Aalto University

MANCHESTER 1824

The University of Manchester

Data analysis with humans

Samuel Kaski

FCAI

Generative probabilistic modelling of data

Model of the world $M(\theta)$

$\downarrow$

Data $D = \{x\}$

Problem types

Most of
machine
learning

AutoML

Most of
data mining

Interesting
region for new
machine learning!

Information
retrieval

Visual analytics

Only little

A lot

User needed on-line

Contents

User as a data source

Model of the world $M(\theta)$

Model of the user $M^u(\theta^u)$

$\downarrow$

Data $D = \{x\}$

Data $D^u = \{x^u\}$

Task: prediction for high-dim data

Predictors ***x*** | Outputs ***y***

Precision medicine:	Molecular measurements and other data	Efficacy of treatments

Problem: too little data for estimating the predictor ("small *n*, large *p*")

Sundin et al., Bioinformatics 2018

Interactive expert knowledge elicitation

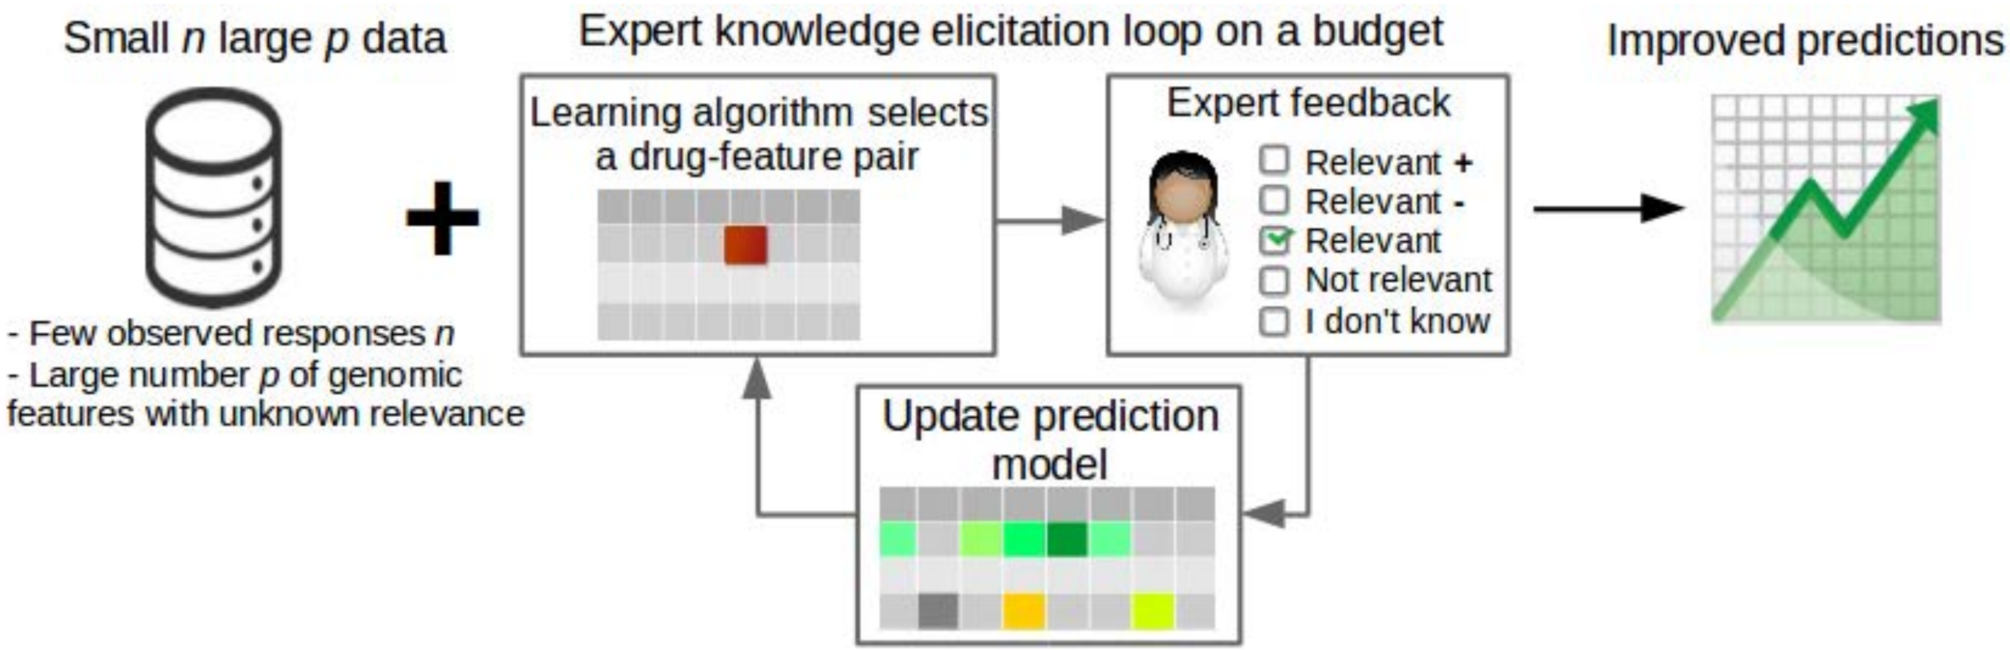

Interactive system brings an expert to the loop

Sundin et al., Bioinformatics 2018

User model

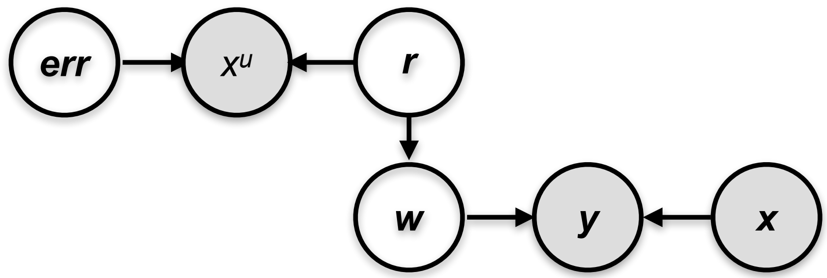

Joint model of user's responses x^u and drug responses $\boldsymbol{y}|\boldsymbol{x}$, coupled by the latent variable r which tells is a variable relevant

User model: assume the user is uncertain about the relevance, and can make an error with probability *err*

Decide the query by sequential experimental design: query the variable that *maximizes expected utility.* Utility is here the expected information gain in $\boldsymbol{y}$

