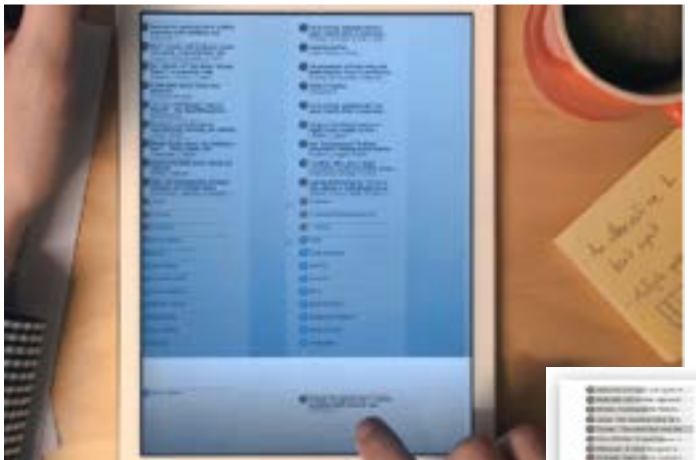

Klouche, K., Ruotsalo, T., Cabral, D., Andolina, S., Belluci, A. and Jacucci, G. Designing For Exploratory Search On Touch Devices. In Proceedings of the 33rd annual ACM conference on Human factors in computing systems (CHI '15). ACM (full paper) (to appear).

Andolina, S., Klouche, K., Peltonen, J., Hoque, M., Ruotsalo, T., Cabral, D., Klami, A., Glowacka, D., Floréen, P. and Jacucci, G. IntentStreams: smart parallel search streams for branching exploratory search. In Proceedings of the 2015 international conference on Intelligent User Interfaces (IUI '25). ACM (short paper) (to appear).

Contents

1. Interactive intent modelling for information discovery
2. Interactive knowledge elicitation
3. Multimodal feedback
4. Inferring cognitive user models with ABC

Inferring Cognitive Models from Data using Approximate Bayesian Computation

Antti Kangasrääsiö[1], Kumaripaba Athukorala[1], Andrew Howes[2], Jukka Corander[3], Samuel Kaski[1], Antti Oulasvirta[4]
[1]Helsinki Institute for Information Technology HIIT,
Department of Computer Science, Aalto University, Finland
[2]School of Computer Science, University of Birmingham, UK
[3]Department of Biostatistics, University of Oslo, Norway
[4]Helsinki Institute for Information Technology HIIT,
Department of Communications and Networking, Aalto University, Finland

CHI 2017

Inverse Reinforcement Learning from Summary Data

Antti Kangasrääsiö [1] Samuel Kaski [1]

