Preface: The 1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval (AMIR)

Joeran Beel[1] and Lars Kotthoff[2]

[1] Trinity College Dublin – School of Computer Science & Statistics – Artificial Intelligence Discipline – ADAPT Centre – Ireland
[2] University of Wyoming – Department of Computer Science – Meta-Algorithmics, Learning and Large-scale Empirical Testing Lab – USA

beelj@tcd.ie | larsko@uwyo.edu

Abstract. Algorithm selection is a key challenge for most, if not all, computational problems. Typically, there are several potential algorithms that can solve a problem, but which algorithm would perform best (e.g. in terms of runtime or accuracy) is often unclear. In many domains, particularly artificial intelligence, the algorithm selection problem is well-studied, and various approaches and tools exist to tackle it in practice. Especially through meta-learning, impressive performance improvements have been achieved. The information retrieval (IR) community, however, has paid relatively little attention to the algorithm selection problem. The *1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval* (AMIR) brought together researchers from the IR community as well as from the machine learning (ML) and meta-learning community. Our goal was to raise the awareness in the IR community of the algorithm selection problem; identify the potential for automatic algorithm selection in information retrieval; and explore possible solutions for this context. AMIR was co-located with the *41st European Conference on Information Retrieval* (ECIR) in Cologne, Germany, and held on the 14th of April 2019. Out of ten submissions, five (50%) were accepted at AMIR, and an estimated 25 researchers attended the workshop.

1 Motivation[1]

There is a plethora of algorithms for information retrieval applications, such as search engines and recommender systems. There are about 100 approaches to recommend research papers alone [1]. The question that researchers and practitioners alike are faced with is which one of these approaches to choose for their particular problem. This is a

[1] Major text passages in this and the following sections are taken from J. Beel and L. Kotthoff, “Proposal for the 1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval (AMIR),” in Proceedings of the 41st European Conference on Information Retrieval (ECIR), 2019, vol. 11438, pp. 383–388. DOI 10.1007/978-3-030-15719-7_53.

The 1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval (AMIR), 14 April 2019, Cologne, Germany. Editors: Joeran Beel and Lars Kotthoff. Co-located with the 41st European Conference on Information Retrieval (ECIR). http://amir-workshop.org/

difficult choice even for experts, compounded by ongoing research that develops ever more approaches.

The challenge of identifying the best algorithm for a given application is not new. The so-called "algorithm selection problem" was first mentioned in the 1970s [2] and has attracted significant attention in various disciplines since then, especially in the last decade. Particularly in artificial intelligence, impressive performance achievements have been enabled by algorithm selection systems. A prominent example is the award-winning SATzilla system [3].

More generally, algorithm selection is an example of meta-learning, where the experience gained from solving problems informs how to solve future problems. Meta-learning and automating modelling processes has gained significant traction in the machine learning community, in particular with so-called AutoML approaches that aim to automate the entire machine learning and data mining process from ingesting the data to making predictions. An example of such a system is Auto-WEKA [4]. There have also been multiple competitions [5], [6] and workshops, symposia and tutorials [7–11] including a Dagstuhl seminar [9]. The OpenML platform was developed to facilitate the exchange of data and machine learning models to enable research into meta-learning [12].

Despite the significance of the algorithm selection problem and notable advances in solving it in many domains, the information retrieval community has paid relatively little attention to it. There are a few papers that investigate the algorithm selection problem in the context of information retrieval, for example in the field of recommender systems [13–21]. Also, the field of query performance prediction (QPP) has investigated how to predict algorithm performance in information retrieval [22], [23]. However, the number of researchers interested in this topic is limited and results so far have been not as impressive as in other domains.

There is potential for applying IR techniques in meta-learning as well. The algorithm selection problem can be seen as a traditional information retrieval task, i.e. the task of identifying the most relevant item (an algorithm) from a large corpus (thousands of potential algorithms and parameters) for a given information need (e.g. classifying photos or making recommendations). We see great potential for the information retrieval community contributing to solving the algorithm selection problem.

2 The 1st AMIR Workshop

The 1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval (AMIR)[2] was accepted to be held at the 41st European Conference on Information Retrieval (ECIR) in Cologne, Germany, on the 14th of April 2019 [24]. AMIR aimed at achieving the following goals:

- Raise awareness in the information retrieval community of the algorithm selection problem.