Sundin et al., Bioinformatics 2018

Case: Predict cholesterol levels

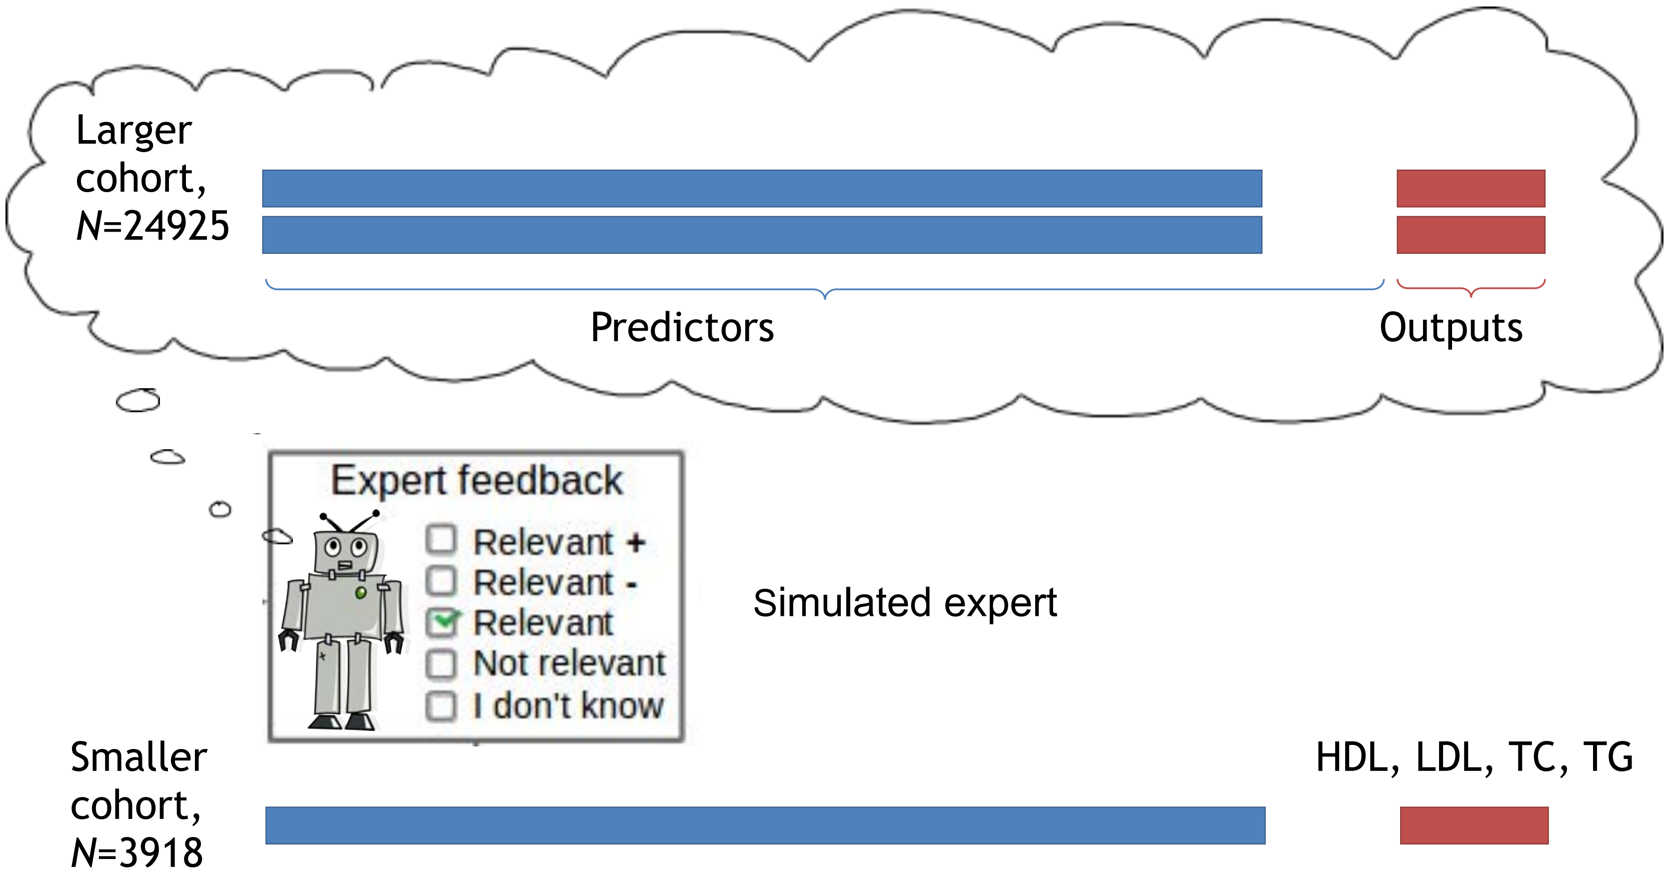

Sundin et al., Bioinformatics 2018

Case: Predict cholesterol levels

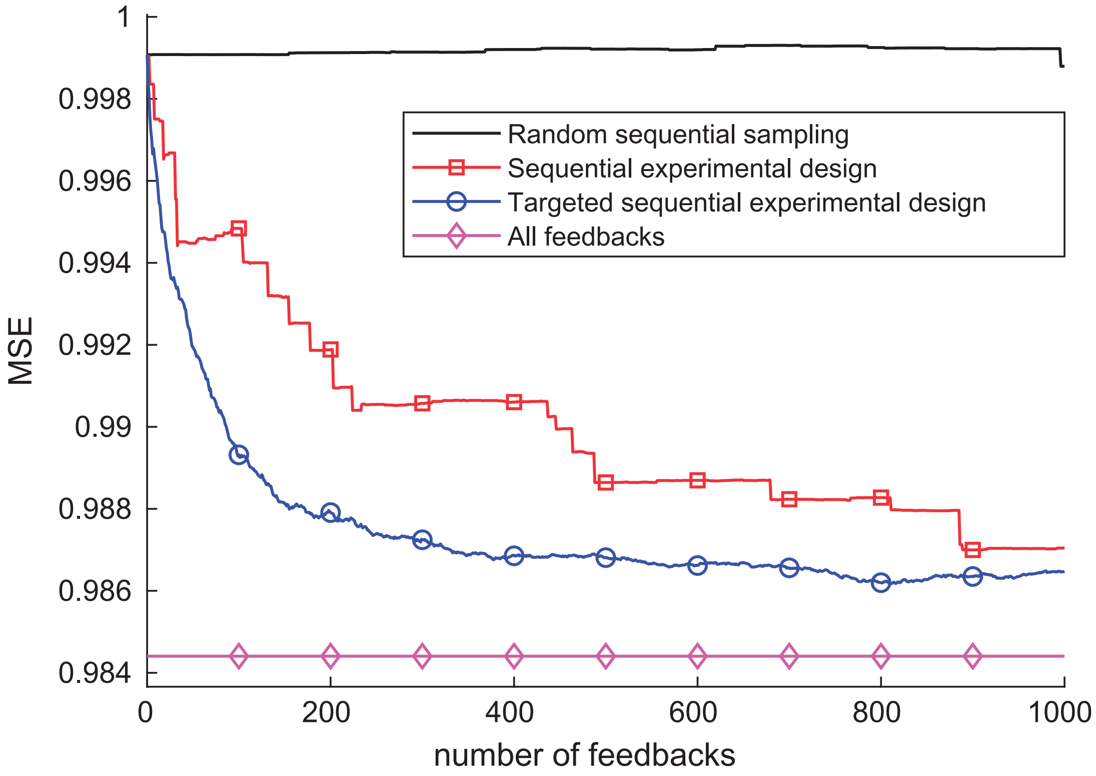

FCAI

fcai.fi

Sundin et al., Bioinformatics 2018

IntentRadar for information retrieval

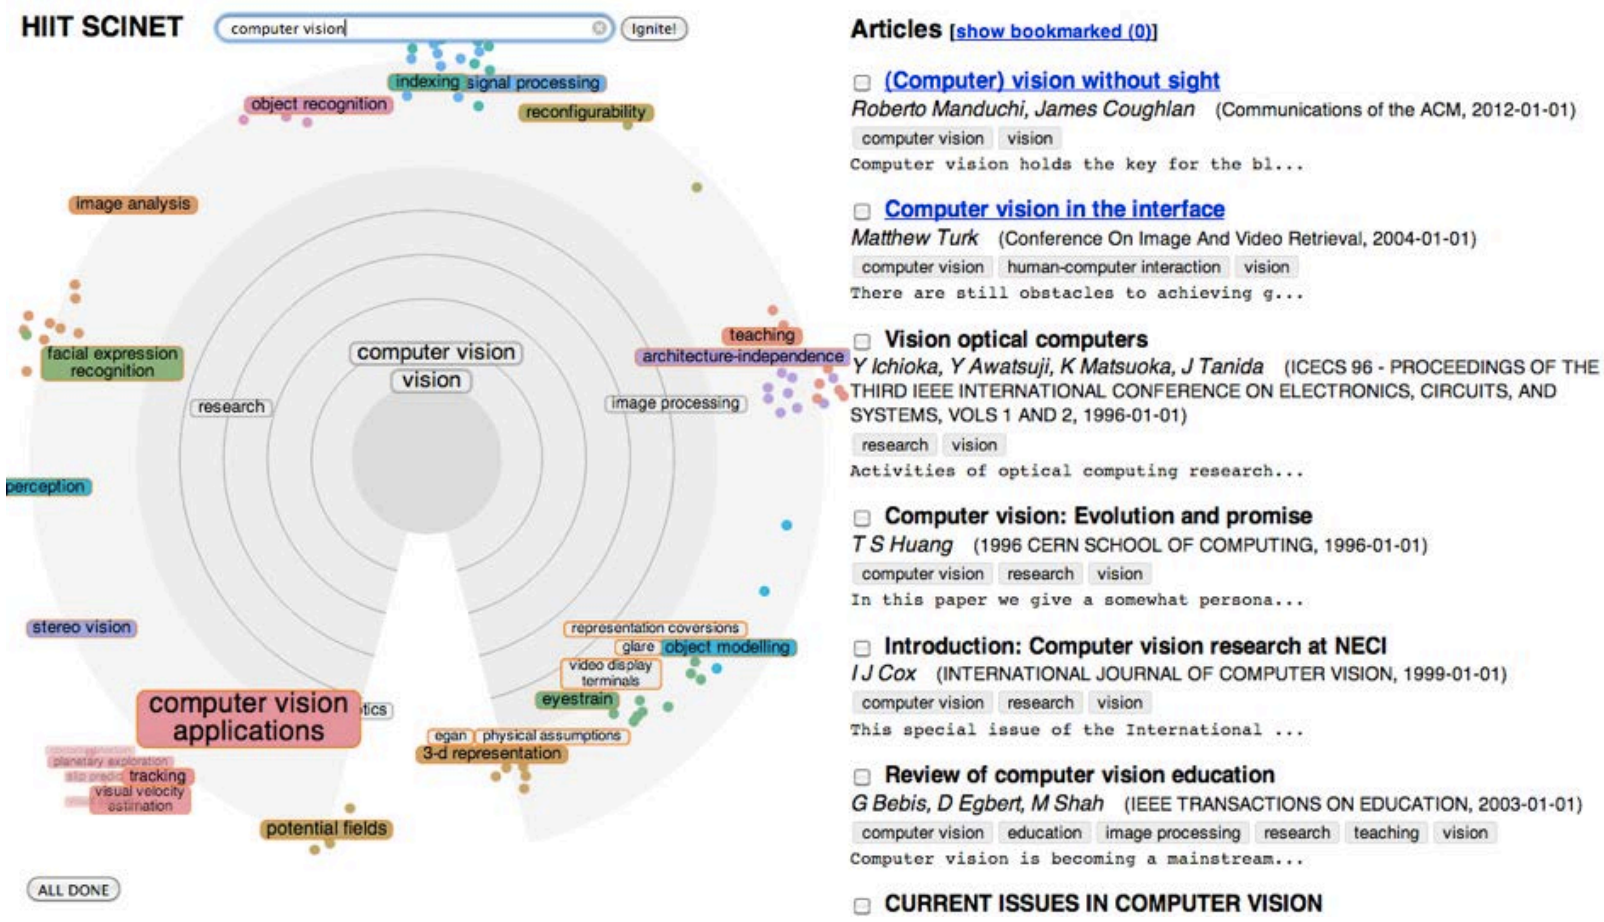

Glowacka et al. IUI 2013,
Ruotsalo et al. Commun ACM 2015, ACM TOIS 2018

Materials science: molecule on a surface