arXiv:1703.09700, 2017

Inverse modelling of complex interactive behavior with ABC

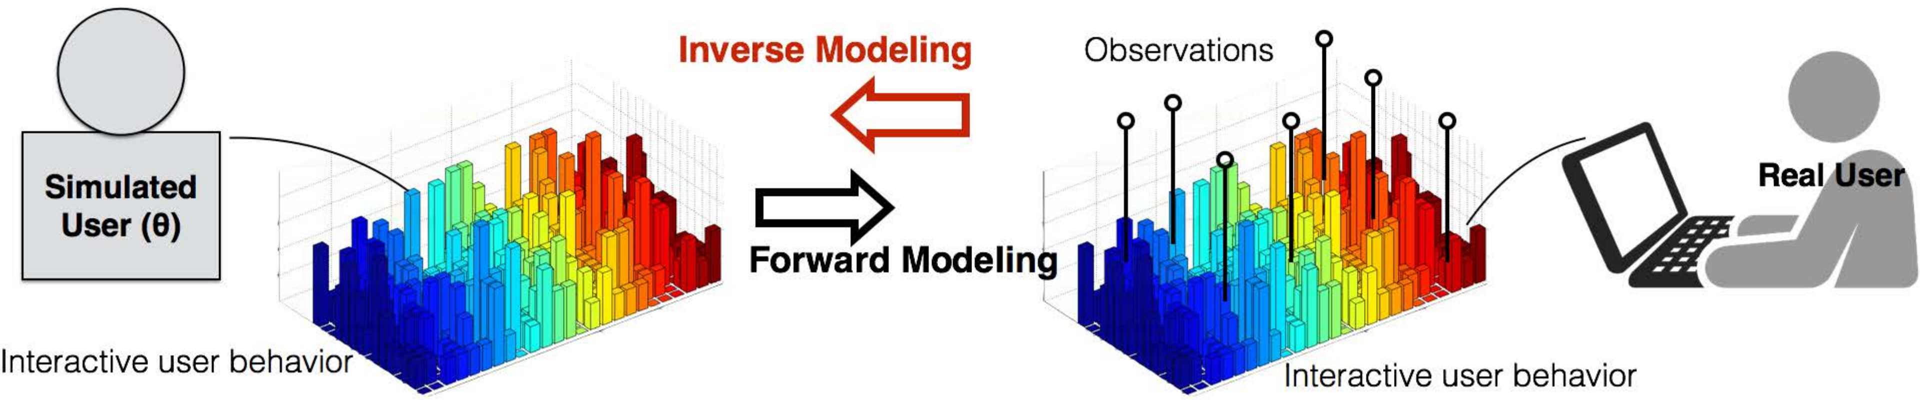

STA in parameter inference: (i) simplified models, (ii) find parameters from literature, or (iii) fit parameters by manual iteration

Big dream: Instead of having to run a laborious user experiment every time a new interface design is tried, run a simulated user experiment.

In other words: Modelling-driven user interface design

Computational rationality

Instead of trying to model *all* aspects of human behaviour, make an assumption:

Computational rationality*:* Assume users behave (approximately) to maximize utility given constraints coming from

- the environment (the interface)
- the goal and
- their own limited (cognitive) capacity.

The simulator is given the constraints. It solves the optimal behavioural policy by reinforcement learning, and then simulates behaviour according to the policy.

