Department of Computer Science

Interactive Knowledge Elicitation for Decision-Support Models in Precision Medicine

Iiris Sundin

Aalto University

DOCTORAL THESES

Aalto University publication series
DOCTORAL THESES 78/2023

Interactive Knowledge Elicitation for Decision-Support Models in Precision Medicine

Iiris Sundin

A doctoral thesis completed for the degree of Doctor of Science (Technology) to be defended, with the permission of the Aalto University School of Science, at a public examination held at the lecture hall AS1 of the school on 20 June 2023 at 12:00.

Aalto University
School of Science
Department of Computer Science
Probabilistic Machine Learning

Supervising professor
Professor Samuel Kaski, Aalto University, Department of Computer Science, Finland

Thesis advisor
Professor Samuel Kaski, Aalto University, Department of Computer Science, Finland

Preliminary examiners
Associate Professor Tapio Pahikkala, University of Turku, Finland
Associate Professor Laura Ruotsalainen, University of Helsinki, Finland

Opponent
Doctor Danielle Belgrave, DeepMind, UK

Aalto University publication series
DOCTORAL THESES 78/2023

ISBN 978-952-64-1284-9 (printed)
ISBN 978-952-64-1285-6 (pdf)
ISSN 1799-4934 (printed)
ISSN 1799-4942 (pdf)
http://urn.fi/URN:ISBN:978-952-64-1285-6

Images: Cover Image: Iiris Sundin

Unigrafia Oy
Helsinki 2023

Finland

NORDIC SWAN ECOLABEL
Printed matter
4041-0619

Author
Iiris Sundin

Name of the doctoral thesis
Interactive Knowledge Elicitation for Decision-Support Models in Precision Medicine

Publisher School of Science

Unit Department of Computer Science

Series Aalto University publication series DOCTORAL THESES 78/2023

Field of research Machine Learning

Manuscript submitted 20 August 2022 **Date of the defence** 20 June 2023

Permission for public defence granted (date) 15 November 2022 **Language** English

☐ **Monograph** ☒ **Article thesis** ☐ **Essay thesis**

Abstract

This thesis develops human-in-the-loop machine learning methods that aim at improving the performance of a machine learning model in precision medicine tasks. Many problems in precision medicine are still difficult for machine learning due to lack of data, and human experts' knowledge can provide a valuable source of information to reduce a model's prediction error and uncertainty. Such expert knowledge elicitation requires methods that address the following problems: How to leverage indirect expert knowledge instead of querying labels as in active learning, how to make the interaction less laborious to the expert than in traditional prior elicitation, and how to select the interaction so that it is the most beneficial to the prospective task of the model.

The first contribution of the thesis is to develop an interactive knowledge elicitation method for "small n large p" problems where data is insufficient, that allows even a small amount of sequentially chosen noisy, indirect feedback from an expert to complement the data and improve the accuracy of the model's predictions. The effectiveness of the method is evaluated in a user-study. The method is further extended to a high-dimensional genomics prediction task where we demonstrate, for the first time, how sequentially selected domain expert's feedback improves personalized prediction of the cancer cell's sensitivity to drugs.

The second main contribution of the thesis is to introduce two goal-oriented data acquisition strategies that aim at selecting queries that are maximally useful for a prospective task where the model is to be used: First, targeted Bayesian optimal experimental design to increase the accuracy of a single personalized prediction, and second, active learning that takes the down-the-line decision-making task into account by modeling the probability of a wrong decision.

The last part of this thesis applies human-in-the-loop methods to a new, promising and yet unexplored application domain of *de novo* molecular design. The last contribution is how the goal of molecule generation can be inferred via human-in-the-loop interaction, to make an adaptive objective function to a reinforcement learning algorithm, so that the resulting system generates more molecules that match the user's goal.

Keywords interactive machine learning, human-in-the-loop, Bayesian modeling

ISBN (printed) 978-952-64-1284-9 **ISBN (pdf)** 978-952-64-1285-6

ISSN (printed) 1799-4934 **ISSN (pdf)** 1799-4942

Location of publisher Helsinki **Location of printing** Helsinki **Year** 2023

Pages 163 **urn** http://urn.fi/URN:ISBN:978-952-64-1285-6

Tekijä
Iiris Sundin

Väitöskirjan nimi
Vuorovaikutteisia koneoppimismenetelmiä asiantuntijatiedon tuomiseen laskennallisiin malleihin täsmälääketieteen päätöksenteon tueksi

Julkaisija Perustieteiden korkeakoulu

Yksikkö Tietotekniikan laitos

Sarja Aalto University publication series DOCTORAL THESES 78/2023

Tutkimusala Koneoppiminen

Käsikirjoituksen pvm 20.08.2022 **Väitöspäivä** 20.06.2023

Väittelyluvan myöntämispäivä 15.11.2022 **Kieli** Englanti

☐ **Monografia** ☒ **Artikkeliväitöskirja** ☐ **Esseeväitöskirja**

Tiivistelmä

Tässä väitöskirjassa kehitetään vuorovaikutteisia koneoppimismenetelmiä, joilla pyritään parantamaan koneoppimismallin suorituskykyä täsmälääketieteen tehtävissä. Monet täsmälääketieteen ongelmat ovat edelleen vaikeita koneoppimiselle datan puutteen vuoksi, mutta asiantuntijoiden näkemys tarjoaa arvokkaan tietolähteen, jolla voidaan vähentää mallien ennustevirhettä ja epävarmuutta. Tällaisen asiantuntijatiedon hankkiminen edellyttää menetelmiä, jotka ratkaisevat seuraavat ongelmat: Miten hyödyntää myös epäsuoraa asiantuntijatietoa sen sijaan, että ihminen paljastaa mallille oikeat vastaukset kuten aktiivioppimisessa, miten tehdä vuorovaikutuksesta asiantuntijalle vähemmän työlästä kuin perinteinen priorijakaumien määrittely, ja miten valita asiantuntijalle esitettävät kysymykset siten, että niistä on eniten hyötyä mallin tulevan tehtävän kannalta.

Väitöskirjassa kehitetään interaktiivinen menetelmä asiantuntijatiedon tuomiseen malliin, jolla voidaan ratkaista ns. "pieni n suuri p" -ongelmia, joissa dataa ei ole riittävästi. Menetelmän avulla pienikin määrä aktiivisesti valittua asiantuntijan antamaa palautetta parantaa mallin ennusteiden tarkkuutta kun se yhdistetään data-analyysiin. Menetelmän tehokkuus osoitetaan käyttäjätutkimuksessa. Tämä menetelmä laajennetaan myös korkeaulotteiseen genomiikan ennustustehtävään ja näytetään ensimmäistä kertaa, että asiantuntijan palaute parantaa yksilöllisiä ennusteita syöpäsolujen lääkeherkkyydestä.

Väitöskirjan toinen keskeinen tulos on kahden tavoitteellisen aktiivioppimisstrategian kehittäminen, joilla pyritään valitsemaan asiantuntijalle esitettävät kysymykset niin, että ne ovat mahdollisimman hyödyllisiä tehtävässä, jossa koneoppimismallia on tarkoitus käyttää. Ensimmäinen menetelmä on kohdennettu Bayesilainen optimaalinen koesuunnittelu, jolla pyritään lisäämään yksittäisen yksilöllistetyn ennusteen tarkkuutta. Toinen on päätöksentekoon keskittyvä aktiivioppimismenetelmä, joka ottaa huomioon edessä olevan päätöksentekotehtävän mallintamalla väärän päätöksen todennäköisyyden ja minimoimalla sitä.

Väitöskirjan viimeisessä osassa sovelletaan vuorovaikutteisia koneoppimismenetelmiä uuteen, lupaavaan ja vielä tutkimattomaan sovellusalueeseen, *de novo* -molekyylisuunnitteluun. Tuloksena on menetelmä, jossa algoritmi vuorovaikuttaa kemistin kanssa ja päättelee tämän palautteen perusteella molekyylisuunnittelutehtävän tavoitteen. Tällöin vahvistusoppimisalgoritmille voidaan luoda adaptiivinen tavoitefunktio niin, että järjestelmä tuottaa enemmän kemistin tavoitetta vastaavia molekyylejä.

Avainsanat interaktiivinen koneoppiminen, ihminen silmukassa, Bayesilainen mallinnus

ISBN (painettu) 978-952-64-1284-9 **ISBN (pdf)** 978-952-64-1285-6

ISSN (painettu) 1799-4934 **ISSN (pdf)** 1799-4942

Julkaisupaikka Helsinki **Painopaikka** Helsinki **Vuosi** 2023

Sivumäärä 163 **urn** http://urn.fi/URN:ISBN:978-952-64-1285-6

Preface

This thesis has been an exciting journey, and there are many who have travelled with me. Some have walked with me from the very beginning of my doctoral studies, or even earlier, and with some our paths have joined only for a short while. In any case, I've enjoyed working with these brilliant people, and I want to thank them all for our shared work and time!

The research work presented in this thesis has been carried out at the Department of Computer Science of Aalto University during years 2015 – 2022. I thank Prof. Samuel Kaski for taking me in to the Probabilistic Machine Learning (PML) group and giving me the most exciting projects to work on. Thank you for teaching me how to think in machine learning ways, respecting me and my ideas, and for constructive feedback and challenging questions. They taught me what academic discussion and research are at their best: sharp and challenging arguments accompanied with mutual respect and a common goal. I also wish to express my gratitude to the preliminary examiners of this thesis, Prof. Laura Ruotsalainen and Prof. Tapio Pahikkala, and to Dr. Danielle Belgrave who has agreed to act as an opponent in the upcoming public defence.

I have had a chance to work with many professors during my studies and learn so much from them. Prof. Aki Vehtari taught me the fundamentals of Bayesian data analysis, and Prof. Ola Engkvist introduced me to the exciting world of *de novo* molecular design. I also enjoyed working with Prof. Giulio Jacucci, who knows how to create an atmosphere of happiness and scientific inspiration in a research team. I also wish to thank Prof. Suchi Saria for our collaboration, where she taught me how to hone research ideas into effective hypotheses.

Prof. Pekka Marttinen, Dr. Marta Soare, Prof. Luana Micallef and Dr. Tomi Peltola were some of the first researchers I encountered when I started my probabilistic machine learning journey. Thanks Pekka for being a great colleague throughout these years! I'm grateful to Marta for her support as a post-doc to me, a new Ph.D. student. My warm thanks also to Tomi for the years we worked together and all the help I got from him. Luana was a dear colleague to me, and I was heartbroken to learn that she passed away a few years ago. I miss her. She was a great colleague, empathetic and helpful, and a brilliant researcher.

I'm grateful to Dr. Homayun Afrabandpey, Dr. Pedram Daee, Dr. Muhammad Ammad-ud-din, Dr. Eero Siivola, Petrus Mikkola and many others, who have shared the Ph.D. student journey with me and made the work in the PML group enjoyable. My warm thanks also to Dr. Louis Filstroff, who has been a colleague and a friend during the challenging times of the pandemic. Thank you Louis and Petrus also for proofreading parts of this thesis. I wish to thank my co-authors Dr. Baris Serim, Dr. Mamun Majumder, Dr. Caroline Heckman, Dr. Chen He, Prof. Aki Havulinna, Dr. Peter Schulam, Juuso Kylmäoja, Prof. Aleksei Tiulpin, Dr. Alexey Voronov, Haoping Xiao, Dr. Kostas Papadopoulos, Dr. Esben Jannik Bjerrum, Dr. Markus Heinonen and Dr. Atanas Patronov for the fruitful collaboration and their invaluable insights on the research problems we studied.

I have been fortunate to be a part of the the Academy of Finland Flagship Finnish Center for Artificial Intelligence (FCAI), as well as the Finnish Centre of Excellence in Computational Inference Research (COIN). During these years, I have also received personal grants to fund my research from Jenny and Antti Wihuri foundation, Alfred Kordelin foundation and Emil Aaltonen foundation. I express my deep gratitude for their trust and support to this work. I also thank the computational resources provided by the Aalto Science-IT Project, without which much of this work would have been difficult to execute.

I thank the current and former members of the PML group for scientific and non-scientific discussions and fun coffee break activities, such as Go and plank challenges. In addition, I wish to thank our (apparently very successful) peer support Ph.D. student group: Dr. Antti Silver, Dr. Hannu Leppinen, Dr. Eero Laukkanen and Dr. Olli Kilkki.

My friends and family have been invaluable support to me during these years. I give my warmest thanks and love to my parents Hanna and Tuomas Routa, who have always believed in me. I also thank Mea Sinersaari for being like a sister to me and for all the discussions and emotional support. Last, my deepest thanks to my beloved husband Juho, who has supported me in so many ways and with great effort, and to our dear children Jooa and Varpu. You three make my life perfect.

Espoo, May 4, 2023,

Iiris Sundin

Contents

List of Publications

This thesis consists of an overview and of the following publications which are referred to in the text by their Roman numerals.

I Luana Micallef, Iiris Sundin, Pekka Marttinen, Muhammad Ammad-ud-din, Tomi Peltola, Marta Soare, Giulio Jacucci, and Samuel Kaski. Interactive elicitation of knowledge on feature relevance improves predictions in small data sets. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces (IUI '17)*, Limassol, Cyprus, pp. 547–552, March 2017.

II Iiris Sundin, Tomi Peltola, Luana Micallef, Homayun Afrabandpey, Marta Soare, Muntasir Mamun Majumder, Pedram Daee, Chen He, Baris Serim, Aki Havulinna, Caroline Heckman, Giulio Jacucci, Pekka Marttinen, and Samuel Kaski. Improving genomics-based predictions for precision medicine through active elicitation of expert knowledge. *Bioinformatics*, Volume 34, Issue 13, pp. i395–i403, July 2018.

III Iiris Sundin, Peter Schulam, Eero Siivola, Aki Vehtari, Suchi Saria, and Samuel Kaski. Active Learning for Decision-Making from Imbalanced Observational Data. In *Proceedings of the 36th International Conference on Machine Learning*, Long Beach, United States, PMLR 97, pp. 6046-6055, June 2019.

IV Louis Filstroff, Iiris Sundin, Petrus Mikkola, Aleksei Tiulpin, Juuso Kylmäoja, and Samuel Kaski. Targeted Active Learning for Bayesian Decision-Making. Submitted to *a journal*, 24 pages, February 2022.

V Iiris Sundin, Alexey Voronov, Haoping Xiao, Kostas Papadopoulos, Esben Jannik Bjerrum, Markus Heinonen, Atanas Patronov, Samuel Kaski, Ola Engkvist. Human-in-the-Loop Assisted de Novo Molecular Design. *Journal of Cheminformatics*, 14, 86, December 2022.