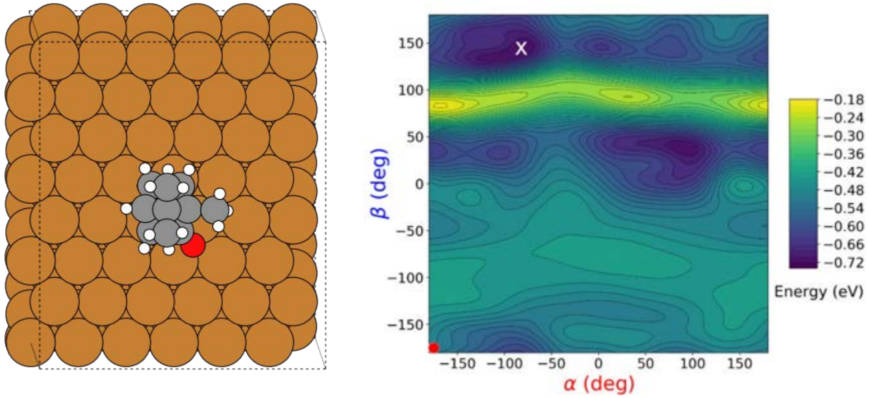

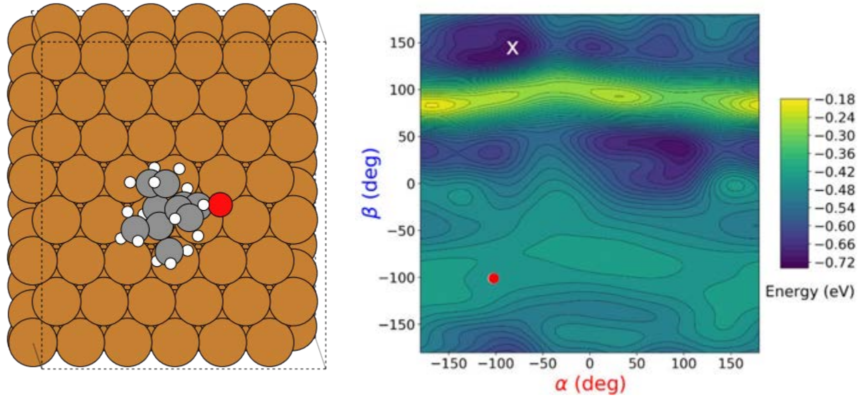

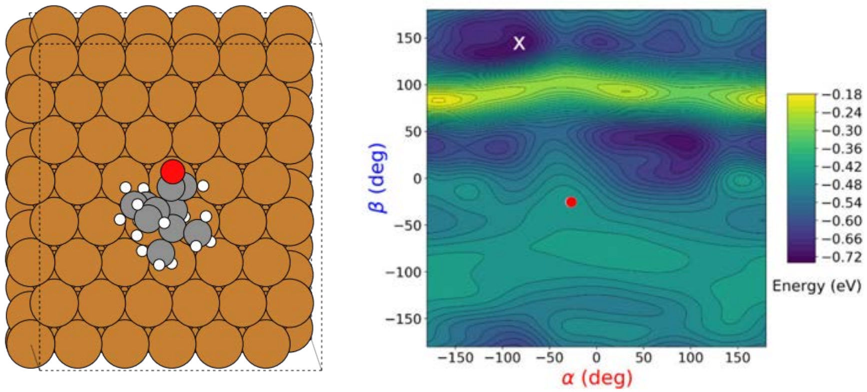

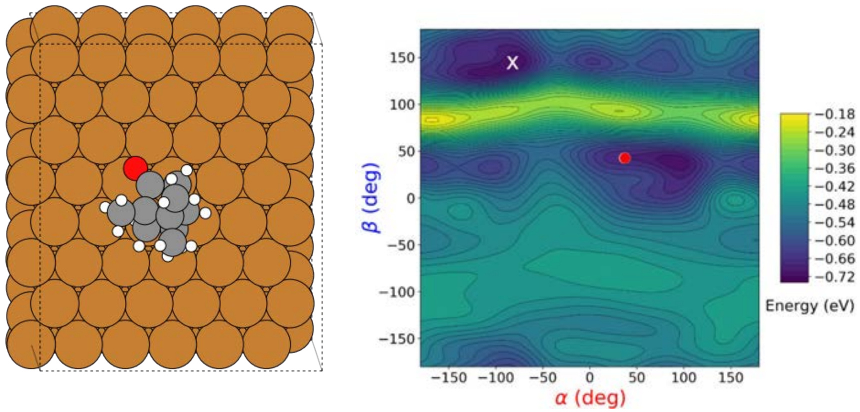

Materials science: molecule on a surface

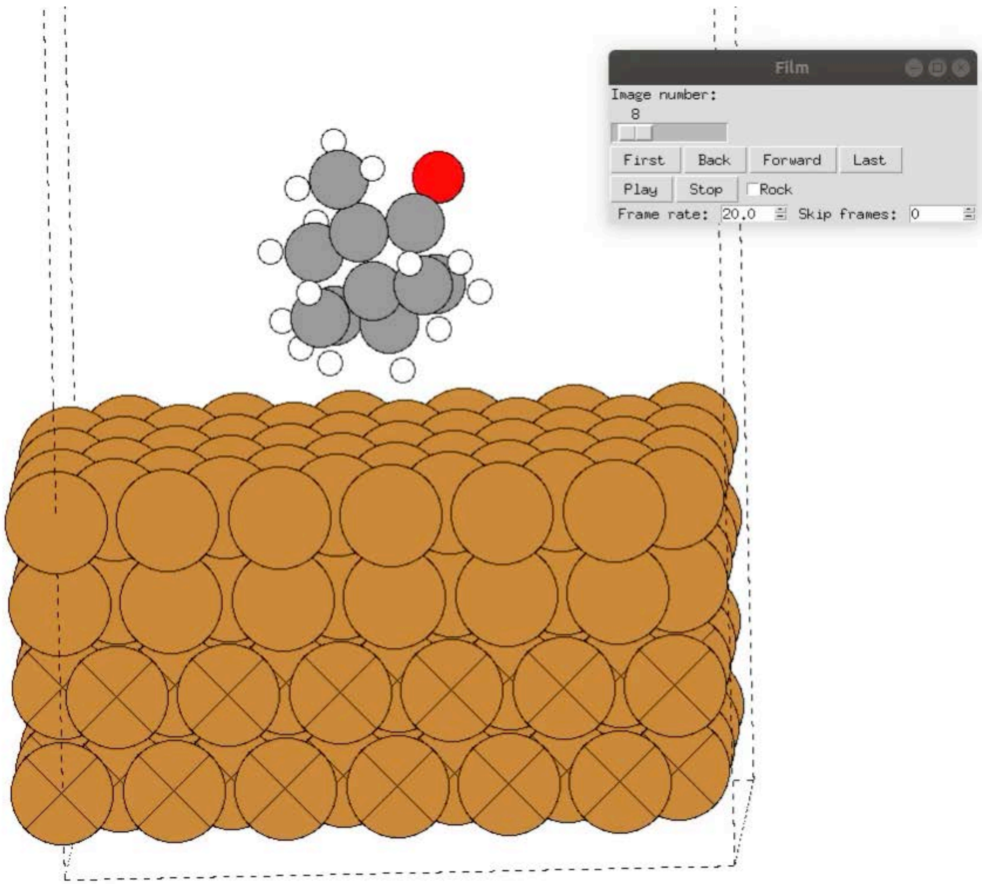

Projective Preferential Bayesian Optimization

Bayesian optimization (BO): Optimization of a black-box function.

Query $f(x)$ at x, to find max $f(x)$

Preferential BO: Query "$f(x) < f(y)$" ?