Example user task: Menu search

Task: Find a given entry from a menu

Actions: fixate on an item, select the item, quit

Reward for: time used (negative), menu item found / not found

Data: Click time data + possibly eye tracking

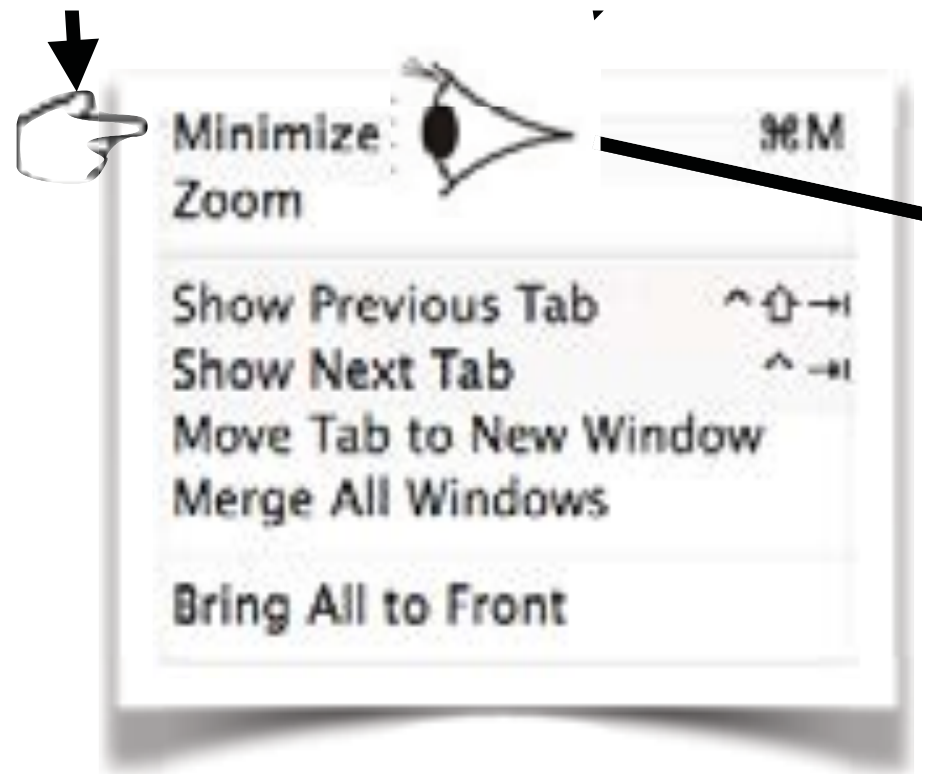

Our task: inverse reinforcement learning given summary data

Infer the parameters given behavioural data: intent, cognitive parameters

More generally: **Inverse reinforcement learning** (IRL) given summary data

- Existing IRL solutions require fully observed state-action sequences
- For summary data would need to integrate over all unobserved paths, which gets intractable.

Approximate Bayesian Computation

- Allows inference when likelihood is difficult or unavailable
- Based on the intuition that similar data are likely to originate from similar processes or parameters
- Observed data compared to simulated