Author's Contributions

Publication I: "Interactive elicitation of knowledge on feature relevance improves predictions in small data sets"

The first three authors, Micallef, Sundin and Marttinen contributed equally. Micallef designed and prepared the user-study and the graphical user interface. Marttinen and Kaski developed the original idea. Sundin participated in the model and algorithm development together with Marttinen and Soare. Sundin and Micallef conducted the user study, and Sundin analyzed the results. Ammad-ud-din, Peltola and Sundin implemented the method (about 30 per cent contribution). Kaski and Jacucci supervised the research. The manuscript was written jointly.

Publication II: "Improving genomics-based predictions for precision medicine through active elicitation of expert knowledge"

The main authors Sundin and Peltola developed the original idea from Kaski and invented the targeted elicitation method together with Marttinen. Sundin designed, implemented and ran the preliminary experiments on the drug sensitivity prediction with a human expert feedback, with help from Soare and Micallef. Peltola implemented and analyzed the metabolite prediction experiment, and wrote the first version of the article after which all authors joined in revisions. Afrabandpey, Daee and He designed and implemented the simulated user experiment and analyzed the results. Majumder, Havulinna, and Heckman provided the experimental data.

Publication III: "Active Learning for Decision-Making from Imbalanced Observational Data"

Sundin is the main author of the article. Schulam and Saria helped Sundin to develop the original idea from Kaski, and Vehtari suggested using Type S error for analysis. Sundin derived the proofs with the help of Kaski, designed the experiments, ran them with support from Siivola, and analyzed the results. Sundin implemented about 60 per cent of the program code for the experiments, and Siivola and Schulam the rest. Vehtari, Saria and Kaski supervised the research. Sundin wrote the first version of the manuscript and all authors contributed in revisions.

Publication IV: "Targeted Active Learning for Bayesian Decision-Making"

Filstroff was the main author of the article. The original idea was Sundin and Kaski's, and Filstroff developed it further together with Kaski, Kylmäoja and Mikkola and Sundin (about 20 per cent contribution). Sundin implemented and ran preliminary experiments and participated in writing (about 10 per cent contribution). Filstroff implemented the method, ran the experiments and analyzed the results, Tiulpin designed the knee osteoarthritis diagnosis experiment and prepared the data.

Publication V: "Human-in-the-Loop Assisted de Novo Molecular Design"

Sundin performed the research together with Voronov and Xiao. Engkvist, Kaski and Heinonen proposed the original idea and supervised the project. Sundin developed and implemented the methods, and designed the experimental setup together with all co-authors. Voronov designed and implemented the graphical user-interface. Xiao implemented the chemist's component experiments. Papadopoulos prepared the data. Sundin and Xiao ran the experiments and analyzed the results. Sundin wrote the manuscript with help from co-authors, and all authors revised and approved the manuscript.

Abbreviations

AI	Artificial Intelligence
AL	Active Learning
BALD	Bayesian Active Learning by Disagreement
D-EIG	Expected Information Gain to Decision-making
DRD2	Dopamine Receptor D_2
DSS	Quantitative Drug Sensitivity Score
EIG	Expected Information Gain
GP	Gaussian Process
HITL	Human-In-The-Loop
IHDP	Infant Health and Development Program
KL-divergence	Kullback–Leibler divergence
LOO-CV	Leave-One-Out Cross-Validation
MCMC	Markov Chain Monte Carlo
MI	Mutual Information
MOCU	Mean Objective Cost of Uncertainty
MPO	Multi-Parameter Optimization
MSE	Mean Squared Error

OAI	Osteoarthritis Initiative
SEM	Standard Error of Mean
T-EIG	Targeted Expected Information Gain
UCB	Upper Confidence Bound
QED	Quantitative Estimate of Druglikeness

Symbols

$\alpha(\cdot)$	Acquisition function
γ	Type S error rate
θ	Parameters of a model
ξ	A variable defining a design of an experiment
$\mathrm{Bernoulli}(\theta)$	Bernoulli distribution with the probability of outcome 1 θ
$\mathcal{D}$	A dataset
$\mathbb{E}[\cdot]$	Expected value of a random variable
$\mathcal{F}$	Set of human feedback
$\mathcal{H}(\cdot)$	Entropy of a random variable
$I(\cdot;\cdot)$	Mutual information of two variables
$\mathbb{I}(C)$	Indicator function about condition C
$k(\cdot,\cdot)$	A kernel function of a Gaussian Process
$\mathrm{N}(\mu,\sigma^2)$	Normal distribution with mean μ and standard deviation σ
n	Sample size
p	Number of parameters in a model
$p(x)$	Probability density function of x
$p(x \mid y)$	Conditional probability density function of x given y

Q	A set of potential queries in interactive knowledge elicitation
$\mathbb{R}$	A set of real numbers
$\mathcal{U}$	Set of unlabeled instances in active learning

1. Introduction

Even though machine learning has become a powerful tool in almost all domains, a challenge remains that models do poorly in cases where there is too little data. So-called "small n large p" problems are one example, where the number of features p is larger than the number of samples n in the data, a case often in precision medicine tasks. Furthermore, the condition is prevalent even in the era of big data, if the point of interest lies outside or at the tail of a training distribution. Bayesian models are a natural choice to partly overcome the problem with lack of data, since they provide uncertainty estimates and can rely on prior information. At best, these probabilistic models know when they fail.

Complementary to these approaches, the key insight in this thesis is to alleviate the inherent lack of informative data points by interacting with a human expert to integrate their expertise into the model. This is one type of human-in-the-loop learning, a field that studies the problem of how machine learning models and humans collaborate. Human-in-the-loop methods have shown potential in solving NP-hard problems in medical applications, where data is scarce [1, 2].

The goal of this thesis is to develop human-in-the-loop methods that enable including a human expert's prior knowledge into the model in a principled way, to improve model performance. Traditionally domain experts' knowledge is acquired via prior elicitation, where they specify the full prior distribution or quantiles of model parameters. This is very laborious and infeasible for large models (thousands of parameters). What lacks, are methods for actively selecting the important aspects of the model to query from a human expert. Active learning is a field that studies

strategies of selecting training examples for humans to label, but this new application field of querying about prior knowledge requires new active query selection methods.

Specifically, this thesis aims to address the following problems in interactive knowledge elicitation in "small n large p" regression tasks: Even though active learning and prior elicitation are known to improve the accuracy of a model, it is unclear how much involving of a human having a limited budget will help, and how to elicit their knowledge interactively. An important aspect is that few interactions should lead to perceivable benefits. This is especially important when we want to minimize the query budget, as is the case when interacting with a domain expert. When the query budget is very limited, selecting queries that aim at improving the general accuracy of the model may not be the most efficient way to improve the performance in a prospective task where the model is to be used. Lastly, one specific regime where expert knowledge is crucially needed is in choosing a personalized treatment to patients based on observational data, such as electronic health records. Due to its imbalanced nature, observational data may increase the probability of making wrong treatment decisions in a population. In this case, new experiments are not an option and, instead, we seek to augment data with expert's domain knowledge.

This work is at the intersection of three fields that study acquiring new information to a model: active learning, where a human labels data, Bayesian experimental design, which aims at maximizing the information gain from an experiment, and prior elicitation, where a domain expert specifies prior distributions to capture their knowledge.

Human-in-the-loop methods developed in this thesis differ from the existing line of work in the following ways. Compared to standard active learning, they have a very small interaction budget, because a feasible number of interactions for a domain expert is tens or hundreds (Publications I, II, III, IV and V). In addition, they are specialized so that if the prospective task of the model is known they take it into account (Publications II, III and IV). Lastly, the experts do not label instances, instead, they can give indirect indications about predictors (Publications I and II). On the other hand, in contrast to traditional prior elicitation where the expert needs to specify the prior distribution of each feature, this work develops

active Bayesian methods to reduce the workload in two ways. The first is to use user-modeling and experimental design methods to select which priors are elicited, and second, to use feedback types that are natural for human communication. The feedback types in this thesis are relevance and effect of a feature (Publications I and II), comparison of potential outcomes (Publication III), and how well a molecule matches to a chemist's goal (Publication V). Inferring priors from the indirect feedback not only reduces the workload in specifying the prior distributions, but also requires less statistical knowledge from the domain expert.

The last part of the thesis considers artificial intelligence (AI) assisted drug design. AI is already transforming drug development in *de novo* molecular design, where deep generative models are used to explore novel candidate molecules for drugs. However, using such systems requires a lot domain expertise as well as technical expertise from the users, limiting their usability. This thesis makes a new opening towards AI assisted drug design, by introducing human-in-the-loop learning to aid the user in specifying the goal for *de novo* molecular design.

At the core of this work is how to make measurable improvement in the task an expert is aiming at solving with the help of machine learning, by letting the system interact with them. Human-in-the-loop learning can also be developed for other objectives, such as for enhancing the interpretability of the model or building trust of the user to the model, but they are not in the scope of this thesis.

Main claims of the thesis

1. Even with indirect and qualitative feedback, that is easy for humans, the proposed expert knowledge elicitation methods can solve the small n large p regularization problem better than standard methods and achieve significant improvement in model performance. (Publications I and II)

2. With the proposed active expert knowledge elicitation methods, it is possible to increase a model's performance efficiently in a few queries to a domain expert, and further improve the query efficiency using goal-oriented active learning methods developed in this thesis that take into account the prospective task of the model. (Publica-

tions I, II, III and IV)

3. Imbalance in observational data increases the probability of wrong decisions from a causal model, but the decision-making accuracy can be improved via counterfactual elicitation if experts are able to make unbiased guesses about potential outcomes of alternative actions. (Publication III)

4. In case the scoring function in *de novo* molecular design task is unknown, the proposed human-in-the-loop method can still solve the design task by introducing a user-model to learn the expert's goal from interaction. By letting an expert give feedback on actively chosen candidate molecules, the probabilistic user-model can infer the expert's goal in relatively few queries to achieve generated molecules that better match the expert's goal. (Publication V)

1.1 Research Questions and Contributions

This thesis studies the following research questions.

- **RQ1:** Can human-in-the-loop approach to knowledge elicitation solve the small n large p regularization problem through eliciting indirect and qualitative expert knowledge, and through selecting the most informative queries to include domain knowledge interactively?

- **RQ2:** How to select the interaction so that it is the most beneficial to a prospective task where the model is to be used: predicting outcome of a new individual, or making a decision of which action to take based on model's prediction of outcome under alternative actions?

- **RQ3:** How to apply human-in-the-loop methods to *de novo* molecular design, to actively infer the goal of an expert from interaction for adapting a *de novo* design tool's outcome to better match the expert's objectives?

Publications I, II and III provide contributions to RQ1. Publication I shows that asking a human to give feedback indirectly on the priors of

predictors indeed improves the model performance in a controlled user-study. Publication II extends the method to a high-dimensional case and demonstrates it in a genomics-based prediction in a medical application with a human expert feedback. Indirect feedback types that are covered in this thesis are relevance of a feature or predictor, and potential correlation of the feature with the predicted quantity (PI,PII). Publication III introduces the idea about counterfactual elicitation, where including expert knowledge about potential outcomes of alternative actions can improve causal modeling. The feedback can be directly on the value of the potential outcome, or indirect by comparing two potential outcomes. Two types of query selection strategies are proposed: a bandit-based user-model that simplifies the problem by modeling a user's knowledge separately from the data generation process, and a joint model of the user and data. The bandit-based user-model learns about the knowledge of the expert, by modeling the features' relevance and solving exploration-exploitation trade-off by selecting queries that are either relevant or have potential to improve the user-model (Publication I). The second method jointly models user's answers and the data, and selects the queries based on their information gain, which accomplishes trading off exploration and exploitation with respect to the full human-in-the-loop system (Publications II, III, IV).

Publications II, III and IV address RQ2. They develop methods that optimize queries to be informative for a task in two types of prospective tasks: individualized prediction (so-called targeted Bayesian experimental design in Publication II), and a decision-making problem of choosing the best action from a set of options for an individual (information gain on the entropy of the decision in Publications III and IV).

RQ3 is studied in Publication V, which introduces a human-in-the-loop method for reward elicitation in *de novo* molecular design. The feedback is on whether or not a molecule matches the expert's goal in a molecule design task. The user-model captures this goal either in a model of parameterized sub-objectives, or in a flexible non-parametric model for capturing a novel sub-objective. The queries are chosen by optimizing the exploration-exploitation trade-off similarly as in the user-model of Publication I. The user-model is then used for scoring the generated molecules during reinforcement learning training phase, which increases the quality of the

generated molecules w.r.t. the expert's goal.

1.2 Outline of the Thesis

The next chapters of the thesis give general background by first introducing prior elicitation in Chapter 2, and then describing methods for actively selecting new training instances in Chapter 3. Chapter 4 discusses the characteristics of the data acquisition problem in active expert knowledge elicitation and existing methods for it, and presents the first main contribution of the thesis on interactive expert knowledge elicitation (RQ1). Chapter 5 considers the second main contribution of the thesis, goal-oriented query selection (RQ2). Finally, Chapter 6 describes application of such human-in-the-loop methods to new area of *de novo* drug design (RQ3), and Chapter 7 concludes the work and contributions of the thesis.

2. Prior Elicitation

In Bayesian analysis and probabilistic modeling, the aim is to infer a *posterior distribution* of model parameters $\theta \in \Theta$. To do so, the *likelihood* of observing data $\mathbf{y} \in \mathcal{Y}$ given the model parameters, $p(\mathbf{y} \mid \theta)$, is combined via Bayes' rule with a *prior distribution* that represents *a priori* belief about model parameters $p(\theta)$. The resulting posterior distribution is $p(\theta \mid \mathbf{y}) = p(\mathbf{y} \mid \theta)p(\theta)/p(\mathbf{y})$, where the denominator is so-called *evidence* that represents the probability of observing $\mathbf{y}$ given the chosen model-family: $p(\mathbf{y}) = \int p(\mathbf{y} \mid \theta)d\theta$. Domain knowledge can be incorporated to the model in the prior distributions $p(\theta_i)$, $i = 1...p$ where p is the number of parameters or features. The posterior distribution can then be used to compute summaries, such as maximum a posterior (MAP) estimate of θ, or to produce predictions about new observables $\tilde{\mathbf{y}}$. The posterior predictive distribution of $\tilde{\mathbf{y}}$ is obtained by integrating out the posterior uncertainty in model parameters: $p(\tilde{\mathbf{y}} \mid \mathbf{y}) = \int p(\tilde{\mathbf{y}} \mid \theta)p(\theta \mid \mathbf{y})d\theta$, assuming independent and identically distributed observations (i.i.d.). Bayesian modelling and statistics are comprehensively presented e.g. in [3, 4, 5, 6].