Projective preferential BO: Give max $f(x)$ on a 1D subspace, by adjusting a slider

arXiv:2002.03113

Projective Preferential Bayesian Optimization

Petrus Mikkola, Milica Todorović, Jari Järvi, Patrick Rinke, Samuel Kaski

Contents

1. User as a data source
2. User as an active agent
3. User as a decision maker
4. Interactive analysis and design

Task: help AI find max *f*(*x*)

You see *f* but the AI only sees what you tell it

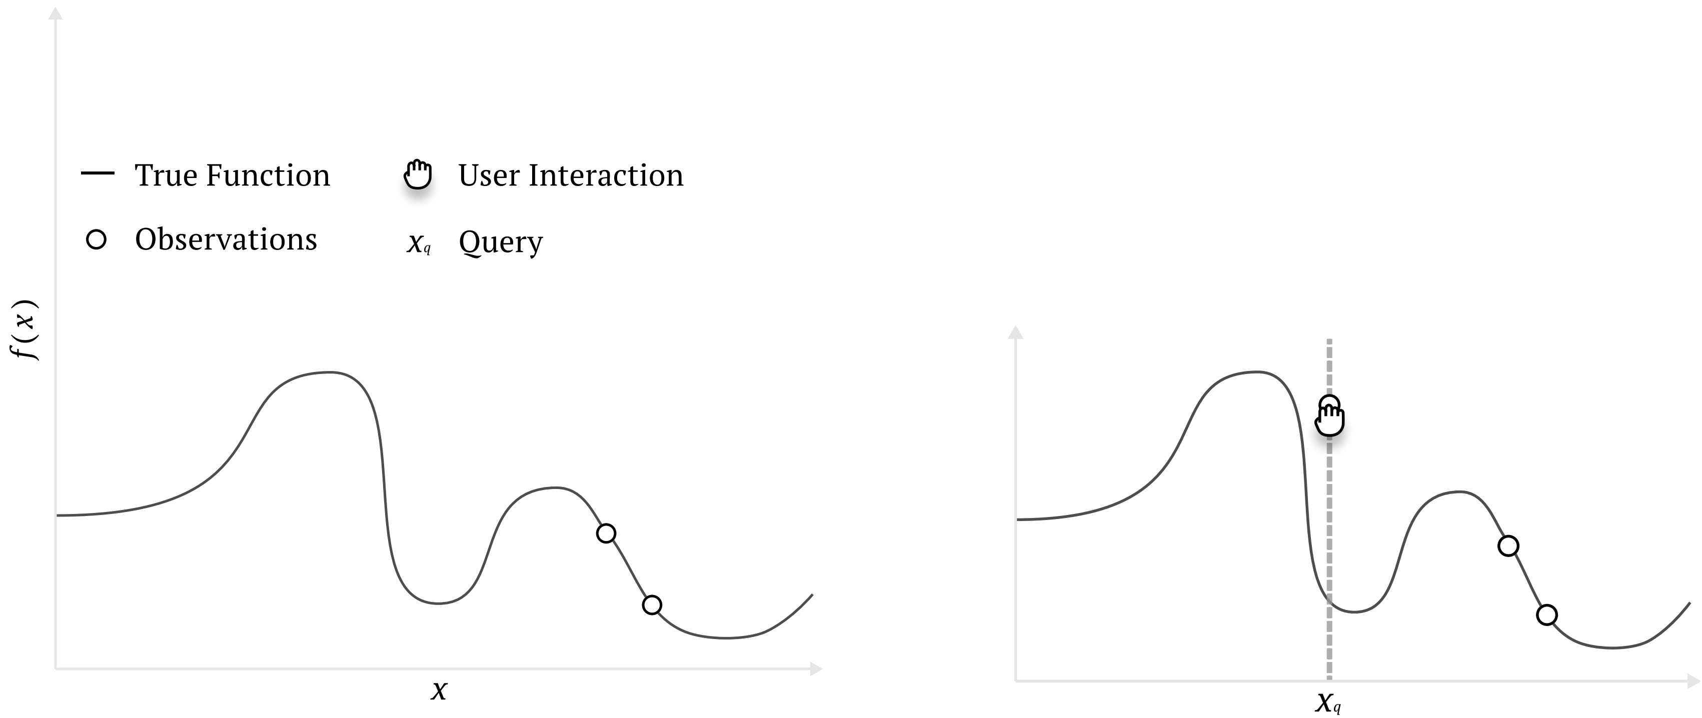

Help AI find max *f*(*x*)