[2] http://amir-workshop.org/

- Identify the potential for automated algorithm selection and meta learning in IR applications.
- Familiarize the IR community with algorithm selection and meta-learning tools and research that has been published in related disciplines such as machine learning.
- Find solutions to address and solve the algorithm selection problem in IR.

Topics of interest to AMIR included:

- Algorithm Configuration
- Algorithm Selection
- Algorithm Selection as User Modeling Task
- Auto* Tools in Practice (e.g. AutoWeka, AutoKeras, librec-auto, auto-sklearn, AutoTensorFlow, …)
- Automated A/B Tests (AutoA/B)
- Automated Evaluations (AutoEval)
- Automated Information Retrieval (AutoIR)
- Automated Machine Learning / Automatic Machine Learning / AutoML
- Automated Natural Language Processing (AutoNLP)
- Automated Recommender Systems (AutoRecSys)
- Automated User Modelling (AutoUM)
- Benchmarking
- CASH Problem (Combined Algorithm Selection and Hyper Parameter Optimization)
- Evaluation Methods and Metrics
- Evolutionary Algorithms
- Hyper-Parameter Optimization and Tuning
- Learning to Learn
- Meta-Heuristics
- Meta-Learning
- Neural Network Architecture Search / Neural Architecture Search (NAS) / Neural Network Search
- Recommender Systems for Algorithms
- Search Engines for Algorithms
- Transfer Learning, Few-Shot Learning, One-Shot Learning, …

Our vision is to establish a regular workshop at ECIR or related venues (e.g. SIGIR, UMAP, RecSys) and eventually – in the long run – solve the algorithm selection problem in information retrieval. We hope to stimulate collaborations between researchers in IR and meta-learning through presentations and discussions at the workshop, which will ultimately lead to joint publications and research proposals.

3 Accepted Papers

We received a total of ten submissions, of which the following five (50%) were accepted to be presented at the workshop [25–29]:

3.1 Algorithm selection with librec-auto

Masoud Mansoury and Robin Burke

Due to the complexity of recommendation algorithms, experimentation on recommender systems has become a challenging task. Current recommendation algorithms, while powerful, involve large numbers of hyperparameters. Tuning hyperparameters for finding the best recommendation outcome often requires execution of large numbers of algorithmic experiments particularly when multiples evaluation metrics are considered. Existing recommender systems platforms fail to provide a basis for systematic experimentation of this type. In this paper, we describe librec-auto, a wrapper for the well-known LibRec library, which provides an environment that supports automated experimentation.

3.2 **Investigating Ad-Hoc Retrieval Method Selection with Features Inspired by IR Axioms**

Siddhant Arora and Andrew Yates

We consider the algorithm selection problem in the context of ad-hoc information retrieval. Given a query and a pair of retrieval methods, we propose a meta-learner that predicts how to combine the methods' relevance scores into an overall relevance score. These predictions are based on features inspired by IR axioms that quantify properties of the query and its top rank documents. We conduct an evaluation on TREC benchmark data and find that the meta-learner often significantly improves over the individual methods in terms of both nDCG@20 and P@30. Finally, we conduct a feature weight analysis to investigate which features the meta-learner uses to make its decisions.

3.3 Augmenting the DonorsChoose.org Corpus for Meta-Learning

Gordian Edenhofer, Andrew Collins, Akiko Aizawa, and Joeran Beel

The DonorsChoose.org dataset of past donations provides a big and feature-rich corpus of users and items. The dataset matches donors to projects in which they might be interested in and hence is intrinsically about recommendations. Due to the availability of detailed item-, user- and transaction-features, this corpus represents a suitable candidate for meta-learning approaches to be tested. This study aims at providing an augmented corpus for further recommender systems studies to test and evaluate meta-learning approaches. In the augmentation, metadata of collaborative and content-based filtering techniques is amended to the corpus. It is further extended with aggregated statistics of users and transactions and an exemplary meta-learning experiment. The performance in the learning subsystem is measured via the recall of recommended items in a Top-N test set. The augmented dataset and the source code are released into the public domain at https://github.com/BeelGroup/Augmented-DonorsChoose.org-Dataset.