Prior elicitation considers the process of acquiring domain knowledge from the expert(s) via questionnaires or interviews, and using them to form a prior distribution of a Bayesian model [7]. A recent survey in prior elicitation can be found in [8]. Usually the experts provide summaries, e.g. quantiles or moments of a probability distribution. Aggregating knowledge from multiple experts is discussed in [9] and [10]. In indirect prior elicitation [11, 12, 13, 14], the goal is to infer the priors of unobservable parameters from observable quantities. The focus in earlier works has been in defining the subjective probability distributions or their quantiles as

accurately as possible. For example, the term overfitting in prior elicitation means eliciting more quantiles than required for fitting a distribution, in order to mitigate the effect of erroneous judgements [7]. These methods are therefore mostly feasible for problems with either relatively small number of model parameters to elicit, or multiple domain experts.

Furthermore, even studies on interactive tools for prior elicitation have so far only considered the question of how to reliably and easily elicit quantiles or full prior distributions of each parameter [15, 16, 17, 18]. For example, [18] developed a web tool for remote prior elicitation, where a domain expert defines quartiles or tertiles of the distribution, or probability that the value lies in a certain range, and a facilitator with statistical knowledge then selects the distribution to be fitted. Indirect predictive elicitation in [12] presents how querying the expert about predicted outcomes is used to elicit priors in low-dimensional ($p = 5$) regression problems interactively [19]. The approach in this thesis is different, with the insight that multiple rough estimates of the priors of well-chosen parameters are more useful in high-dimensional problems where the number of parameters is hundreds or thousands.

The way to solve ill-posed problems with more parameters than observations ("small n large p") in non-probabilistic methods is regularization. In Bayesian analysis, weakly informative priors define the regularization of the problem. For example, in linear regression, the common L2-norm regularization is equivalent to Gaussian priors on the regression coefficients, and L1 ('lasso regularization') is achieved with Laplace priors [3, 5]. More sophisticated sparsity inducing priors include spike and slab [20, 21] and horseshoe prior [22, 23]. From this point-of-view, eliciting prior knowledge of a domain expert can alleviate the small n large p problem via information about sparsity.

Conceptually, prior elicitation is unrelated to active learning in that it acquires *a priori* knowledge, not new data. However, recent works including this thesis show that ideas borrowed from active learning are suitable also for actively eliciting prior distributions [8]. This is especially useful in cases where it's not feasible for a domain expert to carefully define the priors of all model parameters.

3. Active Data Acquisition Methods

3.1 Active Learning

Active learning is a vast field that studies how to choose new training instances for humans to label in order to improve machine learning model's performance [24]. The problem is to select a new instance $\mathbf{x} \in \mathcal{X}$, called a query, to get its label $y \in \mathcal{Y}$. The new labeled instance is then included to the labeled training data $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{n}$ of n examples. Typical active learning scenario is pool-based, where the query is chosen from a pool of unlabeled instances $U \subset \mathcal{X}$. The model's performance usually means how well it generalizes to unseen data, which is evaluated as the prediction accuracy in hold-out test data.

Even though classification is the most common application in active learning, many of the methods are straightforward to extend to regression tasks, where y is continuous. If there are multiple outcomes (i.e. for a single observation y_i is a vector), then the problem is called multi-task active learning. Typically labels are acquired either using relatively cheap crowdsourcing, e.g. in Amazon's Mechanical Turk[1], or expensively by interacting with domain experts, for example to segment and label medical images [25]. This thesis focuses on the latter, with the goal of interacting with domain experts.

The most common methods for selecting the next query rely on uncertainty, the expected error reduction, disagreement in committee methods,

[1] https://www.mturk.com/

and experimental design. Uncertainty-based strategies are in general called uncertainty sampling [26], where the idea is to select the query that the current model is the most uncertain of. For example, for a binary classification task, this is equivalent to selecting queries that are close to the decision boundary. Formally, for a probabilistic model with parameters θ, uncertainty sampling selects the instance from U with maximum entropy:

$$\mathbf{x}^* = \arg\max_{\mathbf{x} \in U} \mathcal{H}(Y \mid \mathbf{x}, \theta),$$

where $\mathcal{H}$ is the entropy [27] of the label Y, or differential entropy in regression, given features $\mathbf{x}$ and the model parameters θ. [24]

Disagreement based methods [28], such as query by committee, alleviate model misspecification problems by selecting the instance for which the committee of models disagree the most in the predicted label. They are suitable for cases when it is cheap to train multiple models, probabilistic methods are too expensive, or model selection is difficult.

Expected error reduction methods aim at minimizing the expected generalization error of the model. Typically the error in the test set is approximated empirically from the training data, and the query is chosen so that it minimizes the expected mean squared error (MSE) in the training data:

$$\mathbf{x}^* = \arg\min_{\mathbf{x} \in U} \mathbb{E}_{p(y|\mathbf{x},\mathcal{D})} \left[\frac{1}{n} \sum_{i=1}^{n} (y_i - \hat{y}_i)^2 \right],$$

where $\hat{y}_i$ is the model's prediction for $\mathbf{x}_i$ after adding a new observation $(\mathbf{x}, y)$ to the training data. The downside of the error reduction methods are their high computation cost [24], which arises from the need to approximate the expectation and to re-fit the model multiple times except for the simplest models.

Experimental design approaches choose queries based on their expected information gain or mutual information between the query and model parameters. Bayesian Active Learning by Disagreement (BALD) [29] uses the mutual information, and computes it in the space of observations:

$$\mathbf{x}^* = \arg\max_{\mathbf{x} \in U} \mathcal{H}(Y \mid \mathbf{x}, \mathcal{D}) - \mathbb{E}_{p(\theta|\mathcal{D})} \left[\mathcal{H}(Y \mid \mathbf{x}, \theta) \right].$$

It selects an instance $\mathbf{x}$ for which the prediction has high entropy, similar to uncertainty sampling, but for which predictions using posterior draws of

parameters have low uncertainty, which can be interpreted as the sampled models disagreeing with each other. Other Bayesian experimental design approaches are described in detail in Section 3.3.

Variance reduction selects the query that will reduce the expected variance of the predicted labels the most. The method can be either seen as a variant of expected error reduction, because reducing variance reduces the error for unbiased estimates, or as an approximation of the expected entropy reduction, where the entropy of a variable is approximated by its variance [24].

3.1.1 Goal-Oriented Active Learning

The second main contribution of this thesis is to develop active learning methods that take the prospective decision-making task into account. Concurrent works have been considering similar problems of optimizing the acquisition for other tasks than model accuracy.

Among these works, the closest to the problem setup in this thesis are [30, 31] and [32]. The work in [30] proposes myopic posterior sampling to solve a goal-oriented experimental design problem of selecting queries that maximize the reward in a task. They extend Thompson sampling [33] to take into account the goal of a prospective task, by defining a reward function of the task, and optimizing the acquisitions for it. On the other hand, [31] introduce a goal-oriented active learning framework that considers the expected, minimum or maximum effect of a query to different model evaluation metrics, such as Fisher information or prediction entropy.

Neither of these methods explicitly account for a decision-making task where the model is to be used as decision-support, which is the main contribution in Secion 5.2. The decision-making is considered in [32], which presents a heuristic uncertainty sampling based method for taking the decision-making task into account in a binary decision problem with binary outcomes. They also show experimentally that, if applicable, the approach of modeling the outcomes of decisions and selecting the optimal decision according to Bayesian decision-theory [5] is superior to the approach of directly modeling the decisions.

Other decision-centric active learning works have proposed ways for

reducing uncertainty to improve model-based decision-making. The idea in [34] is to drive all remaining hypothesis into the same decision region, which is done by modeling the overlapping decision regions as a hypergraph and selecting queries that maximally reduce the probability mass of the edges. Another approach is based on Mean Objective Cost of Uncertainty (MOCU), which is used to reduce the uncertainty that amounts to classification error with with 0-1 loss [35, 36]. The idea in Publications III and IV is similar, to reduce the uncertainty related to a decision, but for a case when the decision comes after learning a multi-task regression model that predicts the outcomes of alternative decisions. The optimal decision is modelled as a categorical random variable, and the queries then aim at reducing its expected entropy, following the Bayesian experimental design paradigm (see Section 5.2 for details).

3.2 Multi-Armed Bandit and Bayesian Optimization

Selecting queries that are meaningful to an expert and also help in improving model performance in the task brings out the exploration-exploitation trade-off: Whether to find the best choice based on current data, or explore further to learn more about the problem. Minimizing cumulative regret in this setup has been extensively studied in multi-armed bandit literature (see e.g. a recent book [37]). The most common approach is to select a query $\mathbf{x}$ that maximizes an upper confidence bound (UCB) [38]. The UCB selection is $\mathbf{x}^* = \arg\max_{\mathbf{x} \in U} r(\mathbf{x}) + c(\mathbf{x})$ where $r(\mathbf{x})$ is the expected reward of selecting $\mathbf{x}$, and $c(\mathbf{x})$ is a confidence bound and U is the set of options, called arms in bandit literature. The different choices of $c(\mathbf{x})$ lead to different regret bounds. Another common method is Thompson sampling, which is originally a heuristic that chooses the query by maximizing the expected reward given a randomly drawn belief from the reward model [39, 33].

Bayesian optimization solves an optimization problem of finding the maximum of an unknown black-box function f with minimum number of steps, by learning a probabilistic surrogate model for the objective function [40, 41]. The query is chosen by maximizing an acquisition function $\alpha(\mathbf{x})$, which can be based on, for example, the expected entropy of the

maximum [42, 43] following Bayesian experimental design approach, or similarly to the bandit problem, the UCB criterion [44]. Bayesian optimization has been successfully applied to many design problems, including generating and optimizing small molecules [45], and optimizing material composition in material sciences [46, 47, 48].

Obtaining the value $f(x)$ often implies running physical experiments (e.g. synthesize material and measure its properties), but they may include human-in-the-loop components as well. For example, in [48] in addition to an automated experiment, an expert processes raw information to compute a material property in answer to a query. This takes 20-25 min per query, which highlights the importance of minimizing the number of queries. Bayesian optimization can in some cases also be used to elicit expert knowledge, for example on the minimum-energy-configuration of a molecule via preferential queries [49].

3.3 Bayesian Experimental Design

Bayesian experimental design considers the question of how to select a set of experiments to be maximally informative, usually for the task of inferring some unknown parameter(s) $\theta \in \Theta$ [50]. The sequential version of experimental design addresses the same problem as active learning: which instance x ("an experiment") to choose to acquire the outcome y next. Traditionally, the experiments are physical experiments to determine e.g. the toxicity of a chemical. However, when the experiment is understood in a wider sense, the experimental design paradigm is applicable also to selecting which queries to pose to an expert.

In Bayesian experimental design, the information gain from design ξ and its observed outcome y is defined as the Kullback–Leibler divergence (KL-divergence) between the prior and posterior distributions [50, 51]

$$\begin{aligned} IG(\xi, y) &= \mathrm{KL}[p(\theta \mid y, \xi) \parallel p(\theta)] \\ &= \int_{\Theta} p(\theta \mid y, \xi) \log \frac{p(\theta \mid y, \xi)}{p(\theta)} d\theta \end{aligned} \tag{3.1}$$

The expected information gain (EIG) from ξ is then

$$U(\xi) = \int_{\mathcal{Y}} IG(\xi, y) p(y \mid \xi) dy$$

$$
\begin{aligned}
&= \int_{\mathcal{Y}} \int_{\Theta} p(\theta, y \mid \xi) \log \frac{p(\theta \mid y, \xi)}{p(\theta)} d\theta dy \\
&= \int_{\mathcal{Y}} \int_{\Theta} p(\theta, y \mid \xi) \log \frac{p(\theta, y \mid \xi)}{p(\theta) p(y \mid \xi)} d\theta dy \\
&= I(\theta; y)
\end{aligned}
$$

which is exactly the mutual information $I(\theta; y)$ between θ and the new observation. An alternative definition for the information gain is the gain in Shannon information (i.e. reduction of Shannon entropy), but when the question is to maximize the expected information gain from design ξ, these two become equivalent [50].

In the active learning setting, the corresponding acquisition function is then

$$\alpha_{\text{EIG}_\theta}(\mathbf{x}) = \int_{\mathcal{Y}} \text{KL}[p(\theta \mid \mathcal{D} \cup \{(\mathbf{x}, y)\}) \parallel p(\theta \mid \mathcal{D})] p(y \mid \mathbf{x}, \mathcal{D}) dy.$$

Alternatively, the utility can be defined as the expected information gain in prediction of the outcome for $\mathbf{x}_i$ in $\mathcal{D}$:

$$\alpha_{\text{EIG}_y}(\mathbf{x}) = \int_{\mathcal{Y}} \left[\sum_{i=1}^{n} \text{KL}[p(\tilde{y} \mid \mathbf{x}_i, \mathcal{D} \cup \{(\mathbf{x}, y)\}) \parallel p(\tilde{y} \mid \mathbf{x}_i, \mathcal{D})] \right] p(y \mid \mathbf{x}, \mathcal{D}) dy,$$

which is used by [52] for sequential knowledge elicitation from experts by extending it to relevance feedback (details in the Section 4.1.2).

Already early works in Bayesian experimental design point out that the information gain should be with respect to the quantity of interest, and inference about θ is only one possible task [50]. However so far this has gained little attention in Bayesian active learning and interactive knowledge elicitation literature, where most works resort to optimizing information to θ, e.g. [53, 54], or predictions y, e.g. [55, 56, 29, 57, 52, 58].

Section 5.2 presents two examples of how to apply the sequential experimental design to other goals, for the interest of interactive knowledge elicitation in precision medicine and decision support. A similar approach has been used in bandit settings by Russo and Van Roy in [59], where they propose information directed sampling that uses information gain to find the optimal action.

4. Interactive Knowledge Elicitation

Interactive knowledge elicitation aims at fast and easy interaction with a human expert, in a way that still enables using the elicited expert knowledge to solve hard problems that arise from the lack of data. Even though active learning and prior elicitation are known to improve the accuracy of models, it is still unclear how to maximally benefit from this kind of indirect and limited-budget human involvement. Then the question is, how to use indirect queries and select them to reduce the labor of humans (**RQ1**).