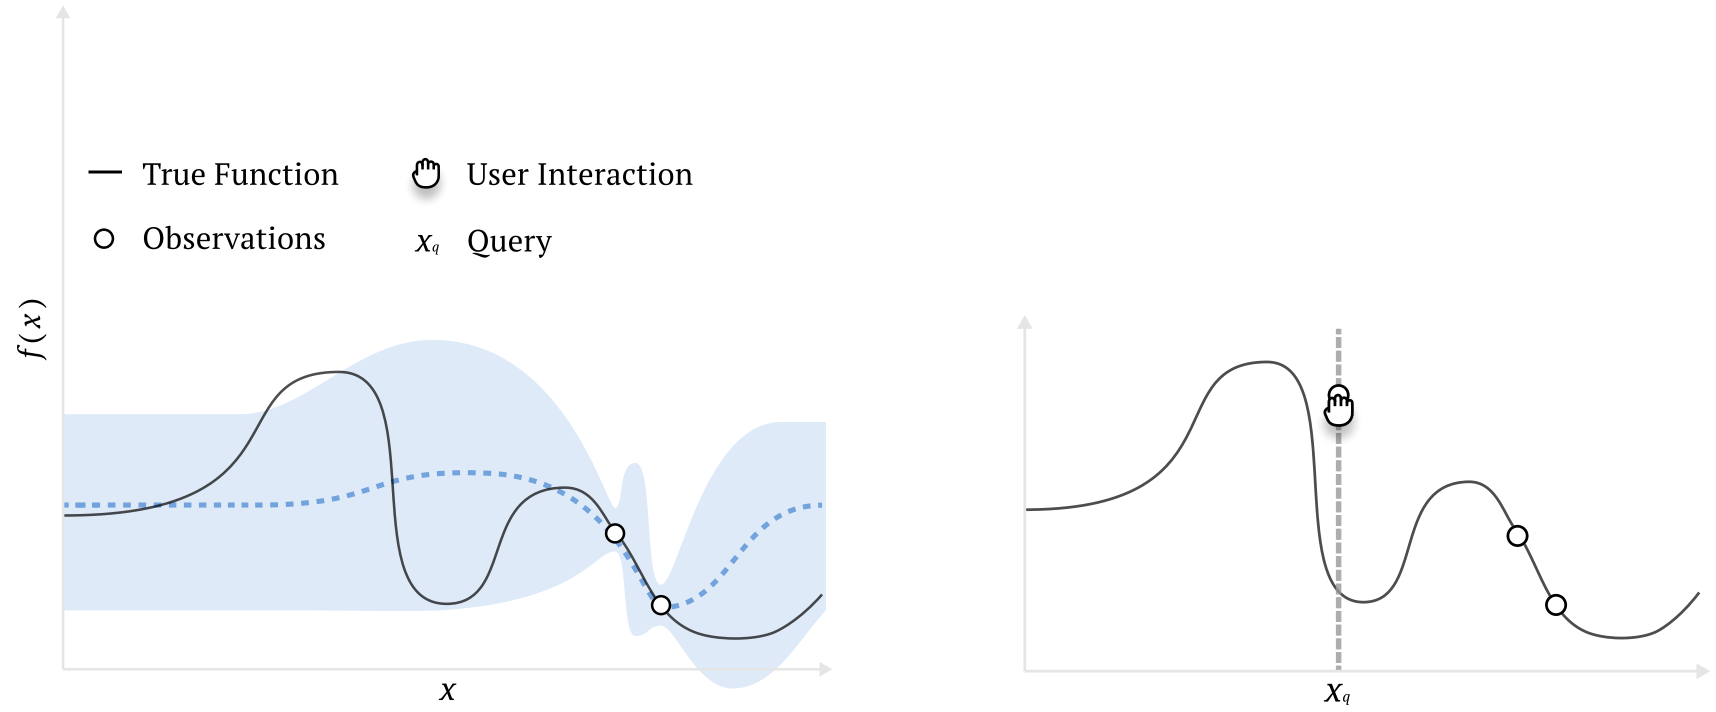

Help AI find max *f*(*x*)

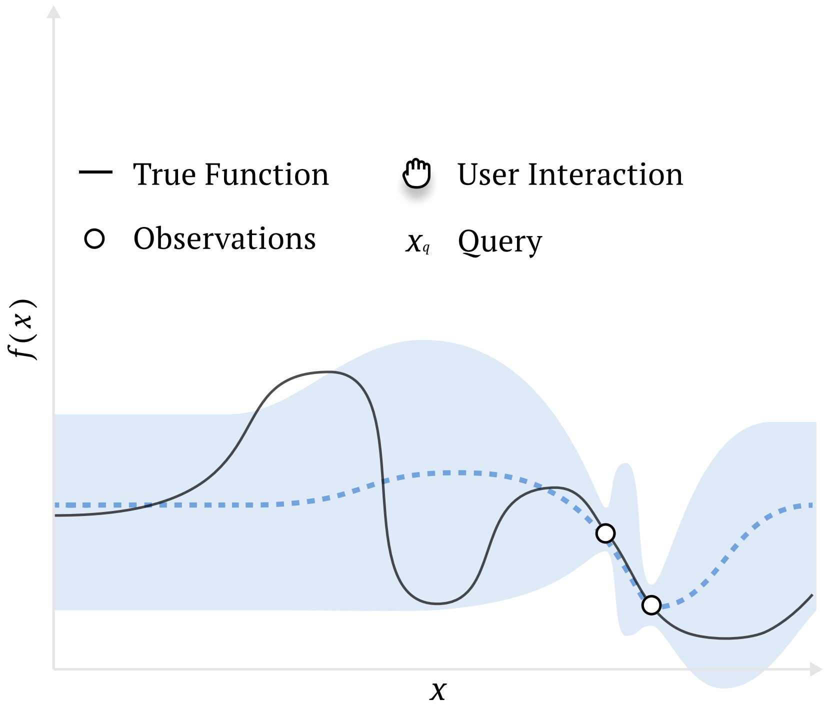

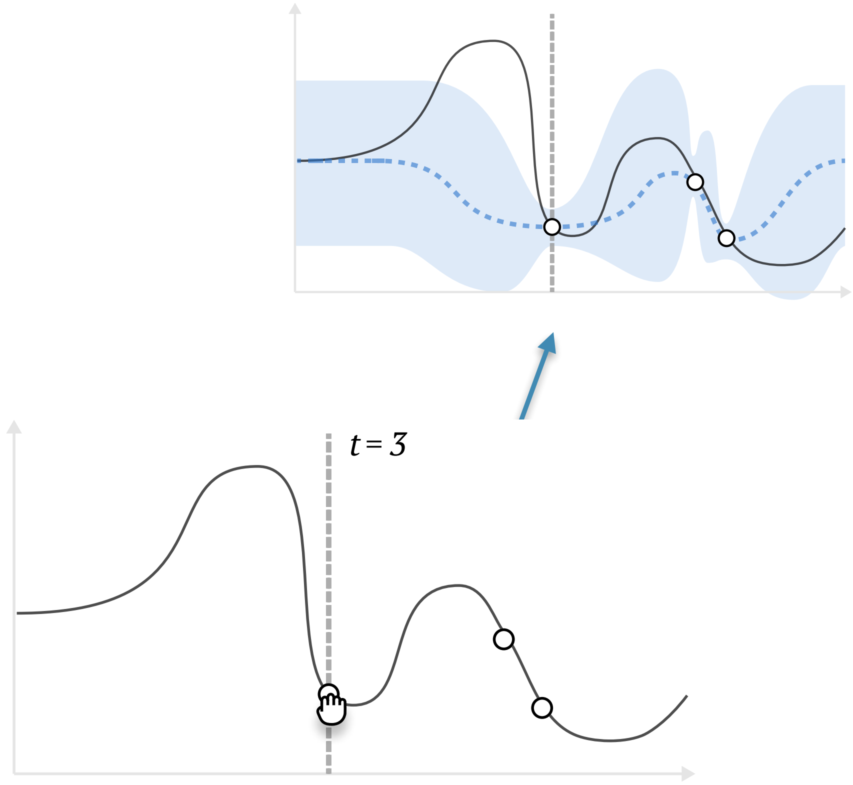

Help AI find max *f*(*x*)