3.4 **RARD II: The 94 Million Related-Article Recommendation Dataset**

Joeran Beel, Barry Smyth and Andrew Collins

The main contribution of this paper is to introduce and describe a new recommender-systems dataset (RARD II). It is based on data from a recommender-system in the digital library and reference management software domain. As such, it complements datasets from other domains such as books, movies, and music. The RARD II dataset encompasses 94m recommendations, delivered in the two years from September 2016 to September 2018. The dataset covers an item-space of 24m unique items. RARD II provides a range of rich recommendation data, beyond conventional ratings. For example, in addition to the usual ratings matrices, RARD II includes the original recommendation logs, which provide a unique insight into many aspects of the algorithms that generated the recommendations. The recommendation logs enable researchers to conduct various analyses about a real-world recommender system. This includes the evaluation of meta-learning approaches for predicting algorithm performance. In this paper, we summarise the key features of this dataset release, describe how it was generated and discuss some of its unique features. Compared to its predecessor RARD, RARD II contains 64% more recommendations, 187% more features (algorithms, parameters, and statistics), 50% more clicks, 140% more documents, and one additional service partner (JabRef).

3.5 An Extensive Checklist for Building AutoML Systems

Thiloshon Nagarajah and Guhanathan Poravi

Automated Machine Learning is a research area which has gained a lot of focus in the recent past. But the required components to build an autoML system is neither properly documented nor very clear due to the differences and the recentness of researches. If the required steps are analyzed and brought under a common survey, it will assist in continuing researches. This paper presents an analysis of the components and technologies in the domains of autoML, hyperparameter tuning and meta learning and, presents a checklist of steps to follow while building an AutoML system. This paper is a part of an ongoing research and the findings presented will assist in developing a novel architecture for an autoML system.

4 Keynote and Hands-on Sessions

We were delighted to hear the keynote [30] from Marius Lindauer and having two hands-on sessions about automated algorithm selection tools [31], [32].

4.1 Automated Algorithm Selection: Predict which algorithm to use!

Marius Lindauer

To achieve state-of-the-art performance, it is often crucial to select a suitable algorithm for a given problem instance. For example, what is the best search algorithm for a given instance of a search problem; or what is the best machine learning algorithm for a given dataset? By trying out many different algorithms on many problem instances, developers learn an intuitive mapping from some characteristics of a given problem instance (e.g., the number of features of a dataset) to a well-performing algorithm (e.g., random forest). The goal of automated algorithm selection is to learn from data, how to automatically select a well-performing algorithm given such characteristics. In this talk, I will give an overview of the key ideas behind algorithm selection and different approaches addressing this problem by using machine learning.

4.2 Hands-on Session with ASlib
Lars Kotthoff

ASlib is a standard format for representing algorithm selection systems and a bechmark library with example problems from many different application domains. I will give an overview of what it is, example analyses available on its website, and the algorithm selection competitions 2015 and 2017 that were based on it. ASlib is available at http://www.aslib.net./

4.3 Hands-On Automated Machine Learning Tools: Auto-Sklearn and Auto-PyTorch
Marius Lindauer

To achieve state-of-the-art performance in machine learning (ML), it is very important to choose the right algorithm and its hyperparameters for a given dataset. Since finding the correct settings needs a lot of time and expert knowledge, we developed AutoML tools that can be used out-of-the-box with minimal expertise in machine learning. In this session, I will present two state-of-the-art tools in this field: (i) auto-sklearn (www.automl.org/auto-sklearn/) for classical machine learning and (ii) AutoPyTorch (www.automl.org/autopytorch/) for deep learning.

5 Organization

5.1 Organizers

Joeran Beel[3] is Assistant Professor in Intelligent Systems at the School of Computer Science and Statistics at Trinity College Dublin. He is also affiliated with the ADAPT Centre, an interdisciplinary research centre that closely cooperates with industry partners including Google, Deutsche Bank, Huawei, and Novartis. Joeran is further a Visiting Professor at the National Institute of Informatics (NII) in Tokyo. His research focuses on information retrieval, recommender systems, algorithm selection, user modelling and machine learning. He has developed novel algorithms in these fields and conducted research on the question of how to evaluate information retrieval systems. Joeran also has industry experience as a product manager and as the founder of three business start-ups he experienced the algorithm selection problem first hand. Joeran is serving as general co-chair of the 26th Irish Conference on Artificial Intelligence and Cognitive Science and served on program committees for major information retrieval venues including SIGIR, ECIR, UMAP, RecSys, and ACM TOIS.