This chapter presents an overview of active knowledge elicitation methods, existing approaches and two example case studies in the thesis. The first is an illustrative case for interactive prior elicitation in a text-based prediction task, evaluated in a user-study. In the second case, the method is extended to apply to a challenging high-dimensional genomics data, and demonstrated with medical domain expert feedback.

4.1 Parts of Interactive Expert Knowledge Elicitation System

This section describes the design and modeling choices involved in interactive knowledge elicitation for improving machine learning model's performance. When developing such a system, there are four aspects to consider:

- Feedback types and interaction modalities
- The integration of feedback to the model
- Query strategy

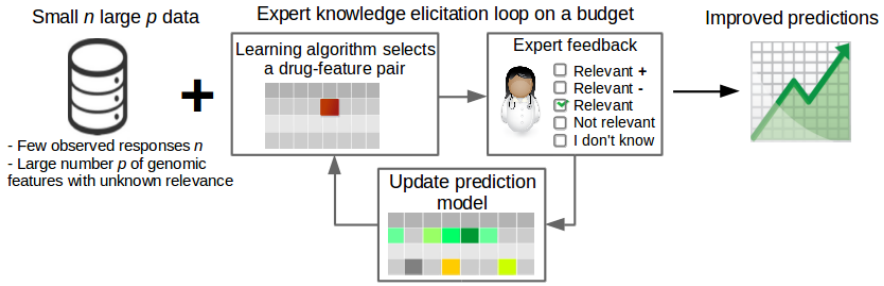

Figure 4.1. Overview of interactive expert knowledge elicitation for improving predictions in small-sample-size problems. Queries are presented sequentially to the expert, whose feedback indicates whether the queried feature is relevant in predicting the outcome, such as cancer cell's sensitivity to a drug. This type of indirect feedback can also include an indication in which direction the effect is likely to be. The query strategy aims at maximizing the performance with a minimal number of queries, i.e., on a budget. Reprinted with permission from Publication II.

- User-modeling

Figure 4.1 illustrates the workflow in interactive expert knowledge elicitation, in the application of small n large p genomics-based prediction in Publication II. The available sample size of drug sensitivity measurements is too small to identify a model that includes all genomic features. The human-in-the-loop approach iteratively queries a new pair of drug and feature from a domain expert, who indicates whether the feature is relevant for predicting cancer cell's sensitivity to the drug. The feedback is integrated to the prediction model, and its predictions consequently improve.

The modeling task is divided conceptually into two parts: a *user-model* and the model of the phenomenon, but depending on the problem the implementation may either consist of two separate models that interact, as in Publications I and V, or a joint model of the user and the data, as in Publications II, and III. The user-model describes the process of how the user gives feedback, for example the noise and potential biases, and models their knowledge about the phenomenon.

4.1.1 Feedback Types and Integration to a Model

Essential part of the interaction are the ways in which the user can communicate their ideas to the model, that is, the type of feedback and more gen-

erally the interaction modalities. The choice of the most suitable feedback type is highly application dependent: in some active learning scenarios the user giving the labels of instances is the most natural, whereas in others, indirect feedback such as feedback on features is easier [55, 60]. Together with the choice of the feedback type, comes the choice of how to incorporate the feedback into the model.

Indirect feedback has been studied more in non-active settings than in the active setting. For example, in [61] Ribeiro et al. developed a method to find the features that best explain the predictions of a classifier, and showed them to humans. They found that even persons who are not experts in machine learning could identify which features are unimportant for a classification task, which consequently improved the model when those features were removed.

Interactive knowledge elicitation for clustering has been studied in [62], where the user is shown potentially important features, and the user may decide to include relevant features or reject irrelevant ones. Strictly rejecting or accepting features is, however, sensitive to errors in feedback. In [63], the user indicates pairs of instances belonging to the same cluster, which is taken into account in the model by pairwise must-link/cannot-link constraints in the likelihood of a discriminative classifier that represents the cluster preferences. In classification, [55] introduces an active learning framework for querying both features and instances simultaneously. The relevant features are incorporated to the model by scaling. In [60], feedback type is a feature's class label, and the feedback is taken into account by modifying a multinomial prior distribution in a classification model to favor the selected class.

For regression problems, similarity feedback about the features is introduced in [64], and [65] assume the feedback about a feature is about the value of the regression coefficient associated with the feature in linear regression. A Bayesian framework for sequential knowledge elicitation with relevance feedback was introduced in [52], and it is applied and extended to richer feedback types in this thesis.

Two of the types of indirect feedback about feature prior distributions developed in this thesis are relevance feedback and directional feedback. Section 4.2 presents *relevance* feedback to the prediction task, where

the priors of relevant features are modified accordingly (Publication I). Section 4.3 extends the relevance feedback to cover both positive and negative association with the outcome using the framework from [52], and introduces the new feedback type about the sign of the correlation (Publication III).

Another type of indirect feedback is *counterfactual elicitation*, introduced in Publication III. It presents two indirect feedback types for an individual treatment effect estimation problem: feedback about counterfactual outcomes, and comparisons of potential outcomes, as described in Section 4.4. Integration of this feedback to a model is described in Section 5.2.1.

Summary of indirect feedback types in literature

- Feedback on instances
  - Similarity of instances in a classification task [63]
  - Per-example feature selection: Inclusion of a feature for predicting an instance x in a classification task [66]
- Feedback on features
  - Feature selection: Inclusion of a feature in a classifier [55, 62, 61]
  - Association of a feature to a class in classification [60]
  - Similarity of features in regression [64]
  - The value of regression coefficient [65, 52]
  - Relevance of a feature, defined as potentially erroneous indication about inclusion of the feature in a probabilistic regression model [52], (PII)
  - Relevance of a feature defined as indication of positive correlation with the outcome (PI)
  - Directional feedback about the sign of the correlation of the feature with the outcome (PII)
- Counterfactual elicitation: asking "what if"-queries (PIII)
  - The value of a counterfactual outcome, e.g. "What if a patient had gotten treatment B instead of the treatment A?"

– Comparison of two potential outcomes, e.g. "Would the patient have benefited more from treatment B than the treatment A?"

4.1.2 Query Strategy

In high-dimensional problems, all potential queries cannot be asked from an expert to elicit their knowledge, and hence the strategy of selecting the queries is important. Two factors that concern the selection of the query are informativeness of queries and the exploration-exploitation trade-off. A specific challenge in interactive knowledge elicitation is the small query budget of tens or at most hundreds of queries.

The queries $q \in Q$ are rated by an acquisition function $\alpha(q) : Q \longrightarrow \mathbb{R}$, where Q is the set of all possible queries, and the next query is chosen by $\arg\max_{q \in Q} \alpha(q)$. The acquisition function can be derived from active learning, experimental design (e.g. [55, 52], and Publication II), Bayesian optimization (e.g. [49]), multi-armed bandit algorithms (e.g. Publication I) or reinforcement learning (e.g. [66]), depending on the purpose of elicitation. The first two approaches focus on modeling, and the queries are chosen by their informativeness to learning a model. In contrast, bandit problems and reinforcement learning explicitly address the exploration-exploitation trade-off by balancing between exploring to learn a model or selecting a query with a high expected reward. Furthermore, Bayesian optimization aims at finding a query that maximizes an unknown function.

It is important to note that for indirect queries, standard active learning does not provide a clear solution as to what to choose as α. Proposed solutions in [55, 60] introduce active learning for querying both features and instances in text classification. The feature selection is done in batches, and the acquisition function in [55] is the mutual information (MI) of a binary feature and the class label

$$\alpha_{\mathrm{MI}}(j) = \sum_{c \in 0,1} \sum_{\tau_j \in 0,1} P(c, \tau_j) \log \frac{P(c, \tau_j)}{P(c)P(\tau_j)},$$

where c is the class label and τ_j indicates whether the feature j is present in an instance.

The above approach does not explicitly account for expert specific factors, such as the noise in the expert's answer being different than the noise in

the data. In this thesis, the solution is to define a user-model that includes these factors.

Daee et al. [52] proposed the approach of jointly modeling the data and the expert, for sequential knowledge elicitation for high-dimensional regression problems. The queries are about binary relevance of a feature: $f \in \{0, 1\}$, where $f_j = 1$ if the j:th feature is relevant, and $f_j = 0$ if not. The information gain is defined as the KL-divergence between predictive distribution before and after observing the feedback

$$\alpha_{\text{EIG}}(j) = \mathbb{E}_{p(\tilde{f}_j|\mathcal{D},\mathcal{F})}\left[\sum_{i=1}^{n} \text{KL}[p(\tilde{y} \mid \mathbf{x}_i, \mathcal{D}, \mathcal{F} \cup \{(j, \tilde{f}_j)\}) \parallel p(\tilde{y} \mid \mathbf{x}_i, \mathcal{D}, \mathcal{F})]\right] \tag{4.1}$$

where $\mathcal{F}$ contains feedback on queries that have already been asked. The expectation is taken w.r.t. the model's posterior predictive distribution of the expert's answer $p(\tilde{f}_j \mid \mathcal{D}, \mathcal{F})$, which comes form the user-model and takes into account the potential errors in the expert's answer.

This thesis proposes selecting queries in two alternative ways. The first is by explicitly modeling the user's knowledge and selecting queries that help to learn about the knowledge (Section 4.2). The second extends the sequential experimental design approach from [52] to multi-output and more descriptive feedback (Section 4.3).

4.1.3 User-Modeling

A core idea in expert knowledge elicitation is to model the feedback-giving process and the expert in a user-model, which allows personalizing the interaction to the expert. Two approaches have been studied in this thesis: First is to simplify the problem by learning an separate model about expert's knowledge as in Publication I. This kind of user-model has been used for suggesting features that are likely to be relevant in the expert's opinion (PI), or finding new potentially good molecules in a molecular design task (PV). The second approach is a full joint model of the expert's answers and the data, as first suggested in [52], which simultaneously optimizes the informativeness of queries and takes into account expert specific aspects. Even though the expert's answers are here interpreted as observations about latent quantities of the data or a model, this kind

of user models also include expert-specific parameters, such as noise in their answer (Publications II and III), or cognitive aspects that help in interpreting the expert's feedback [67].

User-modeling has a striking effect when the expert's bias is taken into account. A user-model based on Bayesian knowledge update by Daee et al. [67] was shown to be able to disentangle the prior knowledge of the expert from a new evidence. This enhanced the usefulness of the elicited prior knowledge by removing overfitting arising from the user having seen model-estimated relevance of features.

An example of user-modeling in non-interactive context is [68], which introduced learning human-like kernels for Gaussian processes based on experimental data, where human subjects had been asked to extrapolate a one-dimensional function based on previous observations from it. They showed that human-kernels captured inductive biases of humans about function shapes, and provided a model of how humans extrapolate functions.

4.2 Interactive Prior Elicitation

The goal of interactive prior elicitation is to reduce the effort of a domain expert compared to standard prior elicitation. In this thesis, it is done by querying indirect feedback about the prior distribution and using a user-model to choose which prior distributions to elicit (**RQ1**).

This section presents a case-study on interactive prior knowledge elicitation to solve "small n large p" regression problem of predicting outcome $y \in \mathbb{R}$ based on n observations of instances $x \in \mathbb{R}^p$ in training data $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$, where $p > n$. The interactive method elicits knowledge about feature relevance, and it is shown to improve model prediction accuracy significantly in a user-study where the task is to predict citations of scientific articles based on their keywords.

Publication I addresses the high-dimensional prediction problem by modifying prior distributions according to queried relevance r_j of a feature j, for a subset of all features $j = 1, ..., p$. To select which features get feedback, Publication I develops a multi-armed bandit based user-model for finding

the most relevant features according to the expert's knowledge.

In this case study, relevance is defined as the feature being positively correlated with the outcome, even if not necessarily in the training data. Feedback is included in the model as indication of the type of prior distribution. As a consequence, the relevant features are regularized less than the non-relevant ones, but restricted to be positive. This means that the expert-indicated relevant features will more easily have larger weight than other features, if data provides evidence for it.

Model

The prior distributions of a standard Bayesian linear regression model [3]

$$y_i \sim N(\mathbf{w}^T \mathbf{x}_i, \sigma^2),$$

$$w_j \sim N(0, \sigma_0^2)$$

are modified so that the weighs w_j of relevant features have half-normal prior distributions (Half-N):

$$w_j \sim N(0, \sigma_0^2), \text{ if } r_j = 0,$$

$$w_j \sim \text{Half-}N(0, a\sigma_0^2), \text{ if } r_j = 1. \tag{4.2}$$

The constant a determines the overall ratio of the effect sizes between relevant and non-relevant features; at initialization $r_j = 0$ for all j. Parameters σ is the residual variance in regression and σ_0 the prior variance of the weights. The inference, i.e. the computation of the posterior distribution of w given $\mathcal{D}$, is done via MCMC (Markov Chain Monte Carlo) sampling in a probabilistic programming language Stan [69].

User-Model and the Query Strategy

Publication I introduces a linear multi-armed bandit based user-model, that learns which features the human thinks are relevant, and selects queries based on their upper confidence bound (UCB) [38]. To extrapolate feedback on one feature to other features, the features are described by auxiliary descriptors, $z_j \in \mathcal{R}^{n_z}$, which in our case-study are occurrences of keywords in clusters of scientific documents. The user-model then estimates the relevance of features as

$$\hat{r}_{j,t} = z_j \hat{\mathbf{v}}_t + b \quad \forall \quad j \in 1, \ldots, p,$$

where $t = 1...T$ indexes the interaction iteration, $\hat{\mathbf{v}}_t$ is a model parameter encoding the expert's knowledge about relevant features, estimated from feedback up to the iteration $t-1$, and b defines the default relevance. The queries are chosen by maximizing the UCB

$$r_{j,t}^{UCB} = \hat{r}_{j,t} + c_{j,t},$$

where $c_{j,t}$ is a high probability bound of the uncertainty of $\hat{r}_{j,t}$ computed using SupLinUCB from Theorem 1 in [70]. This means that the algorithm selects features which either are anticipated to be relevant in the expert's opinion (exploitation) or have high uncertainty or potential to be relevant (exploration).

Evaluation

In the user study of Publication I, the machine-learning task is to predict citation counts based on $n = 162$ scientific documents described by presence or absence of $p = 457$ unique keywords. The task of participants was to indicate whether or not they think the proposed keywords are relevant in predicting the citations of an article in artificial intelligence domain. Further details and the interactive visualization and user-interface developed for the user-study are described in Publication I.