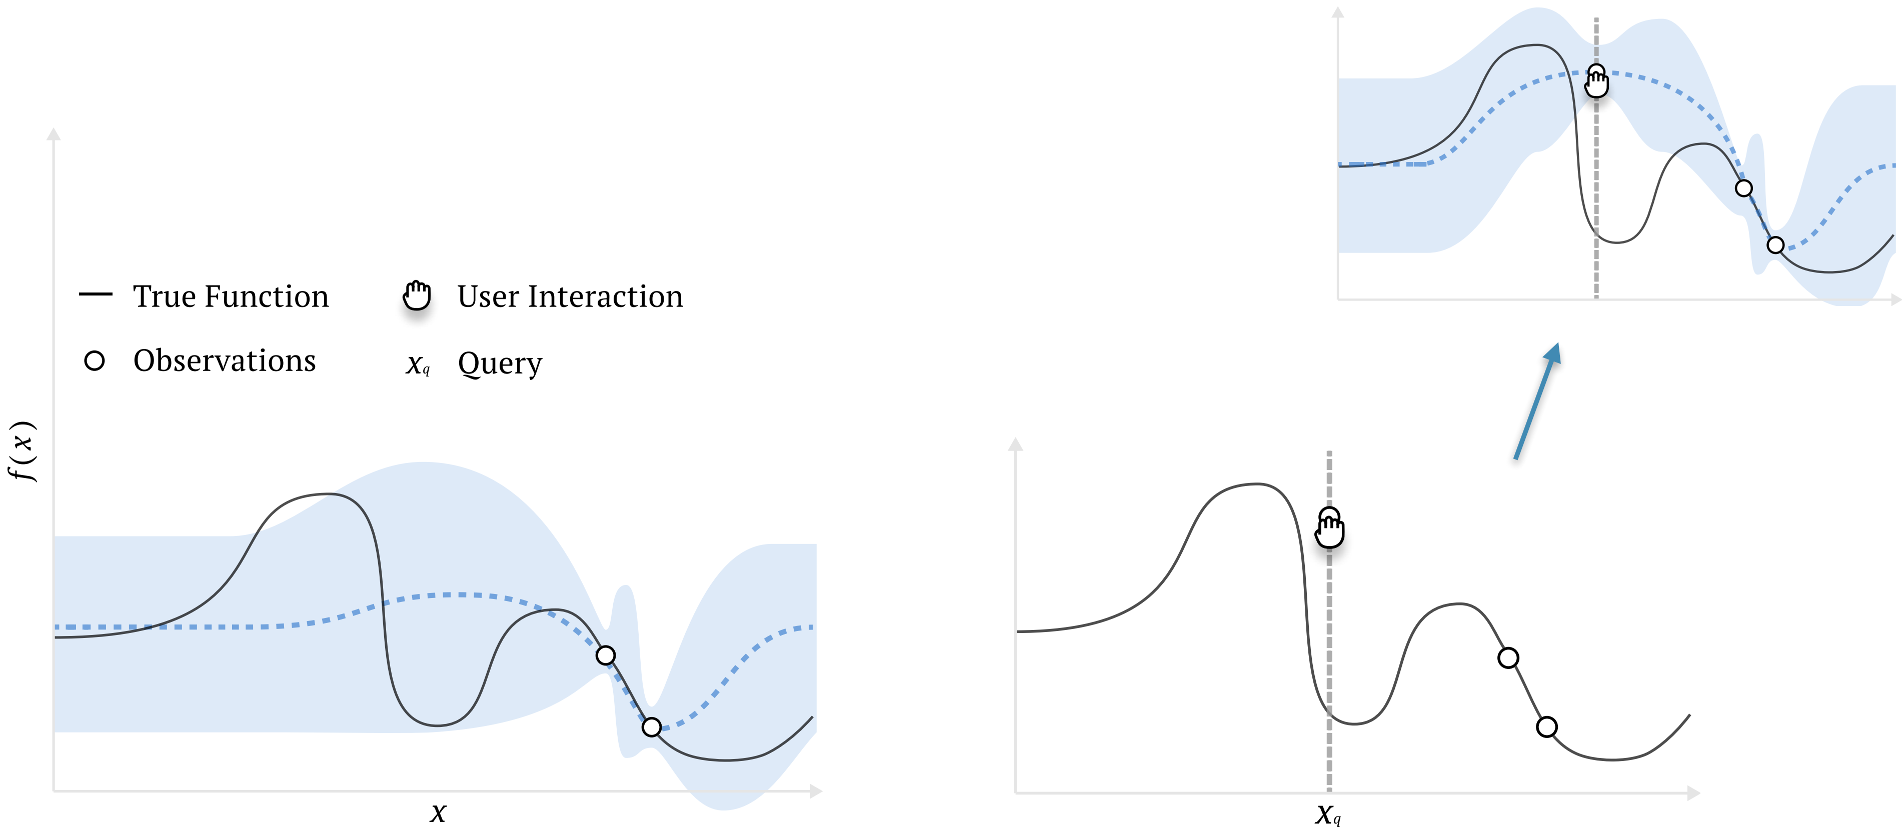

Results: People do not just passively give *f*(*x*)

… and that improves performance!

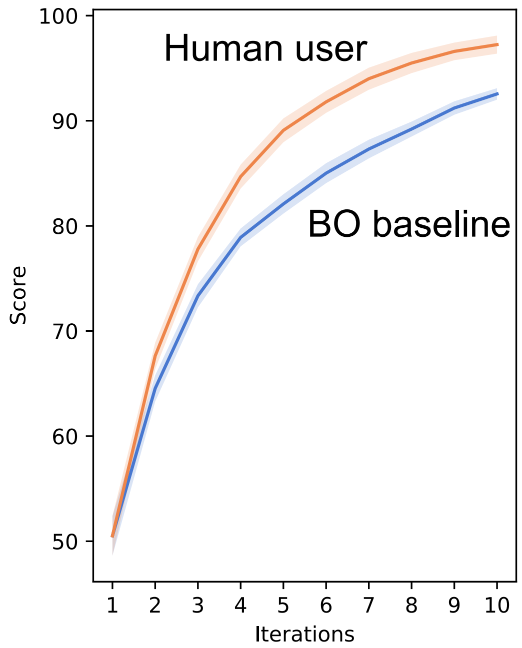

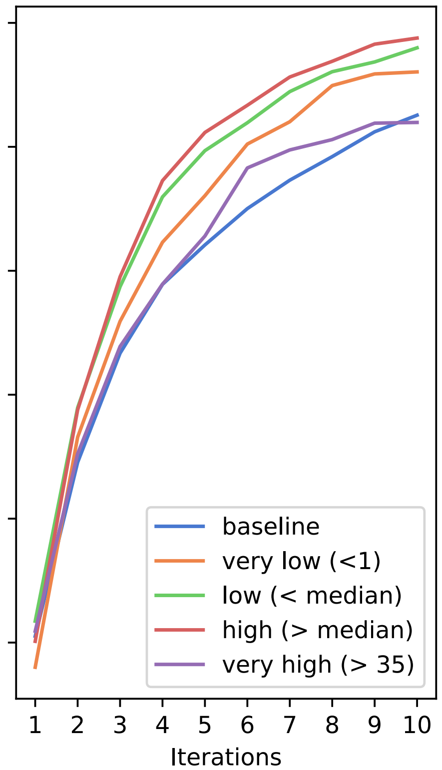

Difference of *f*(*x*) and user's answer

FCAI

Human Strategic Steering Improves Performance of Interactive Optimization

arXiv:2005.01291

Fabio Colella, Pedram Daee, Jussi Jokinen, Antti Oulasvirta, Samuel Kaski

UMAP 2020

AI's assumptions of the user: Levels of user models

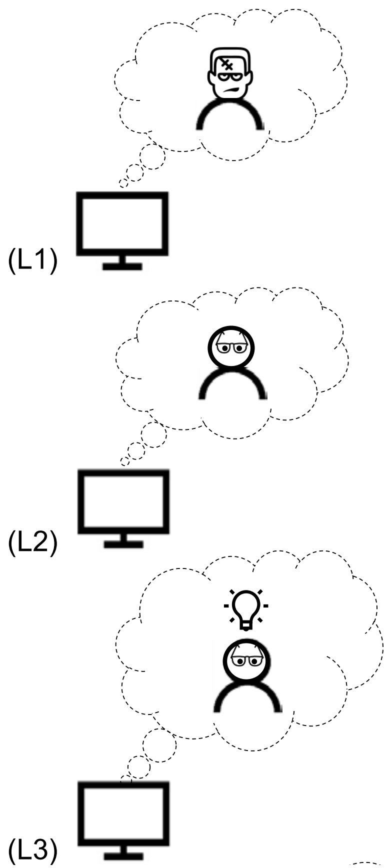

Prescriptive and fixed model of the user

Passive user profile

Active agent

arXiv:1912.05284

Interactive AI with a Theory of Mind

Mustafa Mert Çelikok, Tomi Peltola, Pedram Daee, Samuel Kaski

Cognitive model of the user

Model of the user $M^u(\theta^u)$

$\downarrow$

Data $D^u = \{x^u\}$

Simulator-type models: Can simulate data given parameters

Question (inverse modelling): How to infer the parameters given data?

Active user: sequential decision maker with limitations

Model for sequential decision-making: (Partially Observable) **Markov Decision Process** (PO)MDP

Bounded rationality: With limitations

Forward solution: **Reinforcement learning** (RL)

Inverse modelling: **Inverse Reinforcement Learning** (IRL)
- infer parameters given behavior data: goals, limitations

Inferring cognitive models from data

Our solution:
Approximate Bayesian
Computation (ABC)

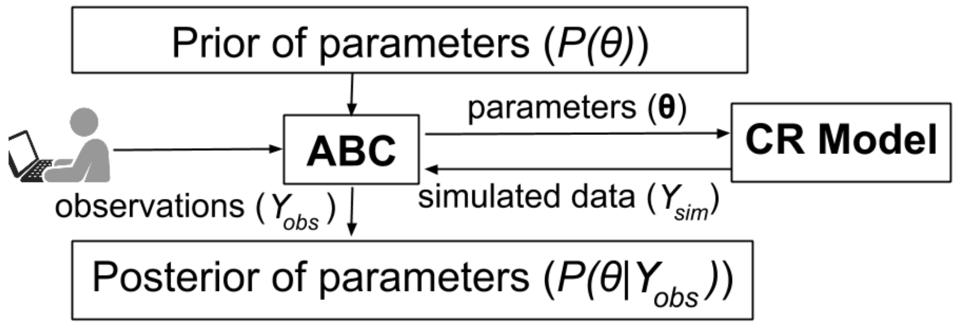

What is special in ABC:
- very heavy simulator for MDP-based models: reinforcement learning in the inner loop
- applicable to also other cognitive models