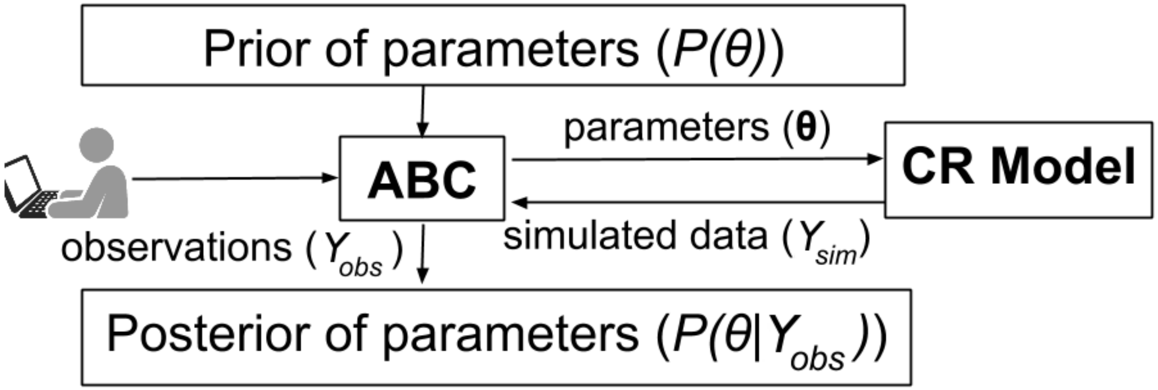

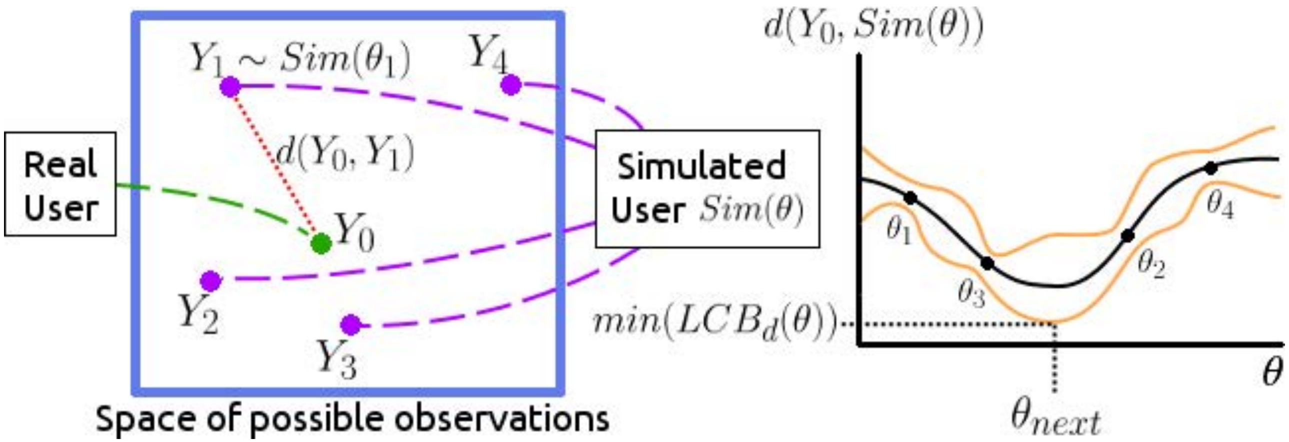

BOLFI: Gutmann & Corander 2016

Results: data distributions

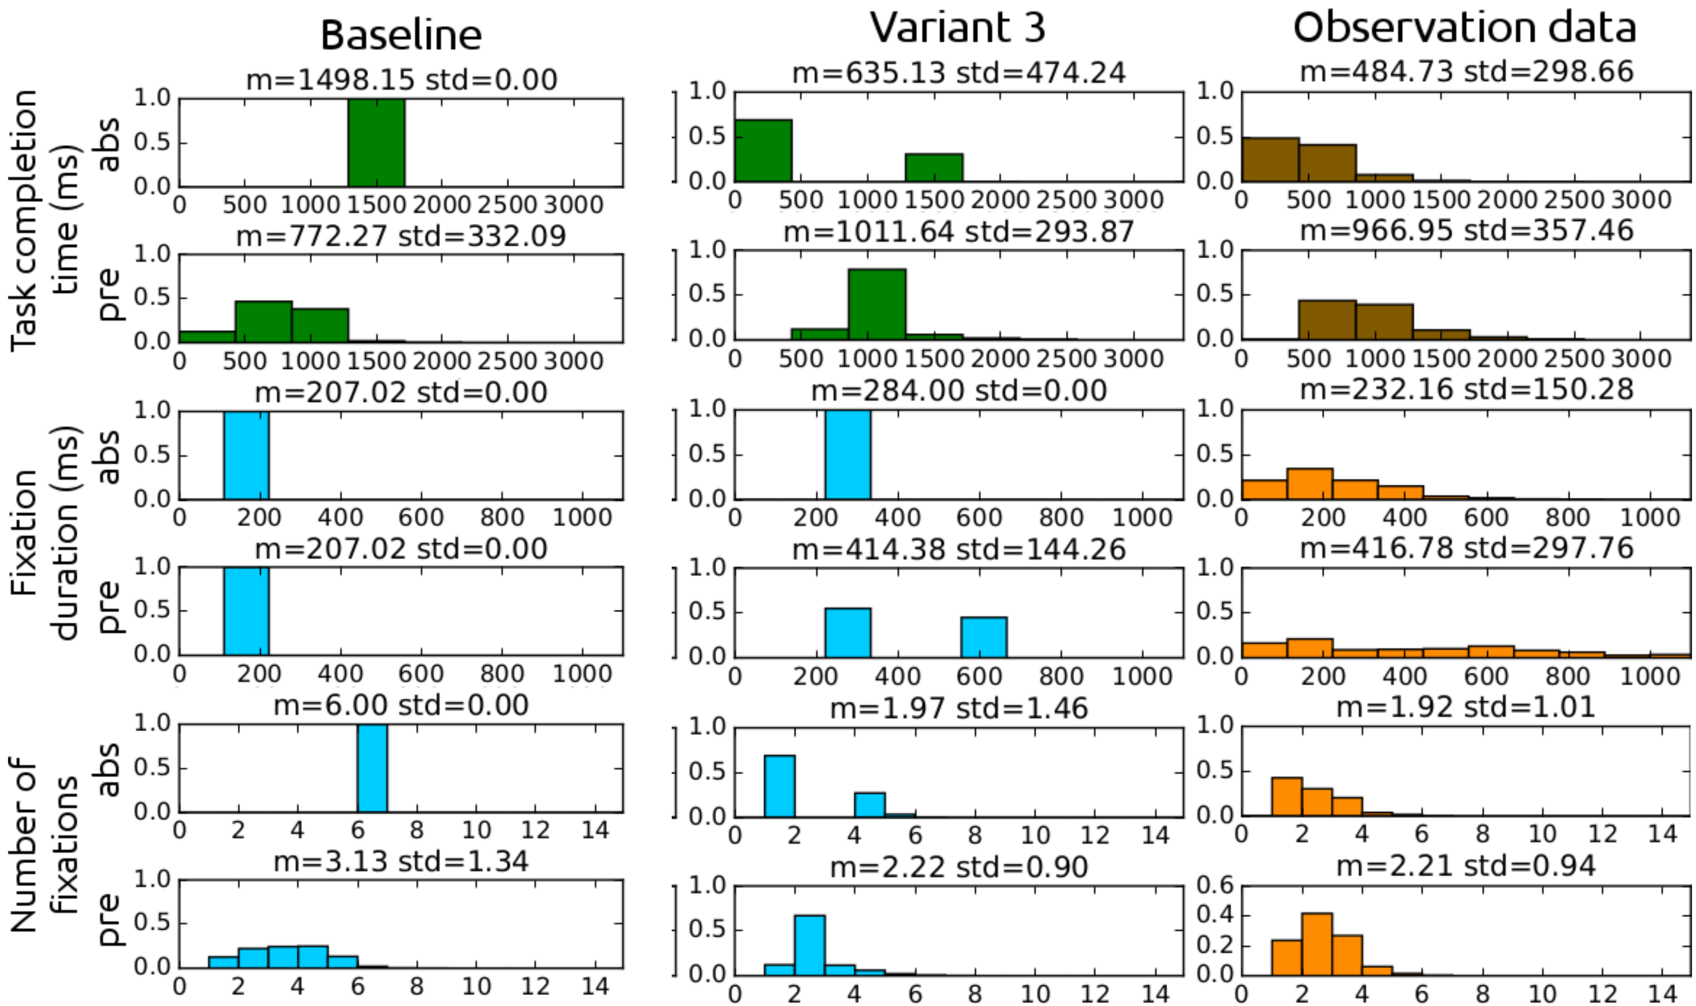

ELFI: ABC for everyone

ELFI = Engine for Likelihood-Free Inference, launched in Dec 2016

Why use ELFI?

- For end users: Bring your own simulator, the engine does the inference, diagnostics and visualization
- For advanced users: Model definition as graphical models; out-of-the-box parallelization; interface in Python
- For developers: Modular community-driven design –> easy to re-use and contribute

ELFI: Engine for Likelihood Free Inference

Jarno Lintusaari[1], Henri Vuollekoski[1], Antti Kangasrääsiö[1], Kusti Skytén[1], Marko Järvenpää[1], Michael Gutmann[2], Aki Vehtari[1*], Jukka Corander[3*], and Samuel Kaski[1*]

arXiv:1708.00707, 2017

elfi.readthedocs.io
pip install elfi

Aalto University

UNIVERSITY OF HELSINKI

Summary

1. Interactive intent modelling for information discovery
   – Simple user model balances exploration-exploitation tradeoff with good results
2. Interactive knowledge elicitation
   – Elicitation was formulated as sequential inference on joint user-prediction model. It improves prediction results on “large p, small n” data.
3. Multimodal feedback
   – Implicit feedback from eye tracking and mind reading is informative but not sufficient to replace explicit feedback yet.

4. Inferring cognitive user models with ABC

– Computational rationality based models require solving a new inverse reinforcement problem, which can be done with ABC & ELFI.

Papers and code available at:

http://research.cs.aalto.fi/pml/

http://augmentedresearch.hiit.fi

Thanks to many students and collaborators, listed earlier in the talk!

ACADEMY OF FINLAND

SEVENTH FRAMEWORK PROGRAMME

Tekes