Results

The proposed method was evaluated in a user-study of 23 participants in Publication I. The results are summarized in Figure 4.2, showing how the error in prediction decreases with increasing feedback. The model predictions improve for all participants, showing that including human feedback is beneficial. Furthermore, the proposed user-model approach is significantly better than randomly selecting the queries in interactive prior elicitation. In addition, the participants found this type of feedback easy to give: self-reported difficulty in a post-study survey was 50% easy and 29% neutral.

Summary

In a user-study in Publication I, the interactive prior elicitation method was shown to improve significantly over a model without interaction, and it was also more query efficient than choosing the queries randomly. One

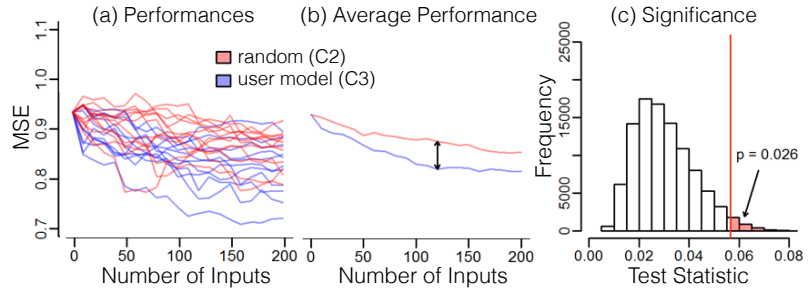

Figure 4.2. Summary of interactive prior elicitation user-study results and significance test showing that the user-model approach is significantly better than random. (a) Mean squared error of the prediction with respect to the number of feedback from a domain expert. Curves correspond to individual participant. (b) The maximum distance between average performance curves is chosen as a test statistic. (c) Distribution of the test statistic in 10^6 permutations. The vertical line indicates the value of the test statistic in the user-experiment. Reprinted with permission from Publication I.

strength of the method is that the choice of prior distributions in eq. (4.2) makes the model less sensitive to errors in features that are deemed irrelevant, because they are not excluded from the model. In contrast, the model is potentially more sensitive to erroneously identified relevant features. Nevertheless, the post-questionnaire results and the empirical performance of the method indicate that relevant features are easy to identify by experts in this kind of prediction tasks.

4.3 Active Expert Knowledge Elicitation for Prediction from Genomic Features

In the interactive prior elicitation method presented in the previous section, expert's opinion modifies the prior distributions of model parameters significantly. As a result, the method is sensitive to erroneous feedback. One elegant way to avoid this is to treat the expert knowledge as another data source, as introduced by Daee et al. in [52]. This section presents an extension of the methods in Publication I and [52] to a medical application case with richer feedback types. The method is demonstrated with domain expert feedback in a case-study of predicting drug sensitivity scores from genomic features in Publication II.

This section addresses RQ1 in following two aspects. How to use active

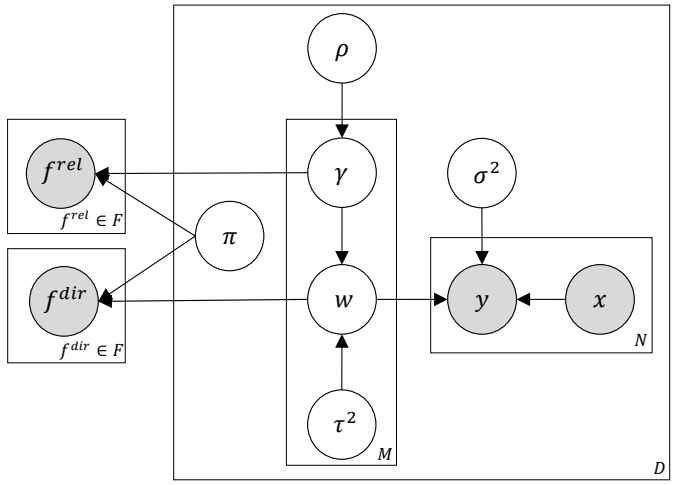

Figure 4.3. Plate diagram of the sparse linear regression prediction model and the feedback observation model, where grey circles represent observed variables, and white are unobserved. Feedback from a human expert, f^{rel}, is a noisy observation about the sparsity γ, and f^{dir} about the association w of the genomic features to predicting sensitivity to a drug. Reprinted with permission from Publication II.

knowledge elicitation to solve the regularization problem in real-world medical small n large p data? How to incorporate indirect feedback on feature relevance and quality of the association to maximally benefit from expert's knowledge?

The extension of feedback types in Publication II from Publication I is to incorporate knowledge of both positive and negative correlation with the outcome, namely "relevant" (correlates with the outcome) "not relevant" (no correlation), and allow "I don't know" answers (treated as no feedback). In addition, for the 'relevant' case two subcategories of directional feedback are introduced: 'positive' and 'negative' for the sign of the correlation, which extends the work from [52].

Model

The model is defined as joint probability of observations of both data and feedback, and the relations of variables are described in a plate-diagram in Figure 4.3. The observation model of data assumes that the outcomes y_d of drug $d = 1, \ldots, n_d$ are independent from each other and depend linearly

on the features $\mathbf{x} \in \mathcal{R}^p$ as

$$y_{i,d} \sim \mathrm{N}(\boldsymbol{w}_d^\top \mathbf{x}_i, \sigma_d^2),$$

where $y_{i,d}$ is a drug sensitivity score of a patient i to drug d, and $\mathbf{x}_i$ are their features. The parameters $\boldsymbol{w}_d$ are regression coefficients and σ_d^2 is the residual variance of the outcome model for drug d. The training set is $\mathcal{D} = \{(\mathbf{x}_i, y_{i,1}, y_{i,2}, \ldots, y_{i,n_d})\}_{i=1}^n$.

Sparsity of features is implemented with spike-and-slab priors [20, 21]:

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

$$\gamma_{d,j} \sim \mathrm{Bernoulli}(\rho_d),$$

where $w_{d,j}$, $j = 1, \ldots, p$, is the regression coefficient of the j:th feature, and $\gamma_{d,j} \in \{0, 1\}$ indicates whether the feature is relevant in predicting outcome of d. A Dirac delta function at 0, δ_0, shrinks the weights of irrelevant features to zero, resulting in a sparse model with the the expected sparsity of ρ_d. Hyperparameter $\tau_{d,j}^2$ defines the prior variance of relevant features. Hyperpriors used in the experiment are described in (PII).

User-Model

Observations of the expert feedback are included to the model via feedback observation models. The observations are $f_{d,j}^{rel} \in \{0, 1\}$ for the relevance of feature j in predicting outcome of drug d (0 for not relevant and 1 for relevant), and $f_{d,j}^{dir} \in \{0, 1\}$ for the directional feedback (0 for negative and 1 for positive correlation). Essentially the expert is regarded as a new data source that yields noisy observations about the relevance parameter $\gamma_{d,j}$ and the sign of the regression coefficients:

$$f_{d,j}^{rel} \sim \gamma_{d,j}\, \mathrm{Bernoulli}(\pi_d^{rel}) + (1 - \gamma_{d,j})\, \mathrm{Bernoulli}(1 - \pi_d^{rel}), \text{and}$$

$$f_{d,j}^{dir} \sim \mathbb{I}(w_{d,j} \geq 0)\, \mathrm{Bernoulli}(\pi_d^{dir}) + \mathbb{I}(w_{d,j} < 0)\, \mathrm{Bernoulli}(1 - \pi_d^{dir}), \quad (4.3)$$

where π_d^{rel} and π_d^{dir} are the probabilities of the expert being correct in relevance and directional feedback respectively. $\mathbb{I}(C) = 1$ when the condition C holds and 0 otherwise.

Query Strategy

With multiple feedback types, it is not obvious how to define a reward of a bandit user-model as in Publication I. Instead, the sequential experimental

design approach also used in [52] provides a more principled solution. The acquisition function for queries are extended from [52] to account for multiple outputs d and two feedback types. The next query, a (drug,feature) pair (d^*, j^*) is selected by maximizing the following expected information acquisition function of the expected information gain:

$$\alpha_{\text{EIG}}((d,j)) = \mathbb{E}_{p(\tilde{f}_{d,j}^{rel}, \tilde{f}_{d,j}^{dir} | \mathcal{D}, \mathcal{F})} \left[\sum_{i=1}^{n} u_{i,d,j} \right], \tag{4.4}$$

where the utility $u_{i,d,j}$ is the information gain of feedback $(\tilde{f}_{d,j}^{rel}, \tilde{f}_{d,j}^{dir})$ in predicting the outcome $\tilde{y}_d$ of patient $\mathbf{x}_i$:

$$u_{i,d,j} = \mathrm{KL}[p(\tilde{y}_d \mid \boldsymbol{x}_i, \mathcal{D}, \mathcal{F} \cup \{(\tilde{f}_{d,j}^{rel}, \tilde{f}_{d,j}^{dir})\}) \parallel p(\tilde{y}_d \mid \boldsymbol{x}_i, \mathcal{D}, \mathcal{F})],$$

and $\mathcal{F}$ are queries for which feedback has been collected up to iteration t. This is the expected information gain from feature j to prediction of the sensitivity to drug d, approximated as an average over the prediction in training data. This approximation optimizes for in-sample accuracy, since it has been shown that maximizing the expected information gain is equivalent to minimizing the expected log-loss (Chapter 6.1 in [24]).

The second proposed query strategy, targeted expected information gain approach (T-EIG) aims at optimizing for personalized prediction. It can be applied in cases where the features of a patient of interest, $\tilde{\mathbf{x}}$, are known. This strategy is described in more detail in Section 5.1, but presented also here for completeness:

$$\alpha_{\text{T-EIG}}((d,j)) = \mathbb{E}_{p(\tilde{f}_{d,j}^{rel}, \tilde{f}_{d,j}^{dir} | \mathcal{D}, \mathcal{F})} \left[\tilde{u}_{d,j} \right], \text{where} \tag{4.5}$$

$$\tilde{u}_{d,j} = \mathrm{KL}[p(\tilde{y}_d \mid \tilde{\mathbf{x}}, \mathcal{D}, \mathcal{F} \cup \{(\tilde{f}_{d,j}^{rel}, \tilde{f}_{d,j}^{dir})\}) \parallel p(\tilde{y}_d \mid \tilde{\mathbf{x}}, \mathcal{D}, \mathcal{F})] \tag{4.6}$$

The only difference between eq.(4.4) and eq.(4.6) is the set of patients in which the utility is evaluated.

In order to compute posterior distributions $p(\boldsymbol{w}_d \mid \mathcal{D}, \mathcal{F})$ and further predictive distributions for $\tilde{y}_d$, we use approximate inference using expectation propagation algorithm [71]. In computation of the expected information gain, the posterior distributions $p(\tilde{f}_{d,j}^{rel}, \tilde{f}_{d,j}^{dir} \mid \mathcal{D}, \mathcal{F})$ and those of $\tilde{y}_d$ are approximated using partial expectation propagation updates [52, 72].

Results

Publication II demonstrated the method with domain expert feedback in a case study of predicting ex vivo drug sensitivity scores of multiple myeloma cancer, based on patient's genomic profile (somatic mutations and karyotype data). For the experiments, feedback was collected offline from two domain experts of multiple myeloma cancer: a doctoral candidate and a senior researcher. The proposed query strategies EIG and T-EIG were evaluated in a simulated experiment, where the sequential queries were answered using the offline collected feedback. The results showed that the active query strategies were efficient in improving prediction accuracy compared to random selection of queries. On average, T-EIG required only 11% and EIG 70% of the number of queries compared to random sampling to gain 80% of the potential improvement. Similar results were obtained in a simulated case study on phenotype prediction in Publication II, with T-EIG reaching 80% of the performance of a model where all potential feedback is included, in just 2% of all possible queries. The model with full feedback was also found to be confidently better than sparse models using elastic net regularization or spike and slab prior that did not include feedback.

Summary

The application of expert knowledge elicitation to high-dimensional genomics data is challenging: even domain experts' knowledge on genomic features is sparse, resulting in most queries to be answered by 'I don't know' (PII). Even more difficult, the query budget to experts in such a scenario is at most hundreds, which means that it is possible to get feedback on less than 10% of all potential features. The user-study in Publication I validated a proof-of-concept of the method, and this example case in Publication II demonstrates that the developed methods improve predictive performance in real-world medical small n large p data with expert feedback on reasonable number of features, that is, hundreds, instead of thousands.

4.4 Counterfactual Elicitation

One fundamental problem in causal inference is that the observations only consist of outcome of actions that were made. In contrast, what would have happened if another action had been taken, that is, the *counterfactual outcomes*, are always unknown. Publication III proposes that in some cases experts could have knowledge about the counterfactual outcomes. For example, that a disease progression would have been faster if a patient had not received a treatment, or that with hindsight another treatment plan would have been better. Therefore, it should be possible to benefit from this expert knowledge, and Publication III introduces an idea of *counterfactual elicitation*, where the expert is queried about potential outcomes of alternative actions in the past. Publication III demonstrates expert knowledge elicitation in counterfactual elicitation in a simulated medical human-in-the-loop scenario.

Counterfactual elicitation could in theory also be applied if there are other methods to access proxies of counterfactuals. Examples of such are new experiments on new instances with similar features, but especially in medical domain this poses severe ethical concerns: It is of course not ethical to administer a drug to a patient in order to improve treatment plan of another. Instead, expert knowledge elicitation provides means to access quantities that would be either practically or ethically impossible to acquire otherwise.

5. Goal-Oriented Active Learning for Expert Knowledge Elicitation

To answer RQ3, this chapter presents new acquisition functions for goal oriented active learning in two prospective tasks: personalized prediction, and decision-support task, where two methods are developed for binary and multi-option decision-making respectively. They aim at reducing drastically the number of interactions while still improving model performance in the prospective task.

Section 5.1 considers the problem of personalized prediction, and presents how knowing the test-task in advance facilitates the elicitation by selecting the most informative queries first. Section 5.2 develops query strategies for improving decision-making in a task of choosing an action to take from a set of discrete choices, based on data about past instances, observed actions and their outcomes.

5.1 Personalized Prediction

This section revisits the goal-oriented query selection strategy from Section 4.3, where the goal was to improve the accuracy of predictions in precision medicine. The example task was to use a machine learning model to predict how different targeted cancer drugs would affect a new patient, based on their genomic profile and data from a screening of *ex vivo* drug sensitivity scores.

As a recap of the setup, the outcomes y_d of drugs $d = 1, \ldots, n_d$ are predicted from individual's p genomic features $\mathbf{x} \in \mathbb{R}^p$, based on data $\mathcal{D}_0 = \{(\mathbf{x}_i, y_{i,1}, \ldots, y_{i,n_d})\}_{i=1}^n$, and $n < p$. Interaction with a domain expert yields feedback $f_{d,j}$ that evaluates qualitatively how feature j affects prediction

of outcome d. Details of the model and feedback types are presented in Section 4.3.