What is special in IRL:
- only coarse summary-level data available
- infer not only reward function but user parameters

Learning cognitive models with ABC

Parameter Inference for Computational Cognitive Models With Approximate Bayesian Computation

Antti Kangasrääsiö,[a] Jussi P. P. Jokinen,[b] Antti Oulasvirta,[b] Andrew Howes,[c] Samuel Kaski[a]

COGNITIVE SCIENCE
A Multidisciplinary Journal
2019

Example task: Menu search
CHI 2017

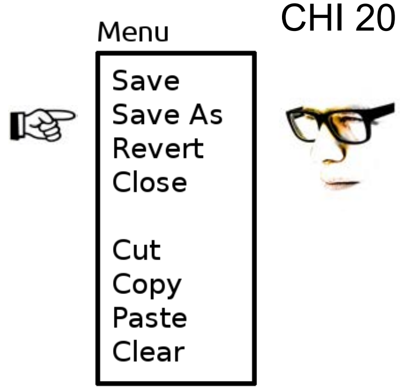

Inverse reinforcement learning from summary data

Antti Kangasrääsiö[1] · Samuel Kaski[1]

Machine Learning (2018)

ELFI: Engine for Likelihood-Free Inference

JMLR

elfi.ai

Jarno Lintusaari, Henri Vuollekoski, Antti Kangasrääsiö, Kusti Skytén, Marko Järvenpää, Pekka Marttinen, Michael U. Gutmann, Aki Vehtari, Jukka Corander, Samuel Kaski; 19(16):1–7, 2018.

AI’s assumptions of the user: Levels of user models

(L1) 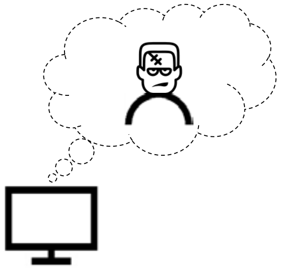

Prescriptive and fixed
model of the user

(L2) 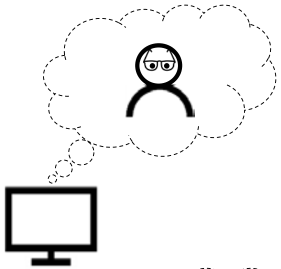

Passive user profile

(L3) 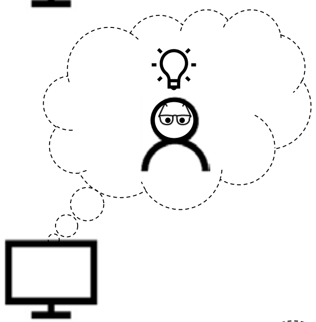

Active agent

(L4) 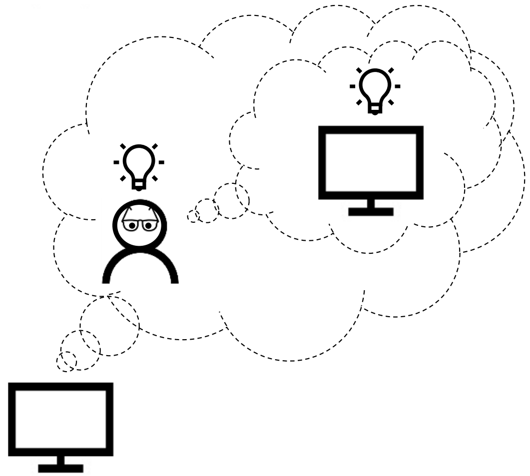

Active agent that has
a model of the AI system

arXiv:1912.05284

Interactive AI with a Theory of Mind

Mustafa Mert Çelikok, Tomi Peltola, Pedram Daee, Samuel Kaski

Computational Modeling in HCI: ACM CHI 2019 Workshop,

AI to understand its user

Ever tried to get an AI system to actually do stuff? That would work much better if it understood that you do have a goal.

- We give the AI learner the capability to understand that the user has a goal and a plan, by learning a user model
- Further, we give an AI teacher the ability to teach an active learner

Case study: information retrieval

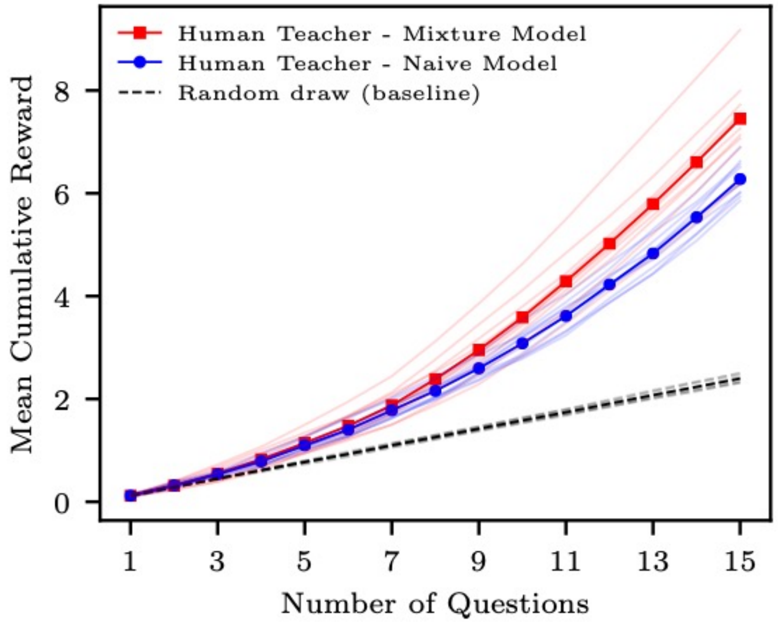

Peltola et al., NeurIPS 2019

Machine teaching of active sequential learners

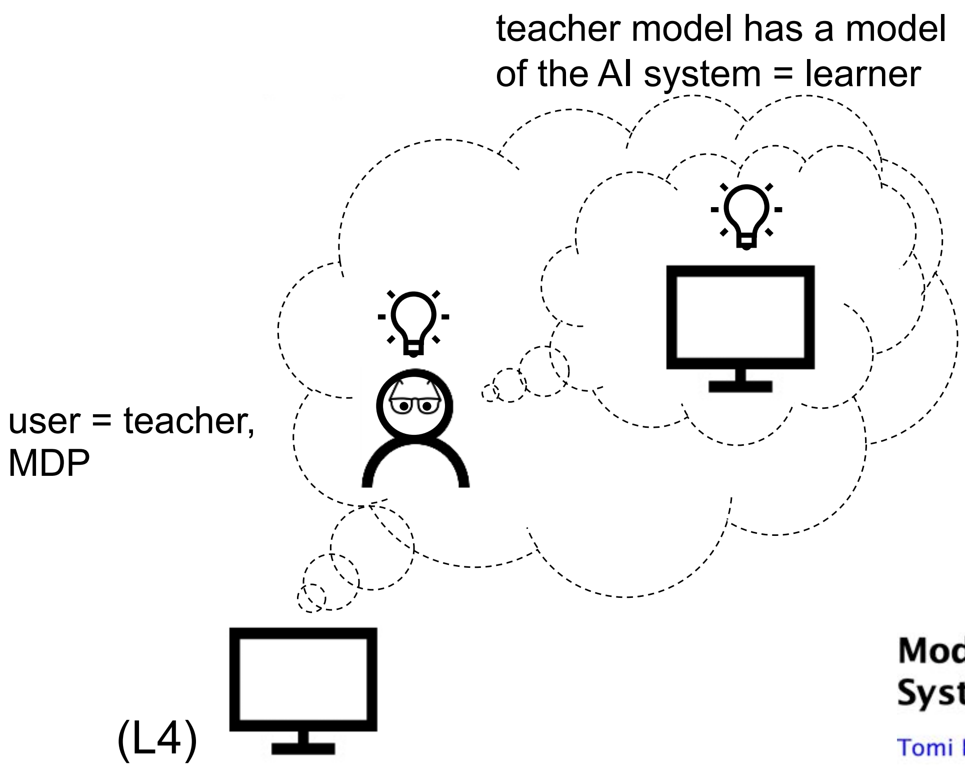

Machine Teaching of Active Sequential Learners

Tomi Peltola
tomi.peltola@aalto.fi

Mustafa Mert Çelikok
mustafa.celikok@aalto.fi

Pedram Daee
pedram.daee@aalto.fi

Samuel Kaski
samuel.kaski@aalto.fi

NeurIPS 2019

Modelling User's Theory of AI's Mind in Interactive Intelligent Systems

Tomi Peltola, Mustafa Mert Çelikok, Pedram Daee, Samuel Kaski

arXiv:1809.02869v2 2018

https://aaltopml.github.io/machine-teaching-of-active-sequential-learners/

Contents

Active learning for decision-making

Task: Decision support for personalized medicine. Improve predictions at $\tilde{x}$