Lars Kotthoff[4] is Assistant Professor at the University of Wyoming. He leads the Meta-Algorithmics, Learning and Large-scale Empirical Testing (MALLET) lab and has acquired more than $400K in external funding to date. Lars is also the PI for the Artificially Intelligent Manufacturing center (AIM) at the University of Wyoming. He

[3] https://www.scss.tcd.ie/joeran.beel/
[4] http://www.cs.uwyo.edu/~larsko/

co-organized multiple workshops on meta-learning and automatic machine learning (e.g. [11]) and the Algorithm Selection Competition Series [5]. He was workshop and masterclass chair at the CPAIOR 2014 conference and organized the ACP summer school on constraint programming in 2018. His research combines artificial intelligence and machine learning to build robust systems with state-of-the-art performance. Lars' more than 60 publications have garnered >1111 citations and his research has been supported by funding agencies and industry in various countries.

5.2 Programme Committee

- Akiko Aizawa, National Institute of Informatics, Tokyo
- Andreas Nürnberger, University of Magdeburg
- Andreas Weiler, ZHAW School of Engineering
- Corinna Breitinger, University of Konstanz
- Dietmar Jannach, University of Klagenfurt
- Douglas Leith, Trinity College Dublin
- Felix Beierle, Technical University of Berlin
- Felix Hamborg, University of Konstanz
- Heike Trautmann, University of Münster
- Johann Schaible, GESIS
- Katharina Eggensperger, University of Freiburg
- Marius Lindauer, University of Freiburg
- Mark Collier, University of Edinburgh
- Matthias Feurer, University of Freiburg
- Moritz Schubotz, University of Konstanz
- Nicola Ferro, University of Padua
- Owen Conlan, Trinity College Dublin
- Pascal Kerschke, University of Münster
- Pavel Brazdil, University of Porto
- Rob Brennan, Trinity College Dublin
- Roman Kern, Know-Center, Austria
- Tiago Cunha, University of Porto
- Vincent Wade, Trinity College Dublin
- Zeljko Carevic, GESIS

6 Acknowledgements

This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) under Grant Number 13/RC/2106 and funding from the European Union and Enterprise Ireland under Grant Number CF 2017 0303-1. Lars Kotthoff is supported by NSF award 1813537.

References

[1] J. Beel, B. Gipp, S. Langer, and C. Breitinger, "Research Paper Recommender Systems: A Literature Survey," *International Journal on Digital Libraries*, no. 4, pp. 305–338, 2016.

[2] J. R. Rice, "The algorithm selection problem," 1975.

[3] L. Xu, F. Hutter, H. H. Hoos, and K. Leyton-Brown, "SATzilla: portfolio-based algorithm selection for SAT," *Journal of artificial intelligence research*, vol. 32, pp. 565–606, 2008.

[4] L. Kotthoff, C. Thornton, H. H. Hoos, F. Hutter, and K. Leyton-Brown, "Auto-WEKA 2.0: Automatic model selection and hyperparameter optimization in WEKA," *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 826–830, 2017.

[5] M. Lindauer, J. N. van Rijn, and L. Kotthoff, "The Algorithm Selection Competition Series 2015-17," *arXiv preprint arXiv:1805.01214*, 2018.

[6] W.-W. Tu, "The 3rd AutoML Challenge: AutoML for Lifelong Machine Learning," in *NIPS 2018 Challenge*, 2018.

[7] P. Brazdil, "Metalearning & Algorithm Selection," *21st European Conference on Artificial Intelligence (ECAI)*, 2014.

[8] R. Calandra, F. Hutter, H. Larochelle, and S. Levine, "Workshop on Meta-Learning (MetaLearn 2017) @NIPS," in *http://metalearning.ml*, 2017.

[9] H. H. Hoos, F. Neumann, and H. Trautmann, "Automated Algorithm Selection and Configuration," *Report from Dagstuhl Seminar 16412*, vol. 6, no. 11, 2016.

[10] R. Miikkulainen, Q. Le, K. Stanley, and C. Fernando, "Metalearning Symposium @NIPS," in *http://metalearning-symposium.ml*, 2017.

[11] J. Vanschoren, P. Brazdil, C. Giraud-Carrier, and L. Kotthoff, "Meta-Learning and Algorithm Selection Workshop at ECMLPKDD," in *CEUR Workshop Proceedings*, 2015.