In active learning, the task is to choose a query $q = (d, j)$, and augment training data with expert's feedback f_q. At each iteration t, the feedback is added to training data $\mathcal{D}_t$, that contains both the observed data $\mathcal{D}_0$ and the feedback up to time t: $\mathcal{F}_t$.

To make active learning goal-oriented for this task, the insight is to select queries which improve maximally the personalized predictive accuracy for a target individual $\tilde{\mathbf{x}}$ (e.g. a new patient who needs treatment), instead of seeking to improve the model for the general population. Therefore, the information gain is only evaluated to the posterior predictive distribution of $\tilde{\mathbf{x}}$, which results in the following targeted expected information gain (T-EIG) criterion:

$$q^* = \underset{q \notin F_{t-1}}{\arg\max}\, \mathbb{E}_{\tilde{f}_q | q, \mathcal{D}_{t-1}} [\tilde{u}_{q,t}] \text{ where}$$

$$\tilde{u}_{q,t} = \mathrm{KL}[p(\tilde{y}_d \mid \tilde{\mathbf{x}}, \mathcal{D}_{t-1} \cup \{\tilde{f}_q\}) \parallel p(\tilde{y} \mid \tilde{\mathbf{x}}, \mathcal{D}_{t-1})].$$

Here $\tilde{f}_q$ is the expert's predicted feedback and $\tilde{y}_d$ is the predicted outcome of the target individual.The only difference to the common information gain approach in eq. (4.1) is simply to evaluate the information gain to predicting the outcome of the target individual ($\tilde{\mathbf{x}}$), instead of approximating the information gain to a general prediction task using the training data $\mathbf{x}_i$, $i = 1, \ldots, n$ or the unlabeled pool of active learning, which would be standard choices [24].

Publication II shows experimentally that if the features $\tilde{\mathbf{x}}$, i.e. the genomic profile of the patient is known in advance, then T-EIG is more efficient than traditional information gain approach in finding relevant queries, because it personalizes the queries to the prediction for $\tilde{\mathbf{x}}$.

5.2 Active Learning for Decision-Making

The goal in interactive knowledge elicitation is not merely to improve the accuracy of models. Its further aim is to make the models more useful in helping experts to make decisions, and the few interactions with the expert should benefit this task maximally. Therefore, this thesis studies how to

take a down-the-line decision-making task into account in active learning.

Specifically, the decision-making task considered in this thesis is: given models that predict outcomes under alternative actions, what is the best action to take for an individual. An example case is how to select a treatment to a patient, given their individual features $\mathbf{x}$, and past data about patients.

This chapter presents two acquisition functions that include the prospective decision-making task, and shows that they effectively improve decision-making even with extremely limited query budget to an expert (Publications III, IV).

5.2.1 Binary Decision Problem

In case of a binary decision problem, the question we consider is whether to treat a patient or not. Patients are represented by their features $\mathbf{x} \in \mathbb{R}^p$, e.g. age, sex, ethnicity and so on, and the data $\mathcal{D} = \{(\mathbf{x}_i, a_i, y_i)\}_{i=1}^n$ logs for previous patients if they were treated ($a = 1$) or not ($a = 0$, so-called 'control'), and their outcome $y_i \in \mathbb{R}$. The important difference to the case in Section 4.3 is that the data does not contain a complete set of measurements for all potential treatments. This type of logged data is called observational data in causal inference. We use the common potential outcomes framework in causal inference from [73] to fit two potential outcome models $p(y[a] \mid \mathbf{x}, \mathcal{D})$ that predict the potential outcome $y[a]$ of action a of patient $\mathbf{x}$. Other works have developed multi-task methods for potential outcome models, which fit a single model and enable sharing statistical strength between the outcomes, see e.g. [74]. Those methods are orthogonal to the methods presented in this thesis, and could be used to improve the performance further.

Making an assumption that there are no unmeasured confounders and that the potential outcomes are consistent with the observed outcomes, the outcome models can be estimated as conditional probabilities $p(y[a] \mid \mathbf{x}, \mathcal{D}) = p(y \mid \mathbf{x}, a, \mathcal{D}) := \hat{p}(y_a \mid \mathbf{x}, \mathcal{D})$. The current probabilistic model is denoted with $\hat{p}$ to emphasize that the probabilities from the model may differ from the true data generation process.

We assume that a decision-maker follows the Von Neumann-Morgenstern

expected utility theorem [75] and selects a that maximizes the expected utility $u(y, a)$ of the predicted outcome for $\mathbf{x}$:

$$a^* = \arg\max_{a \in \{0,1\}} \int_y u(y_a, a) \hat{p}(y_a \mid \mathbf{x}, \mathcal{D}) dy \tag{5.1}$$

This assumption is equivalent to assuming that the final decisions are made automatically by the model using Bayesian decision-theory [5]. A new insight is that when it is possible to collect more data, then it is not anymore inconsequential to separate model inference from decision-making as has been traditionally thought. In this thesis, the utility is assumed to be directly the outcome, $u(y, a) = y$, but the methods can be extended to more complex utilities as well.

Given the above assumptions about decision-making behaviour and the utility, the task can be summarized as inferring a conditional treatment effect $\tau \mid \mathbf{x}$. The treatment effect is the difference between the potential outcomes: $\tau \mid \mathbf{x}, \mathcal{D} \sim \hat{p}(y_1 - y_0 \mid \mathbf{x}, \mathcal{D})$, and the down-the-line decision will be to treat a patient if their expected treatment effect $\mathbb{E}[\tau]$ is positive. This insight will be used for deriving the active learning strategy.

Type S error rate

Publication III shows that, under certain assumptions, imbalance in observational data will lead to increased probability of wrong decisions in this setup. Here imbalance refers to imbalance in covariate distributions in treated and control groups. The main insight is that the rate of wrong decisions can be approximated form the model's posterior, as described next, and we can then choose queries that aim at minimize it.

Type S error, introduced in [76], is a measure for decision-making performance in this kind of binary decision-making tasks (PIII). The Type S error rate γ is the probability of a decision being incorrect, as illustrated in Figure 5.1, or, equivalently, the proportion of wrong decisions in repeated decision-making tasks.

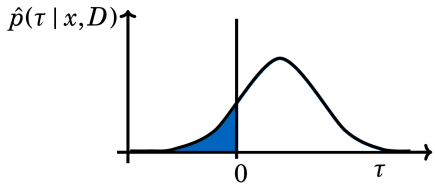

Figure 5.1. Posterior distribution of treatment effect τ given data $\mathcal{D}$ and features $\mathbf{x}$, $\hat{p}(\tau \mid \mathbf{x}, \mathcal{D})$. The optimal decision is to treat a patient $\mathbf{x}$ because the expected treatment effect $\mathbb{E}[\tau]$ is positive. Still, the probability of making a wrong decision is the tail probability of $\tau < 0$ (colored in blue), which is the estimated Type S error rate of the treatment effect model at $\mathbf{x}$.

Decision-Making Aware Acquisition Function

We demonstrate the method in a counterfactual elicitation scenario (see Section 4.4). The pool of queries $\mathcal{U}$ contains counterfactual examples $\{(\mathbf{x}_i, 1 - a_i)\}_{i=1}^n$, that is, the past training instances and the actions that were not observed in the training data. The selected queries $(\mathbf{x}_j, a_j)$ are answered by an expert who guesses the outcome y_j, and added to $\mathcal{F} = \{(\mathbf{x}_j, a_j, y_j)\}_{j=1}^t$ where t is the number of interactions with the expert so far.

For simplicity, we assume that the test individual $\tilde{\mathbf{x}}$ is known similarly to Section 5.1. However, the method can be extended to a test population as well, as shown shortly in Section 5.2.2. Publication III proposes to select queries to minimize the expected estimated Type S error rate:

$$(\mathbf{x}^*, a^*) = \underset{(\mathbf{x},a)\in\mathcal{U}}{\arg\min}\, \mathbb{E}_{\hat{p}(y_a|\mathbf{x},\mathcal{D},\mathcal{F})}\left[\hat{\gamma}_{\mathcal{F}\cup\{(\mathbf{x},a,y)\}}(\tilde{\mathbf{x}})\right], \text{ where} \tag{5.2}$$

$$\hat{\gamma}_{\mathcal{F}\cup\{(\mathbf{x},a,y)\}}(\tilde{\mathbf{x}}) = \hat{p}\left(y_{a_{\tilde{\mathbf{x}}}} < y_{1-a_{\tilde{\mathbf{x}}}} \mid \tilde{\mathbf{x}}, \mathcal{D}, \mathcal{F} \cup \{(\mathbf{x}, a, y)\}\right) \tag{5.3}$$

is the estimated Type S error rate after seeing feedback $(\mathbf{x}, a, y)$, and $a_{\tilde{\mathbf{x}}}$ is the predicted best action for $\tilde{\mathbf{x}}$: $a_{\tilde{\mathbf{x}}} = \arg\max_{a'} \mathbb{E}[y_{a'} \mid \tilde{\mathbf{x}}, \mathcal{D}, \mathcal{F} \cup \{(\mathbf{x}, a, y)\}]$.

In practice, directly minimizing the error rate in eq. (5.2) focuses too much on exploitation and, therefore, Publication III modifies the criterion to add exploration in a heuristic way. Section 5.2.2 presents a more principled method for better exploration.

Note that this method cannot as such be easily extended to multiple actions. Even if the outcomes are independent from each other given $\mathbf{x}$, the comparisons of the expected utilities do not factorize into a product of pairwise comparisons. An example to give intuition is to consider three

independent Gaussian random variables X,Y and Z, for which it is easy to show that $P(X > Y \mid X > Z) \neq P(X > Y)$. The extension to multiple actions in Publication IV is presented in Section 5.2.2.

Outcome models

The potential outcomes are modelled using Gaussian Processes (GP) [77], so that each outcome has a GP prior $f_a \sim GP(0, k)$ with a squared exponential kernel $k(x, x')$ and mixed noise likelihood to account for different noise in the observed data $\mathcal{D}$ and in the answers about counterfactuals $\mathcal{F}$. The likelihood is

$$y \mid \mathbf{x}, a, \mathcal{D}, \mathcal{F} \sim N(f_{a,x}, \sigma_a^2), \text{if}(\mathbf{x}, a, y) \in \mathcal{D}, \text{and}$$
$$y \mid \mathbf{x}, a, \mathcal{D}, \mathcal{F} \sim N(f_{a,x}, \sigma_e^2), \text{if}(\mathbf{x}, a, y) \in \mathcal{F} \tag{5.4}$$

where $f_{a,\mathbf{x}}$ is the expected outcome of the GP at $\mathbf{x}$, fitted with $\mathcal{D}$ and $\mathcal{F}$. An unknown noise variance of the observation process of outcome a is σ_a^2, and σ_e^2 is an unknown noise in the expert's answers. If systematic bias of the experts is known or can be modeled, it could be added to eq. (5.4) to account for the bias.

Results

Publication III derives conditions under which imbalance increases the risk of wrong decisions in a down-the-line decision-making task in the setup described above. Publication III showed experimentally that when the method is applied to counterfactual elicitation with a simulated expert, imbalance decreases and the proportion of correct decision increases with increasing feedback, assuming that the expert's answers are not biased but may contain independent noise.

5.2.2 Deciding Between Multiple Actions

Publication IV extends the method presented in Section 5.2.1 to decision-making between multiple actions $a = 1, \ldots, K$ based on K predicted outcomes $p(y_a \mid x, D)$. Instead of Type S error, the insight is to consider the optimal decision D_{best} as the parameter of interest for experimental design, which also fixes the exploration problem of the previous method.

In order to model the D_{best}, we characterize the posterior probability π_a

of a being the optimal decision at $\tilde{\mathbf{x}}$:

$$\pi_a = P\left(f_{a,\tilde{\mathbf{x}}} = \max_{a'} f_{a',\tilde{\mathbf{x}}}\right), \tag{5.5}$$

$$= P\left(\bigcap_{a' \neq a} \{f_{a,\tilde{\mathbf{x}}} > f_{a',\tilde{\mathbf{x}}}\}\right), \tag{5.6}$$

The connection to Type S error rate is that for each a, quantity $1 - \pi_a$ tells the estimated error rate *without aleatoric uncertainty*, that is, without accounting for the inherent stochasticity in the outcomes. This focuses the criterion to reduce epistemic uncertainty.

The active learning criterion coined D-EIG then maximizes the information gain to $p(D_{\text{best}}(\tilde{\mathbf{x}}) \mid \mathcal{D})$, which is a categorical distribution with parameters π_a, $a = 1, \ldots, K$. Here, we use the result that maximizing the expected information gain of model parameter is equivalent to minimizing its expected entropy after acquiring the query, as discussed in Chapter 3.3. The acquisition is then

$$(\mathbf{x}^\star, a^\star) = \underset{(\mathbf{x}_j, a_j) \in U}{\arg\min} \quad \mathbb{E}_{p(y_{a_j}|\mathbf{x}_j,\mathcal{D})}\left[\sum_i^{n_t} \mathcal{H}[p(D_{\text{best}}(\tilde{\mathbf{x}}_i)|\mathcal{D} \cup \{(\mathbf{x}_j, a_j, y_{a_j})\})]\right], \tag{5.7}$$

where $\{\tilde{\mathbf{x}}_i\}_{i=1}^{n_t}$ is the test population, or in case of targeted decision-making, just a single $\tilde{\mathbf{x}}$. If the test input is not known then the criterion can be approximated using $\mathcal{D}$ or U.

The problem setting of identifying the best action for an individual $\tilde{\mathbf{x}}$ is close to Bayesian optimization, but the difference is that the queries are indirect: Bayesian optimization would seek to find a query that maximizes the utility, ending up with $(a, \mathbf{x})$ with the highest outcome. Instead, D-EIG method search for query that provides maximum information about D_{best} at $\tilde{\mathbf{x}}$, without querying $\tilde{\mathbf{x}}$.