Solution: Include a model of user's decision making task as the active learning criterion

Proposed additionally a new query type: counterfactual elicitation

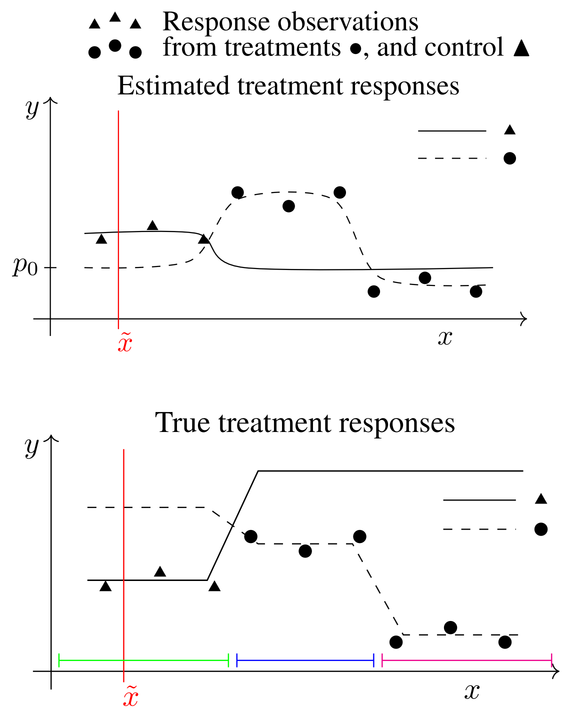

ICML 2019

Active Learning for Decision-Making from Imbalanced Observational Data

Iiris Sundin[1] Peter Schulam[*2] Eero Siivola[*1] Aki Vehtari[1] Suchi Saria[2] Samuel Kaski[1]

Bayesian models that are both understandable and predictive

How to make models both predictive and interpretable?

Widespread approach: Use only simple models

- Bayesian version: prior which favours simple models

Our proposal: Fit an accurate (reference) model, and include interpretability into the utility to be maximized given the model.

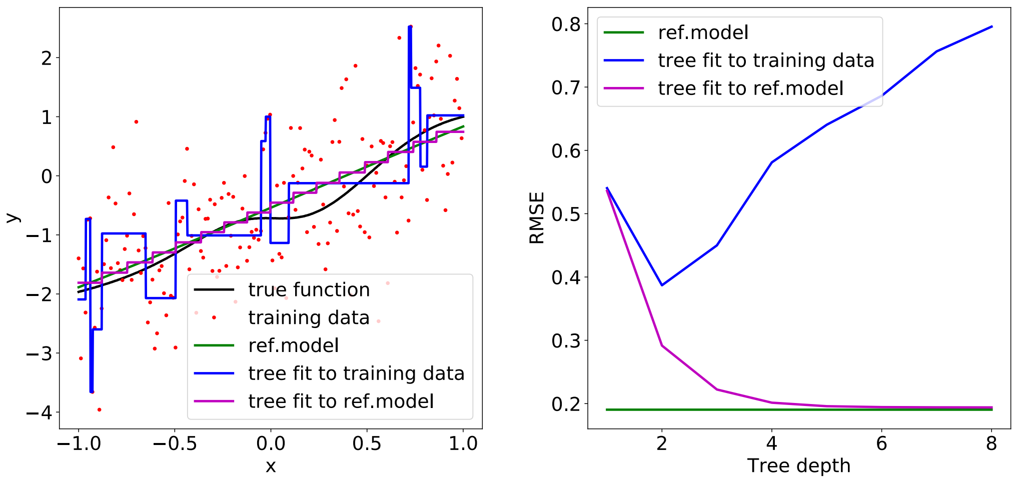

Example: Decision tree for interpretability

Making Bayesian Predictive Models Interpretable: A Decision Theoretic Approach

Homayun Afrabandpey, Tomi Peltola, Juho Piironen, Aki Vehtari, Samuel Kaski

arXiv:1910.09358

Contents

AI-assisted design

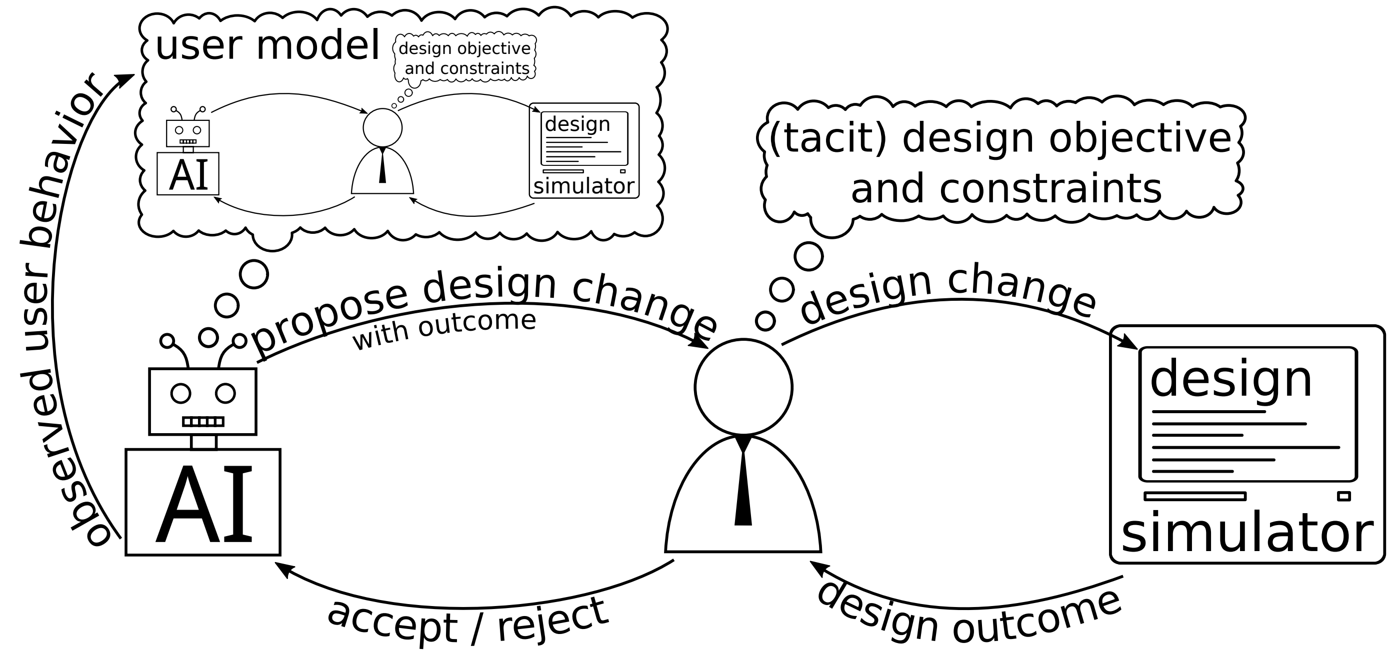

Image: Thanks to Sebastiaan De Peuter

Contents

1. User as a data source
2. User as an active agent
3. User as a decision maker
4. Interactive analysis and design

Collaborators

Aalto Univ: Antti Oulasvirta; Patrick Rinke; Jussi Jokinen, Jari Järvi, Milica Todorovic

Univ Birmingham: Andrew Howes

Univ Helsinki: Giulio Jacucci, Caroline Hecman, Aki Havulinna, Dorota Glowacka, Chen He, Muntasir Mamun Majumder, Tuukka Ruotsalo, Baris Serim

Johns Hopkins Univ: Suchi Saria, Peter Schulam

Univ Oslo: Jukka Corander

More information: fcai.fi , research.cs.aalto.fi/pml