[12] J. Vanschoren, J. N. van Rijn, B. Bischl, and L. Torgo, "OpenML: Networked Science in Machine Learning," *SIGKDD Explorations*, vol. 15, no. 2, pp. 49–60, 2013.

[13] M. Ahsan and L. Ngo-Ye, "A Conceptual Model of Recommender System for Algorithm Selection," *AMCIS 2005 Proceedings*, p. 122, 2005.

[14] J. Beel, "A Macro/Micro Recommender System for Recommendation Algorithms [Proposal]," *ResearchGate https://www.researchgate.net/publication/322138236_A_MacroMicro_Recommender_System_for_Recommendation_Algorithms_Proposal*, 2017.

[15] A. Collins, D. Tkaczyk, and J. Beel, "A Novel Approach to Recommendation Algorithm Selection using Meta-Learning," in *Proceedings of the 26th Irish Conference on Artificial Intelligence and Cognitive Science (AICS)*, 2018, vol. 2259, pp. 210–219.

[16] T. Cunha, C. Soares, and A. C. de Carvalho, "Metalearning and Recommender Systems: A literature review and empirical study on the algorithm selection problem for Collaborative Filtering," *Information Sciences*, vol. 423, pp. 128–144, 2018.

[17] T. Cunha, C. Soares, and A. C. de Carvalho, "CF4CF: recommending collaborative filtering algorithms using collaborative filtering," in *Proceedings of the 12th ACM Conference on Recommender Systems (RecSys)*, 2018, pp. 357–361.

[18] T. Cunha, C. Soares, and A. C. de Carvalho, “Selecting Collaborative Filtering algorithms using Metalearning,” in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 2016, pp. 393–409.

[19] P. Matuszyk and M. Spiliopoulou, “Predicting the performance of collaborative filtering algorithms,” in *Proceedings of the 4th International Conference on Web Intelligence, Mining and Semantics (WIMS14)*, 2014, p. 38.

[20] M. Mısır and M. Sebag, “ALORS: An algorithm recommender system,” *Artificial Intelligence*, vol. 244, pp. 291–314, 2017.

[21] M. Vartak, A. Thiagarajan, C. Miranda, J. Bratman, and H. Larochelle, “A Meta-Learning Perspective on Cold-Start Recommendations for Items,” in *Advances in Neural Information Processing Systems*, 2017, pp. 6907–6917.

[22] B. He and I. Ounis, “Inferring query performance using pre-retrieval predictors,” in *International symposium on string processing and information retrieval*, 2004, pp. 43–54.

[23] C. Macdonald, B. He, and I. Ounis, “Predicting query performance in intranet search,” in *SIGIR 2005 Query Prediction Workshop*, 2005.

[24] J. Beel and L. Kotthoff, “Proposal for the 1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval (AMIR),” in *Proceedings of the 41st European Conference on Information Retrieval (ECIR)*, 2019, vol. 11438, pp. 383–388.

[25] S. Arora and A. Yates, “Investigating Ad-Hoc Retrieval Method Selection with Features Inspired by IR Axioms,” in *Proceedings of The 1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval (AMIR)*, 2019.

[26] J. Beel, B. Smyth, and A. Collins, “RARD II: The 94 Million Related-Article Recommendation Dataset,” in *Proceedings of the 1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval (AMIR)*, 2019.

[27] G. Edenhofer, A. Collins, A. Aizawa, and J. Beel, “Augmenting the DonorsChoose.org Corpus for Meta-Learning,” in *Proceedings of The 1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval (AMIR)*, 2019.

[28] M. Mansoury and R. Burke, “Algorithm selection with librec-auto,” in *Proceedings of The 1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval (AMIR)*, 2019.

[29] T. Nagarajah and G. Poravi, “An Extensive Checklist for Building AutoML Systems,” in *Proceedings of The 1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval (AMIR)*, 2019.

[30] M. Lindauer, “Automated Algorithm Selection: Predict which algorithm to use! (Keynote),” in *1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval (AMIR)*, 2019.

[31] L. Kotthoff, “Hands-on Session with ASlib,” in *1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval (AMIR)*, 2019.

[32] M. Lindauer, “Hands-On Automated Machine Learning Tools: Auto-Sklearn and Auto-PyTorch,” in *1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval (AMIR)*, 2019.