Inference and approximations

The expectation in eq. (5.7) for a single $\tilde{\mathbf{x}}$ is approximated using Monte Carlo estimation with N_s samples as

$$\mathbb{E}_{p(y_{a_j}|\mathbf{x}_j,\mathcal{D})}\left[\mathcal{H}[p(D_{\text{best}}(\tilde{\mathbf{x}})|\mathcal{D} \cup \{(\mathbf{x}_j, a_j, y_{a_j})\})]\right]$$

$$\simeq \frac{1}{N_s} \sum_{l=1}^{N_s} \mathcal{H}[p(D_{\text{best}}(\tilde{\mathbf{x}})|\mathcal{D} \cup \{(\mathbf{x}_j, a_j, y_{a_j}^{(l)})\})]. \tag{5.8}$$

where $y_{a_j}^{(l)}$ are drawn from the posterior $p(y_{a_j}|\mathbf{x}_j, \mathcal{D})$.

Results

Publication IV evaluates the proposed decision-making aware active learning in three scenarios: synthetic data, and two semi-synthetic medical data (Infant Health and Development Program (IHDP) dataset and Osteoarthritis Initiative (OAI) dataset) for treatment recommendation. The results in Figure 5.2 compare different goal-oriented active learning methods in case of deciding for a single test input $\tilde{\mathbf{x}}$. The targeted criterion from Publication II (see Section 5.1) is T-EIG, which illustrates that optimizing the queries for personalized prediction does not directly imply the best decision-making performance. Insted, the decision-making aware active learning D-EIG achieves higher performance both in terms of increasing the proportion of correct decisions (Figure 5.2 a-b) as well as decreasing the entropy of D_{best} (Figure 5.2 c-d).

5.2.3 Summary and Limitations

The aim of the method is to bring an interaction component to a decision-support system, to allow fine-tuning and correcting model predictions for better decision-making. The method assumes that the human expert makes the decision by following Bayesian decision theory, by computing the expected utility with respect to the posterior probability of the outcomes. In reality, people are known not to make decisions in such a simple manner even if the utility is simple to define [78]. Furthermore, they may consider and trade-off multiple factors, have biases and behave optimally only under computational rationality assumption [79, 80]. Section 6 presents one way to learn about expert's unknown multiparameter optimization goal. To account for biases and bounded rationality, a cognitive user-model with computational rationality assumptions could provide a solution, if it can be identified from the interaction [67, 81].

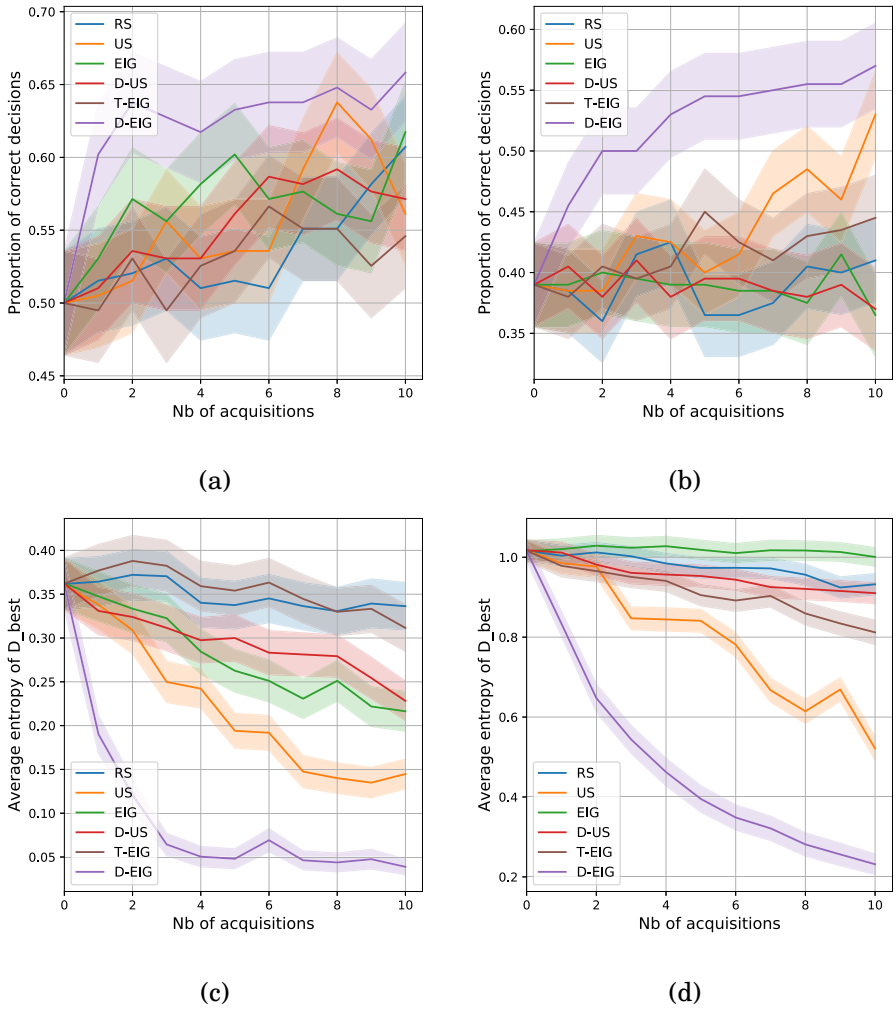

Figure 5.2. Performance of decision-making active learning with respect to the number of queries (x-axis). (a,c) IHDP dataset. (b,d) OAI dataset. The plots show mean and SEM (standard error of the mean) of (a-b) the accuracy score and (c-d) the entropy of the posterior of the optimal decision, in 200 replications. The proposed acquisition method `D-EIG` is able to improve the accuracy of the decision-making even with very few queries, and it also reduces the uncertainty about the optimal decision the most. RS = random sampling, US = uncertainty sampling, D-US = decision uncertainty sampling, T-EIG = targeted expected information gain, D-EIG = decision-making aware active learning. Reprinted with permission from Publication IV.

6. Human-in-the-Loop Assisted Molecular Design

This chapter studies another application field for human-in-the-loop learning: AI assisted molecular design. The research question (RQ3) is what kind of a system will enable a chemist to interact with a *de novo* molecular design tool to improve its performance, measured by how well the molecules it generates match their goal.

De novo molecular design aims at automating generation and exploration of novel molecules to find candidates for new drugs. Even though state-of-the-art *de novo* design tools such as Reinvent [82] can generate valid novel molecules, they still require a human to specify the goal, which consists of several objectives. A common approach is to scalarize this multi-parameter optimization (MPO) problem by combining the (weighted) objectives into a single score [82, 83, 84]. Building this scoring function requires detailed manual work of specifying the objectives, their utility functions and weights. As a result, a chemist often tunes the scoring function iteratively using trial and error. Using human-in-the-loop methods for building effective scoring functions has been identified as an important future direction in molecular design [85]. To this end, Publication V applies interactive machine learning to this new, promising and yet unexplored application domain.

In contrast to the current practice, this thesis proposes adapting the scoring function interactively based on a chemist's feedback, to make it easier for the chemist to specify their goal without requiring them to explicate it mathematically. As a result, the design tool will be able to generate molecules that better match the chemist's goal.

A solution proposed in this thesis consists of a probabilistic model of

the chemist's goal (i.e. a user-model), which acts as the scoring function, an active query selection strategy for eliciting the goal, and a graphical user interface to show candidate molecules to the chemist and let the chemist to give feedback to them. Publication V uses a publicly available *de novo* design tool Reinvent [82] to generate molecules given the scoring function, and empirically shows a proof-of-principle of such a system. From another point of view, the proposed system is an example of a simple reward elicitation task, because the design tool uses the model of the chemist's goal as reward in reinforcement learning, when it is training to generate molecules with desired properties.

Two user-models are proposed in Publication V: a parametric model for adapting the parameters of a scalarized MPO objective function, and a non-parametric model (GP) to directly represent the chemist's goal. The latter approach is similar to that presented in [84], but adds to it an active selection of which molecules are shown to the chemist to maximize the effectiveness of the queries. The most effective query selection strategy was found to be based on Thompson sampling [39, 33], because it balances exploration with exploitation and is therefore suitable for eliciting the goal for molecule generation.

Publication V evaluates the methods in two molecular design tasks with a simulated chemist: for optimizing the druglikeness of a molecule, and for generating molecules that are likely to bind to the dopamine receptor D_2 (DRD2). The results show that by giving feedback to just 50-200 molecules, a simulated chemist achieves meaningful improvement compared to baselines. A pilot demonstration with a medicinal chemist interacting with the system also shows promising results, where 100 evaluations of molecules is sufficient to improve the output score on average twofold compared to no-feedback case.

7. Discussion and Conclusion

As a whole, this thesis set out to study how interaction with a human expert can benefit challenging machine learning tasks with little data in medical applications. The results presented in this thesis show that collaboration between a human and a machine learning model indeed improves performance of the model beyond to what is achievable from data alone, and especially when the interaction is planned by a query selection algorithm.

The research questions studied within the thesis were

- **RQ1:** Can human-in-the-loop approach to knowledge elicitation solve the small n large p regularization problem through eliciting indirect and qualitative expert knowledge, and through selecting the most informative queries to include domain knowledge interactively?
- **RQ2:** How to select the interaction so that it is the most beneficial to a prospective task where the model is to be used: predicting outcome of a new individual, or making a decision of which action to take based on model's prediction of outcome under alternative actions?
- **RQ3:** How to apply human-in-the-loop methods to *de novo* molecular design, to actively infer the goal of an expert from interaction for adapting a *de novo* design tool's outcome to better match the expert's objectives?

For the first research question, two knowledge elicitation methods were developed and experimentally validated in Publications I and II. Both methods use a user-model to anticipate the expert's feedback and model their knowledge. The first method has a separate model for the expert's

knowledge, and the second takes advantage of a joint model of the expert and the data generation process. In both cases, the humans provide qualitative feedback about the relevance or effect of predictors in the model. The results show great potential in mitigating the small n large p problem by using human knowledge, while the query selection strategies make the workload feasible to the experts.

One advantage of the second approach is that the experimental design query selection strategy takes into account both how the expert gives feedback and how the feedback in turn affects the predictive model. These methods have, however, high computational cost, and therefore a separate and simplified user-model, as in Publication I, may be more suitable for online interaction. On the other hand, in that case the connection between the user-model and the prediction model is not well defined, and as a result, the feedback is integrated to the prediction model heuristically.

For the second research question, two new goal-oriented active learning methods were proposed for the tasks of personalizing knowledge elicitation to an individual, and for improving decision-making performance of the model. A key insight is to identify the utility that is suitable for each tasks. After defining the utility and a joint model as described above, queries can be selected in a well-established way by applying Bayesian experimental design to the particular utility.

The third research question is addressed empirically in Publication V in two molecule design tasks, showing for the first time how molecules generated by an AI tool get closer to the chemist's goal when the chemist is allowed to interact with the tool. The method models the goal of the chemist and adapts the objective function in a molecular design tool accordingly. The first main insight is a separate user-model that has the same structure and parameters as the objective function, but its parameters are fitted based on the chemist's feedback instead of requiring the chemist to set them manually. The user-model is then used as the objective function in the molecule generation, resulting in improved performance. The second insight is that the most efficient query selection strategies are based on methods that aim at maximizing the reward of the queries, that is, bandit algorithms or Bayesian optimization algorithms.

A recent survey [86] identified lack of principled methods for knowledge elicitation as a potential problem in applied machine learning. The methods developed within this thesis allow principled human-in-the-loop learning, where the interaction is optimized by a machine learning model that takes into account the expert's behaviour. Future work should merge these new insights within the well-established elicitation protocols that aim at minimizing cognitive biases by combining knowledge of multiple experts [10].

In precision medicine, the goal-oriented active learning methods can ease the workload of experts in two tasks that are important in data-driven medicine: for personalizing queries to a patient, and for concentrating queries to those that make a difference in decision-making tasks of selecting treatment to a patient. They will be useful for example in cases where a clinician wishes to get expert opinions to help decide a treatment for a new patient for which the current predictions of a decision-support model are not reliable.

When applying these methods to precision medicine, ethical issues must be considered if the methods are used for other purposes than knowledge elicitation, e.g. for designing new experiments. Even in knowledge elicitation, the question of privacy is important, to prevent a domain expert from accidentally exposing e.g. private information about patients to the algorithm.

The work in this thesis provides first proof-of-principle studies about how to apply interactive knowledge elicitation, therefore, some simplifying choices were made that limit the generality of the methods. Most importantly, the developed methods rely on the expert's guesses, and a fair question is, can the experts provide answers to the queries. The demonstrations with human experts' feedback show that at least in some cases it is possible, and even easy. Furthermore, the methods developed in this thesis do not assume the experts' answers to be always correct. Overconfidence, for example, is not a critical issue for these methods, since the probabilistic framework allows leveraging the expert knowledge if the experts are on average right (i.e. unbiased) even if they exaggerate the magnitude of their answer. Nevertheless, future work should consider modeling the biases in experts' answers in the user-models. I believe that advancements in

computational cognitive models for capturing human biases could unlock the potential of knowledge elicitation in full.

Another important extension to the method is batch query strategies for knowledge elicitation, since small batches of e.g. ten queries may be even easier for an expert to answer than single queries. Showing queries to the expert in batches, however, raises new modeling questions: the experts are likely to compare the queries within the batch, which may either hinder or increase performance of the method, depending on whether the model factors in this behaviour.

A final point I would like to raise is the distinction of the expert and a user of an AI system - and emphasize the case when they are the same person. In this thesis, the term expert is used for the person whose knowledge is being elicited to improve a machine learning model, and a user refers to an end-user of the said model, which can be a part of a larger AI system. The distinction has been made to emphasize that domain expertise is crucial for these methods to be useful, and that only experts should be allowed to use some of the systems. However, an overarching motivation in my research has been to develop methods for eliciting knowledge online from a user while they are using the system. Then, in case the user of a decision support system is also an expert in that domain, the interactive methods will enable on-demand knowledge elicitation and allow a deployed model to learn continuously. As a result, the methods developed in this thesis will help the experts to collaborate with their AI assistant to reach better results than neither of them could alone.

References

[1] Andreas Holzinger. Interactive machine learning for health informatics: when do we need the human-in-the-loop? *Brain Informatics*, 3(2):119–131, 2016.

[2] Andreas Holzinger, Markus Plass, Michael Kickmeier-Rust, Katharina Holzinger, Gloria Cerasela Crişan, Camelia-M Pintea, and Vasile Palade. Interactive machine learning: experimental evidence for the human in the algorithmic loop. *Applied Intelligence*, 49(7):2401–2414, 2019.

[3] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[4] David Barber. *Bayesian reasoning and machine learning*. Cambridge University Press, 2012.

[5] Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. *Bayesian Data Analysis*. CRC Press, 2013.

[6] Rens van de Schoot, Sarah Depaoli, Ruth King, Bianca Kramer, Kaspar Märtens, Mahlet G Tadesse, Marina Vannucci, Andrew Gelman, Duco Veen, Joukje Willemsen, et al. Bayesian statistics and modelling. *Nature Reviews Methods Primers*, 1(1):1–26, 2021.

[7] Anthony O'Hagan, Caitlin E. Buck, Alireza Daneshkhah, J. Richard Eiser, Paul H. Garthwaite, David J. Jenkinson, Jeremy E. Oakley, and Tim Rakow. *Uncertain Judgements: Eliciting Experts' Probabilities*. Wiley, Chichester, England, 2006.

[8] Petrus Mikkola, Osvaldo A Martin, Suyog Chandramouli, Marcelo Hartmann, Oriol Abril Pla, Owen Thomas, Henri Pesonen, Jukka Corander, Aki Vehtari, Samuel Kaski, et al. Prior knowledge elicitation: The past, present, and future. *arXiv preprint arXiv:2112.01380*, 2021.

[9] Robert T Clemen and Robert L Winkler. Combining probability distributions from experts in risk analysis. *Risk analysis*, 19(2):187–203, 1999.

[10] Anthony O'Hagan. Expert knowledge elicitation: subjective but scientific. *The American Statistician*, 73(sup1):69–81, 2019.

[11] Robert L Winkler. The assessment of prior distributions in bayesian analysis. *Journal of the American Statistical association*, 62(319):776–800, 1967.

[12] Samantha Low Choy, Rebecca O'Leary, and Kerrie Mengersen. Elicitation by design in ecology: using expert opinion to inform priors for bayesian statistical models. *Ecology*, 90(1):265–277, 2009.

[13] K. Coussement, D.F. Benoit, and M. Antioco. A bayesian approach for incorporating expert opinions into decision support systems: A case study of online consumer-satisfaction detection. *Decision Support Systems*, 79:24–32, 2015.

[14] Robert Denham and Kerrie Mengersen. Geographically assisted elicitation of expert opinion for regression models. *Bayesian Analysis*, 2(1):99–135, 2007.

[15] Joseph B. Kadane, James M. Dickey, Robert L. Winkler, Wayne S. Smith, and Stephen C. Peters. Interactive elicitation of opinion for a normal linear model. *Journal of the American Statistical Association*, 75(372):845–854, 1980.

[16] Paul H. Garthwaite, Shafeeqah A. Al-Awadhi, Fadlalla G. Elfadaly, and David J. Jenkinson. Prior distribution elicitation for generalized linear and piecewise-linear models. *Journal of Applied Statistics*, 40(1):59–75, 2013.

[17] Geoffrey Jones and Wesley O. Johnson. Prior elicitation: Interactive spreadsheet graphics with sliders can be fun, and informative. *The American Statistician*, 68(1):42–51, 2014.

[18] David E Morris, Jeremy E Oakley, and John A Crowe. A web-based tool for eliciting probability distributions from experts. *Environmental Modelling & Software*, 52:1–4, 2014.

[19] Justine V Murray, Anne W Goldizen, Rebecca A O'Leary, Clive A McAlpine, Hugh P Possingham, and Samantha Low Choy. How useful is expert opinion for predicting the distribution of a species within and beyond the region of expertise? a case study using brush-tailed rock-wallabies petrogale penicillata. *Journal of Applied Ecology*, 46(4):842–851, 2009.

[20] Toby J. Mitchell and John J. Beauchamp. Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, 83(404):1023–1032, 1988.

[21] Edward I. George and Robert E. McCulloch. Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, 88(423):881–889, 1993.

[22] Carlos M Carvalho, Nicholas G Polson, and James G Scott. The horseshoe estimator for sparse signals. *Biometrika*, 97(2):465–480, 2010.

[23] Juho Piironen and Aki Vehtari. Sparsity information and regularization in the horseshoe and other shrinkage priors. *Electronic Journal of Statistics*, 11(2):5018–5051, 2017.

[24] Burr Settles. *Active learning*. Morgan & Claypool, 2012.

[25] Samuel Budd, Emma C Robinson, and Bernhard Kainz. A survey on active learning and human-in-the-loop deep learning for medical image analysis. *Medical Image Analysis*, page 102062, 2021.

[26] David D Lewis and Jason Catlett. Heterogeneous uncertainty sampling for supervised learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 148–156, 1994.

[27] Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.

[28] Ido Dagan and Sean P. Engelson. Committee-based sampling for training probabilistic classifiers. In *In Proceedings of the Twelfth International Conference on Machine Learning*, pages 150–157. Morgan Kaufmann, 1995.

[29] Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Máté Lengyel. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*, 2011.

[30] Kirthevasan Kandasamy, Willie Neiswanger, Reed Zhang, Akshay Krishnamurthy, Jeff Schneider, and Barnabas Poczos. Myopic posterior sampling for adaptive goal oriented design of experiments. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 3222–3232, 2019.

[31] Minjie Xu and Gary Kazantsev. Understanding goal-oriented active learning via influence functions. In *Proceedings of the NeurIPS workshop on Machine Learning with Guarantees*, 2019.

[32] Maytal Saar-Tsechansky and Foster Provost. Decision-centric active learning of binary-outcome models. *Information systems research*, 18(1):4–22, 2007.

[33] Daniel J Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, and Zheng Wen. A tutorial on thompson sampling. *Foundations and Trends® in Machine Learning*, 11(1):1–96, 2018.

[34] Shervin Javdani, Yuxin Chen, Amin Karbasi, Andreas Krause, Drew Bagnell, and Siddhartha Srinivasa. Near optimal bayesian active learning for decision making. In *Artificial Intelligence and Statistics*, pages 430–438. PMLR, 2014.

[35] Guang Zhao, Edward Dougherty, Byung-Jun Yoon, Francis Alexander, and Xiaoning Qian. Uncertainty-aware active learning for optimal bayesian classifier. In *International Conference on Learning Representations (ICLR)*, 2021.

[36] Guang Zhao, Edward Dougherty, Byung-Jun Yoon, Francis J. Alexander, and Xiaoning Qian. Bayesian active learning by soft mean objective cost of uncertainty. In Arindam Banerjee and Kenji Fukumizu, editors, *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 3970–3978. PMLR, 13–15 Apr 2021.

[37] Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.

[38] Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3:397–422, 2003.

[39] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.

[40] D. Jones, Matthias Schonlau, and W. J. Welch. Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization*, 13:455–492, 1998.

[41] Eric Brochu, Vlad M Cora, and Nando De Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.

[42] José Miguel Hernández-Lobato, Matthew W Hoffman, and Zoubin Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.

[43] Zi Wang and Stefanie Jegelka. Max-value entropy search for efficient bayesian optimization. In *International Conference on Machine Learning*, pages 3627–3635. PMLR, 2017.

[44] Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: no regret and experimental design. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, pages 1015–1022, 2010.

[45] Ksenia Korovina, Sailun Xu, Kirthevasan Kandasamy, Willie Neiswanger, Barnabas Poczos, Jeff Schneider, and Eric Xing. Chembo: Bayesian optimization of small organic molecules with synthesizable recommendations. In *International Conference on Artificial Intelligence and Statistics*, pages 3393–3403. PMLR, 2020.

[46] Daniel Packwood. *Bayesian Optimization for Materials Science*. Springer, 2017.

[47] Yichi Zhang, Daniel W Apley, and Wei Chen. Bayesian optimization for materials design with mixed quantitative and qualitative variables. *Scientific reports*, 10(1):1–13, 2020.

[48] A Gilad Kusne, Heshan Yu, Changming Wu, Huairuo Zhang, Jason Hattrick-Simpers, Brian DeCost, Suchismita Sarker, Corey Oses, Cormac Toher, Stefano Curtarolo, et al. On-the-fly closed-loop materials discovery via bayesian active learning. *Nature communications*, 11(1):1–11, 2020.

[49] Petrus Mikkola, Milica Todorović, Jari Järvi, Patrick Rinke, and Samuel Kaski. Projective preferential bayesian optimization. In *International Conference on Machine Learning*, pages 6884–6892. PMLR, 2020.

[50] Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, pages 273–304, 1995.

[51] Dennis V Lindley. On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, pages 986–1005, 1956.

[52] Pedram Daee, Tomi Peltola, Marta Soare, and Samuel Kaski. Knowledge elicitation via sequential probabilistic inference for high-dimensional prediction. *Machine Learning*, 106(9):1599–1620, Oct 2017.

[53] David JC MacKay. Information-based objective functions for active data selection. *Neural computation*, 4(4):590–604, 1992.

[54] Kasra Arnavaz, Aasa Feragen, Oswin Krause, and Marco Loog. Bayesian active learning for maximal information gain on model parameters. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 10524–10531. IEEE, 2021.

[55] Hema Raghavan, Omid Madani, and Rosie Jones. Active learning with feedback on features and instances. *Journal of Machine Learning Research*, 7(61):1655–1686, 2006.

[56] Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9(2), 2008.

[57] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *International Conference on Machine Learning*, pages 1183–1192. PMLR, 2017.

[58] Andreas Kirsch, Joost van Amersfoort, and Yarin Gal. Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

[59] Daniel Russo and Benjamin Van Roy. Learning to optimize via information-directed sampling. *Advances in Neural Information Processing Systems*, 27:1583–1591, 2014.

[60] Burr Settles. Closing the loop: Fast, interactive semi-supervised annotation with queries on features and instances. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1467–1478, 2011.

[61] Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. "Why should I trust you?": Explaining the predictions of any classifier. In *Proc of ACM SIGKDD*, 2016.

[62] Michael Brooks, Saleema Amershi, Bongshin Lee, Steven M Drucker, Ashish Kapoor, and Patrice Simard. FeatureInsight: Visual support for error-driven feature ideation in text classification. In *Proc of IEEE VAST*, pages 105–112, 2015.

[63] Zhengdong Lu and Todd K. Leen. Semi-supervised clustering with pairwise constraints: A discriminative approach. In *Proc of AISTATS*, pages 299–306, 2007.

[64] Homayun Afrabandpey, Tomi Peltola, and Samuel Kaski. Human-in-the-loop active covariance learning for improving prediction in small data sets. In *International Joint Conference on Artificial Intelligence*, pages 1959–1966, 2019.

[65] Marta Soare, Muhammad Ammad-Ud-Din, and Samuel Kaski. Regression with n$\rightarrow$ 1 by expert knowledge elicitation. In *2016 15th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 734–739. IEEE, 2016.

[66] Alvaro HC Correia and Freddy Lecue. Human-in-the-loop feature selection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 2438–2445, 2019.

[67] Pedram Daee, Tomi Peltola, Aki Vehtari, and Samuel Kaski. User modelling for avoiding overfitting in interactive knowledge elicitation for prediction. In *23rd International Conference on Intelligent User Interfaces*, pages 305–310, 2018.

[68] Andrew G Wilson, Christoph Dann, Chris Lucas, and Eric P Xing. The human kernel. *Advances in Neural Information Processing Systems*, 28:2854–2862, 2015.

[69] Bob Carpenter, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A probabilistic programming language. *Journal of statistical software*, 76(1):1–32, 2017.

[70] Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 208–214. JMLR Workshop and Conference Proceedings, 2011.

[71] Thomas P. Minka and John Lafferty. Expectation-propagation for the generative aspect model. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*, pages 352–359, 2002.

[72] Matthias W Seeger. Bayesian inference and optimal design for the sparse linear model. *Journal of Machine Learning Research*, 9:759–813, 2008.

[73] D.B. Rubin. Bayesian inference for causal effects: The role of randomization. *The Annals of statistics*, pages 34–58, 1978.

[74] Ahmed M Alaa and Mihaela van der Schaar. Bayesian inference of individualized treatment effects using multi-task gaussian processes. *arXiv preprint arXiv:1704.02801*, 2017.

[75] Oskar Morgenstern and John Von Neumann. *Theory of games and economic behavior*. Princeton university press, 1953.

[76] Andrew Gelman and Francis Tuerlinckx. Type s error rates for classical and bayesian single and multiple comparison procedures. *Computational statistics*, 15(3):373–390, 2000.

[77] C.E. Rasmussen and C.K.I. Williams. *Gaussian processes for machine learning*. the MIT Press, 2006.

[78] Mark J Machina. Choice under uncertainty: Problems solved and unsolved. *Journal of Economic Perspectives*, 1(1):121–154, 1987.

[79] Richard L Lewis, Andrew Howes, and Satinder Singh. Computational rationality: Linking mechanism and behavior through bounded utility maximization. *Topics in cognitive science*, 6(2):279–311, 2014.

[80] Samuel J Gershman, Eric J Horvitz, and Joshua B Tenenbaum. Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, 349(6245):273–278, 2015.

[81] Antti Kangasrääsiö, Jussi PP Jokinen, Antti Oulasvirta, Andrew Howes, and Samuel Kaski. Parameter inference for computational cognitive models with approximate bayesian computation. *Cognitive science*, 43(6):e12738, 2019.

[82] Thomas Blaschke, Josep Arús-Pous, Hongming Chen, Christian Margreitter, Christian Tyrchan, Ola Engkvist, Kostas Papadopoulos, and Atanas Patronov. Reinvent 2.0: an ai tool for de novo drug design. *Journal of chemical information and modeling*, 60(12):5918–5922, 2020.

[83] Christoph Grebner, Hans Matter, Alleyn T Plowright, and Gerhard Hessler. Automated de novo design in medicinal chemistry: which types of chemistry does a generative neural network learn? *Journal of Medicinal Chemistry*, 63(16):8809–8823, 2020.

[84] Robin Winter, Joren Retel, Frank Noé, Djork-Arné Clevert, and Andreas Steffen. grünifai: interactive multiparameter optimization of molecules in a continuous vector space. *Bioinformatics*, 36(13):4093–4094, 05 2020.

[85] Joshua Meyers, Benedek Fabian, and Nathan Brown. De novo molecular design and generative models. *Drug Discovery Today*, 26(11):2707–2715, 2021.

[86] Daniel Kerrigan, Jessica Hullman, and Enrico Bertini. A survey of domain knowledge elicitation in applied machine learning. *Multimodal Technologies and Interaction*, 5(12):73, 2021.

Aalto-DT 78/2023

ISBN 978-952-64-1284-9 (printed)
ISBN 978-952-64-1285-6 (pdf)
ISSN 1799-4934 (printed)
ISSN 1799-4942 (pdf)

Aalto University
School of Science
Department of Computer Science
www.aalto.fi

BUSINESS + ECONOMY

ART + DESIGN + ARCHITECTURE

SCIENCE + TECHNOLOGY

CROSSOVER

DOCTORAL THESES